UMEÅ UNIVERSITY

Essays on Labor Supply, Pension Policy, and Inequality

Johan Gustafsson

Department of Economics
Umeå School of Business, Economics and Statistics
Umeå 2021

Doctoral thesis
ISBN: 978-91-7855-599-4
ISSN: 0348-1018
Umeå Economic Studies No. 989
Cover:"*The labor–leisure trade-off*". Photos by Björn Sundqvist, 2021. Design by Ida Åberg, Inhousebyrån, Umeå University.
Electronic version available at: http://umu.diva-portal.org/
Printed by: UmU Print Service at Umeå University
Umeå, Sweden 2021

"A mistake is always forgivable, rarely excusable and always unacceptable."

- Robert Fripp

Abstract

Paper [I] analyzes the dynamic properties of life-cycle earnings in Sweden using micro-data. We study the evolution of permanent and transitory earnings inequality over the period 2002–2015. Our data comes from administrative records gathered in the ASTRID database. We find that some features of the data do not match the predictions of the heterogeneous or restricted income profile models commonly applied in the earning dynamics literature. Instead, we estimate an alternative permanent-transitory (PT) error components model. Analyzing the covariance structure of both male and female earnings, controlling for educational background, we find that the upward trend in permanent earnings inequality observed in Sweden during the 1990s does not continue in the 2000s, and the financial crisis of 2008 did not have any major impact on the variability of earnings. We further simulate the accumulation of income pension entitlements and find that variations in pension entitlements are smaller among college-educated workers.

Paper [II] studies the life-cycle effects of favorable marginal tax treatment on older workers' optimal life-cycle labor supply, retirement timing, and savings. I develop a structural model in continuous time where the life cycle of a representative agent is divided into three distinct phases: pre-treatment, post-treatment, and retirement. Solutions for consumption and savings, labor supply and leisure, and retirement timing are then obtained by solving the model as a salvage value problem. I then calibrate the model to Swedish earnings data and find that the increased extensive margin labor supply is partially offset by a reduction in hours worked during the pre-treatment period. The total effect, however, is an increase in life-cycle labor supply and consumption.

Paper [III] studies the implications of the structure of public pension programs for the trade-offs determining economic behavior over the life cycle. The economy is modeled as a continuous-time overlapping generations model with endogenous labor supply, savings, and human capital formation. Individuals differ in ability and are free to choose how much to work at each period in time and when to enter and exit the labor market. Numerical simulations provide qualitative insights that a redistributive pension system introduces opposite effects on the incentives for retirement for high- and low-skilled individuals, which leads to increased earnings inequality. This effect can, in turn, dominate reduced pension inequality such that lifetime and population-wide income inequality increase. Ultimately, the equity–efficiency trade-off is found to be difficult to characterize.

Paper [IV] explores the effects of pension illiteracy on aggregate labor supply and the redistributive performance of public pension systems. I consider an OLG model in continuous time populated with individuals who differ in labor productivity and pension literacy. Agents suffering from pension illiteracy fail to fully account for the structure of the pension system when planning their economic behavior over the life cycle. I find that pension illiteracy can negatively impact aggregate labor supply and increase earnings inequality and lifetime income inequality. This suggests that pension illiteracy may limit the efficiency gains of an increased correlation between individual contributions and benefits, making the equity–efficiency trade-off difficult to characterize in the context of pension reforms.

Keywords: income inequality, labor supply, life cycle, retirement, public pension

Acknowledgments

Let me begin by assuring any concerned reader that my rejection of a career path in accounting was not in any way guided by an intolerance toward 15th-century Italian monks. In fact, I have several friends who happen to be 15th-century Italian monks. It was merely my inability to find solitude in accounting charts that convinced me that their guild is not the fate to which I would like to aspire.

Having said that, I wish to gratefully acknowledge everyone who has contributed either directly or indirectly to the realization of this thesis. The writing process has been far more of a soli than a cadenza. In particular, the first paper would not have seen the light of day without the splendid work done by my colleague and dear friend Johan Holmberg. A friendship rooted in a feverdreamish childhood in Bergsbyn, blossomed in the haze of laissez-faire snaps consumption, and set to wilt only at death.

This thesis has greatly benefited from a confluence of valuable input from my surroundings. I owe special gratitude to my supervisors Gauthier Lanot, Thomas Aronsson, and Katharina Jenderny. You are truly competing with Sonny Rollins's legendary late 1950s lineup in terms of excellent trios. I cannot thank you enough for helping me make sense of my demented drafts. Also, I am very thankful for the opportunity to discuss my essays with Tomas Sjögren, Sonal Yadav, Fei Xu, Magnus Wikström, Olle Westerlund, Raoul Theler, Johan Lundberg, and many others at seminars, presentations, and in the coffee room.

I also wish to disclose that if it had not been for Mattias, Runar, Aemiro and Niklas's support during my master's studies, I most likely would not have pursued a PhD to begin with.

A PhD journey should preferably not be a period of isolation. Unfortunately, the pandemic had other plans. While lamenting this unfortunate fact of life, I rejoice in the many occasions I got to share with my colleagues. A special thanks to Sef, Linn, Hanna, Marco, Tomas R, Balder, Stina, Peter B, Alejandro E, y Alejandro V, who all made my PhD experience a bit more studentesque with LAN-parties, barbecues and bar-crawls.

These 3.6667 years would have been much less memorable without my friends outside of the econ corridor. Thanks to Erik W, Peter G, and Maria of Crayfish Colony for not only being some of my closest friends but also allies in corrupting the essence of harmony. Jesper, Viktor, Erik S, Adrian, Björn, Marta, Yasemin, Pontus—my immense fondness for you is perpetual.

Last but not least, I want to thank my family, who have been ever so patient with my career choice and exact position in the universe.

To express my sincere gratitude (while being aware that this gesture will definitely fall short of the compensation you truly deserve), the next page contains one of my favorite weekday recipes. I hope that it sparks much joy.

Long may this recognition resonate through the lands.

G

L'ombre au court-bouillon (Harrkok)

Ingredients

2 graylings (freshly caught!)

2 dl of a dry white wine (preferably not aged in oak barrels)

1 liter of water

1 yellow onion

1 leek

1 carrot

2 sprigs of parsley

2 sprigs of lovage

2 cloves

8 white peppercorns

Instructions

1. Scale, gut, and clean the fish. Preferably leave the head on, but remember to discard the gills, as they add a bitter taste.
2. Add the water and wine to a large (wide) pot and bring to a boil. Reduce heat to achieve and maintain a gentle simmer.
3. Thinly slice the onion, carrots, and the white part of the leeks. Add to the pot together with the herbs, cloves, white peppercorns, and a generous pinch of salt. Cover and let simmer for 30 minutes.
4. Carefully add the fish. Cover and let simmer (very gently!) until a thermometer reads 53–55 degrees Celsius next to the backbone. After discarding the vegetables and herbs, the bouillon can be reused.
5. Remove the skin of the fish at the table and serve with a slice of lemon, boiled new potatoes tossed with some olive oil, and chopped parsley.
6. ...And as always, enjoy!

This thesis consists of an introduction, a summary, suggestions for future research, and the following four self-contained papers:

Paper [I]
Gustafsson, J., Holmberg, J. (2019). Earnings Dynamics in Sweden: The Recent Evolution of Permanent Inequality and Earnings Volatility. *Umeå Economic Studies*, No. 963. Revise and resubmit.

Paper [II]
Gustafsson, J. (2021). Age-Targeted Income Taxation, Labor Supply, and Retirement. *CESifo working paper*, No. 8988-2021. *Umeå Economic Studies*, No. 985.

Paper [III]
Gustafsson, J. (2021). Public Pension Reform and the Equity–Efficiency Tradeoff. *Umeå Economic Studies*, No. 992.

Paper [IV]
Gustafsson, J. (2021). Implications of Pension Illiteracy for Labor Supply and Redistribution. *Umeå Economic Studies*, No. 993.

1 Introduction

This thesis studies the relationships between life-cycle behavior, public pension policy, and inequality. Paper [I] examines earnings and pension inequality in Sweden through econometric and simulation-based techniques. Papers [II]–[III] model behavioral adjustments in terms of life-cycle labor supply and savings to tax and retirement policies. Finally, paper [IV] studies the implications of pension illiteracy for aggregate labor supply and the redistributive performance of public pensions.

1.1 Public Pension and Retirement Policy

The overarching objectives of public pension systems are to protect individuals against the risk of poverty in old age and facilitate consumption smoothing as individuals transition from working life to retirement. Following the conceptual framework outlined by the World Bank,[1] the different modalities for achieving these objectives can be defined within three distinct pillars. The first pillar is redistributive and essentially functions as a tax-transfer scheme with non-contributory basic benefits. The second pillar works as a compulsory savings scheme in which benefit entitlements bear some positive correlation to individual contributions. The third pillar is made up of voluntary, private savings. As a result, public pension systems provide several means of redistributing income over the life cycle and between individuals. While the multi-pillar structure is ubiquitous, the relative weights given to each of these pillars can vary substantially between countries (Holzmann et al., 2008).[2]

The mandatory pillars of most pension systems today operate through a pay-as-you-go (PAYG) principle. This implies that the pension benefits of today's retirees are financed by contributions from today's workers, who, in return, are promised to receive the same treatment when they retire.

While the PAYG mechanism introduces little risk and is easily administered, the long-run credibility of such contracts is sensitive to demographic changes. In particular, current and forecasted increases in longevity combined with declining fertility rates in developed economies are making it increasingly difficult for governments to secure enough contributions to finance the accrued benefit obligations.[3] Ultimately, a relevant question for policymakers and scholars alike is how policy can be designed to safeguard the fiscal sustainability of public pension systems.

[1] See e.g., World Bank (1994).

[2] On the one hand, the public pension systems in Germany, Italy, and France are essentially earnings-based with little intragenerational redistribution. On the other hand, the systems in New Zealand, Canada, and Denmark are characterized by a substantial basic benefit component.

[3] The historical implementation of PAYG schemes as opposed to capitalized schemes is often referred to as an "original sin" (see e.g., Cesaratto, 2006). In particular, as contributions to a PAYG scheme are realized as concurrent benefits, the existence of such a scheme crowds out capital in the economy (Feldstein, 1974).

A reform option of more limited ambition is to increase labor market participation by modifying the rules for claiming pension benefits and increasing the retirement eligibility age. While straightforward, such reforms seldom enjoy the support of public opinion (Lacomba and Lagos, 2006).

A potentially less controversial option is to induce longer working lives by modifying the tax code such that it becomes beneficial for individuals to delay their retirement (e.g., Biggs, 2012; Laitner and Silverman, 2012; Alpert and Powell, 2013). An example of such a policy was introduced in Sweden in 2007 in the form of more generous earned income tax credits for older workers. By lowering the effective income tax for these workers, the policy reinforces the financial incentives to delay retirement in terms of increasing the opportunity cost of leisure in old age. By using a reduced-form static model, Laun (2017) finds that the policy had positive effects on the labor market participation of older Swedish workers and concludes that this policy constitutes a viable option for promoting a longer working life.

Paper [II] recognizes that such a policy modifies economic incentives over the whole life cycle in terms of the intensity of labor supply and savings behavior. In particular, while the total labor supply increases and workers delay their retirement date in response to the policy change, this effect is partially offset by an intertemporal substitution of labor supply from the pre-treatment working life phase to the post-treatment years. Hence, the evidence based on a reduced-form static model misses important margins of adjustment to the implementation of such a policy.

A more extensive reform is a complete or partial refinancing of compulsory pension systems from PAYG to fully funded (FF) schemes. Under FF schemes, contributions are capitalized through actual investments. Such reforms would balance contributions and benefit obligations by construction, perfectly hedging against demographic changes. Another argument in favor of such a reform is that any contributions will constitute a perfect substitute for private savings (Castanheira and Galasso, 2011). On the one hand, if the private savings rate is higher than the mandatory contribution rate, the total financial investment in the economy will not be affected by the existence of a public pension system. On the other hand, if the private savings rate is lower than the mandatory contribution rate, the scheme is expected to increase the national savings rate, thus promoting economic growth.

However, reforming pension systems from PAYG to FF introduces a new set of potential problems. In particular, it is difficult to make the transition from PAYG to FF without violating the Pareto criterion. Essentially, the transitional generation will have to carry the financial burden of contributing to both the FF and PAYG mechanisms but will only receive benefits from the FF transfer (Cesaratto, 2006). This realization explains the skepticism raised against such a reform.

An alternative reform that emerged in the 1990s through its implementation in countries such as Italy, Sweden, and Poland is the establishment of notional

pension accounts. The notional defined contributions (NDC) scheme seeks to mimic the capitalization of contributions without abandoning the PAYG transfer mechanism. Instead of being invested in actual funds, individual contributions are registered as entitled pension wealth and appreciate at a fictitious interest rate. At retirement, the notional wealth is generally converted to an entitlement to a stream of annuities.

Many proponents of the NDC model stress the similar potential efficiency gains of notional accounts compared to those expected to emerge from funded accounts. By increasing the earnings dependence of pension benefits and periodically informing individuals about the status of their entitled pension wealth, it is more likely that individuals will view contributions as savings rather than income tax (Williamson and Williams, 2003). This reduces the implicit tax wedge on the labor supply and promotes a longer working life.

An advantage of the NDC scheme is that its implementation does not require any generation to bear the transitional cost. A disadvantage is that pension savings remain illiquid and therefore do not directly contribute to the national savings rate and economic growth (Sundén et al., 2000). In addition, the NDC scheme does not ensure solvency under all conceivable demographics scenarios since the transfer mechanism remains PAYG.

Critics of both NDC and FF schemes often highlight the lack of redistribution from high- to low-income earners as a major weakness of the designs. From an egalitarian perspective, this attribute violates the central rationale for social welfare programs. To allow for intragenerational redistribution, many economies have settled for a scheme where a fraction of contributions are allocated toward the NDC account, and the residual fraction is allocated toward a basic pension income scheme. The policymaker seemingly faces a trade-off between the efficiency gains of individualized pension accounts on the one hand, and the redistributiveness of common or basic benefits on the other (Cigno, 2008).

However, this equity–efficiency trade-off may be difficult to characterize. Sommacal (2006) and Hachon (2008) find that a redistributive pension system introduces labor supply effects that work to increase pension inequality. Sommacal concludes that a flat-rate pension system introduces larger disincentives for labor supply among low-income individuals relative to high-income individuals, as they will be the net beneficiaries of such a transfer, while high-income individuals will be the net contributors. As a result, earnings inequality increases. Under the assumption of a perfectly competitive labor market, the increased earnings inequality is found to completely offset the reduced pension inequality such that lifetime inequality remains unchanged. This suggests that a reform from redistributive to individualized pension accounts may increase efficiency without necessarily increasing economic inequality.

Paper [III] builds on these insights but recognizes that the structure of public pension systems also modifies incentives for labor market participation. By adding behavioral adjustments on the extensive margins of labor supply to the analysis, this study draws a remarkable conclusion about redistribution through

the pension system: namely, that a redistributive pension system may increase lifetime inequality in the economy following a combination of increased dispersion in hours worked and retirement age. This argues against the conventional equity–efficiency trade-off and suggests that the relationship between the benefit formula and inequality is more complex than it appears when abstracting from extensive margin labor supply effects.

In paper [IV], I show that the efficiency gains of reforming a pension system from redistributive to earnings-based are sensitive to the assumption that individuals fully understand the benefit formula. If individuals fail to acknowledge the correlation between contributions and benefits, they will view the contribution rate as a pure labor income tax, irrespective of whether the scheme is earnings-based or not, thereby reducing the intensive margin labor–leisure trade-off to that of only instantaneous effects. Reforming the benefit formula of public pension systems is then predicted to only affect labor supply behavior via changes in replacement income: a low (high) income household will face less (more) beneficial retirement income under an earnings-based scheme relative to a pure redistributive system. Since this affects the cost of retiring, it highlights the importance of including the extensive margin labor supply response in the analysis.

The interaction between earnings and realized pension income is also studied in paper [I]. This paper links earnings inequality to pension inequality in Sweden and illustrates that if earnings become more volatile, pension inequality decreases. The proposed explanation for this finding concerns the upper and lower limits of pension contributions. If the variability of earnings shocks increases, individuals whose permanent income is either too high or too low to qualify as pensionable income will experience, on a more frequent basis, shocks that render their income qualifiable as pensionable income. Individuals with low (high) permanent earnings will more frequently experience positive (negative) shocks large enough to qualify their income as pensionable.

1.2 The Permanent Income Hypothesis and Inequality

In their seminal work on the permanent income hypothesis (PIH), Friedman and Kuznets (1945) and Friedman (1957) introduce the idea that income consists of a permanent and transitory component. Permanent income refers to the anticipated or expected income over an extended period of time. In the short run, income is, in turn, subject to transitory variation (i.e., risk) following unexpected events imperfectly hedged against by the individual. These characteristics provide important insights into consumption and savings behavior over the life cycle.

Income inequality can correspondingly be decomposed into its permanent and transitory components (see, e.g., Lillard and Weiss, 1979; Baker and Solon, 2003; Kässi, 2014). Permanent inequality describes the anticipated gap between income groups. Transitory inequality, in turn, represents the part of inequality

realized through income fluctuations. Making a distinction between the two determines the adequacy of policies targeting income inequality. On the one hand, if observed inequality is driven by human capital dispersion or permanent demand shocks following a technological change or outsourcing, rendering certain competencies obsolete, implementing retraining programs could be suitable. On the other hand, if income inequality follows from higher income risk, such as health shocks and more volatile labor markets, social security benefits could be an appropriate option.

Following the introduction of longitudinal analyses of income inequality, an extensive literature examines the relative size of these components using the covariance structure of income. This approach has its origin in the seminal contributions of Lillard and Willis (1976), Hause (1977), and Lillard and Weiss (1979). Observed income at any instant in time is thought of as partly determined by individual, time-invariant, permanent characteristics and partly through serially correlated income shocks. For long enough panels, the researcher should be able to recover the parameters of the permanent component from the higher order autocovariances, as any contributions from transitory shocks to these moments should be negligible. This approach allows the researcher to test hypotheses regarding the level of heterogeneity and stationarity in the income process.

Two essential schools of model specifications have prevailed within this literature: the heterogeneous income process (HIP) model and the restricted income process (RIP) model. The HIP specification assumes that the income process is the sum of a deterministic individual trend and a serially correlated shock component. Such a model is stationary, implying that income shocks have low to moderate persistence. The RIP model assumes that the permanent component satisfies the statistical properties of a random walk, suggesting that income shocks are truly permanent. If higher order autocovariances are small, negative, or both, the researcher generally rejects the inclusion of long-run trends in the income process in favor of an RIP specification.

This literature typically bases inference on the generalized method of moments (GMM). GMM is convenient because it relies on few assumptions regarding the distribution of shocks (Doris et al., 2011). The estimation procedure follows from matching a set of theoretical moment conditions, derived from the statistical model, with empirically observed moments from the data. Within this context, the distance between the variance–covariance moments predicted by the statistical model and the moments retrieved from the autocovariance matrices is minimized. The resulting system of equations constitutes the GMM estimator. Moffitt and Gottschalk (1995) extend the GMM approach by introducing time factor loadings to the methodology. Hence, we can study how the structure of inequality has evolved over time to account for the effects of labor market policies and macro events. The factor analysis approach has since been further extended to include cohort factor loadings as in, e.g., Doris et al. (2010), to account for differences in the structure of inequality between birth cohorts.

In paper [I], we identify the recent (2002–2015) evolution of permanent and

transitory inequality in Sweden following the above methodological approach. Previous research on permanent and transitory earnings inequality by Gustavsson (2007) and Gustavsson (2008) covers the period of 1960–1999. He finds that permanent inequality increased throughout the 1960s and 1970s, while increased transitory inequality throughout the 1980s increased total earnings inequality. The distinct increase in earnings inequality following the recession of 1991–94 was, in turn, driven by an increase in permanent inequality. Little is known about the post-millennial evolution of earnings inequality in Sweden, and from a descriptive perspective, it is important to study how the 2008 financial crisis has affected the structure of inequality. We find that the increased permanent inequality of the 1990s did not continue in the 2000s and that the financial crisis had little effect on the composition structure of inequality.

Previous studies have focused almost exclusively on male earnings inequality. Instead, we analyze both women and men. There are several motivations for such a study. First, women make up half of the labor force. Second, women in Sweden are overrepresented in the public sector. It is therefore compelling to study whether we could observe any gender differences in the impact of the financial crisis, which is mainly argued to have affected the private export sector, on the structure of inequality. We find no apparent effect of the financial crisis on either permanent or transitory earnings inequality among women.

Since the main pillar of the Swedish pension system is earnings-based, we can use the model to simulate pension income entitlements in a straightforward manner. This exercise allows us to study the impact of the properties of earnings inequality on realized pension inequality, thus extending the analysis beyond labor income to lifetime income inequality. A main finding is a higher variability of pension entitlements for low-educated individuals than for those who have completed higher education, as measured by years in upper secondary and tertiary education. One explanation is the upper limit on contributions to the pension system, which becomes a binding constraint for people with income above a certain level. This contributes to compressing the distribution of realized benefits among more educated groups.

Papers [III] and [IV] relate to the PIH by studying how the interaction between behavioral adjustments and public pension policy affects permanent inequality. In paper [IV], permanent inequality is endogenous to individuals' labor supply decisions. In paper [III], the wage differential is also endogenous to the decision on human capital formation.

1.3 The Life-cycle Perspective: Which Margins Matter?

In a stylized fashion, the life cycle can be decomposed into distinct phases based on consumption and labor supply behavior: human capital formation, working life, and retirement. To smooth consumption over the life cycle, individuals tend to borrow when young, accumulate some wealth while working, and dissave when retired. In addition to decisions on the consumption–savings margin,

individuals can typically decide when to enter and exit the labor market following decisions on education and retirement, and they have some freedom to determine how much they work in a given period of time. Even when abstracting from various sources of uncertainty, these decisions follow a complex set of economic trade-offs.

Regarding labor supply behavior over the life cycle, pension policy can be argued to affect both the intensive margin labor–leisure trade-off and the active–retired trade-off. Since benefits are generally financed through earnings-based contributions, the contribution rate to public pensions, net of the present value of any incremental benefits realized from such transfers, constitutes a labor income tax. The implicit income tax wedge will lower the opportunity cost of leisure, and any realized replacement income from public pensions will make retirement cheaper. If the scheme is redistributive, it may allow low (high) income households to accumulate sufficient pension wealth given less (more) work.

Indirect effects are also present: since public pension systems typically modify incentives for retirement, pension policy will also indirectly affect human capital incentives. Consider the decision to enroll in tertiary education as analyzed within the Mincerian framework: an individual enrolling in tertiary education faces a trade-off between earnings when young and higher future wage prospects following enhanced productivity. The financial returns to human capital investments are realized over the working life of the individual, with labor market entry and exit defining the investment horizon. An early retirement makes human capital less financially attractive. Conversely, time spent in tertiary education is expected to make retirement more expensive in terms of foregone earnings. As a result of this complementarity, the extensive margins of labor supply are jointly determined, making the trade-offs affecting labor market participation more intricate than if the entry and exit margins had been analyzed in isolation (Jacobs, 2009).

This insight has important implications for studying the effects of public pension policy. Contributions to the pension system will reduce the instantaneous labor payoff, making education less costly in terms of foregone earnings opportunities and thus promoting a delayed labor market entry. However, since pension income reduces the cost of retiring by taxing earnings and introducing pension benefits, it increases incentives for an earlier retirement. Following Jacobs's complementarity argument, this makes human capital investments less attractive and thus provides fewer incentives to spend time in education. Studying the redistributive properties of public pension schemes will thereby overlook important behavioral effects.

Sommacal (2006) illustrates that when the assumption of an inelastic labor supply is relaxed in favor of an operable intensive margin, the labor supply incentives introduced by a redistributive scheme lead to higher earnings inequality.

In paper [III], I make the argument that the implicit labor income taxation of public pension systems also directly affects the active–retired trade-off. Dif-

ferent structures of public pension schemes will have distinct implications for the within-cohort distribution of replacement income. In particular, the cost of retirement will increase (decrease) for high-earning (low-earning) individuals when pension contributions are allocated toward the redistributive pillar. The increased difference in labor market participation contributes to an increased earnings inequality in the economy. When accounting for these effects through endogenous extensive margins of labor supply in the model framework, I find that a redistributive pension system may, in fact, increase lifetime income inequality. This conclusion suggests that the equity–efficiency trade-off, within the context of public pension design, is not straightforward.

1.4 Modeling Approach

In papers [II]–[IV], I make extensive use of optimal control theory. When modeling the economic behavior of individuals, I decompose the life cycle into the stylized phases described in section 1.3. In papers [II] and [IV], I model working life and retirement, while paper [III] also includes an initial phase of human capital formation. These phases or stages can be thought of as subdomains of the life cycle, each characterized by unique features in either the specification of the objective function or resource constraint(s).

To illustrate the modeling approach in this thesis, consider a model of life-cycle consumption and labor supply in the spirit of Heckman (1974), as follows:

$$\begin{aligned} &\max_{c(t),l(t)} \int_0^T u(c(t),l(t))dt \\ &\textit{subject to:} \\ &\dot{k} \equiv \frac{dk(t)}{dt} = rk(t) + (\bar{l} - l(t))w(t) - c(t), \\ &k(0) = k_0, \\ &k(T) = k_T. \end{aligned} \tag{1}$$

In the above specification, t is time and $t \in [0, T]$ defines the finite control horizon. $c(t)$ denotes the consumption of non-durable goods, $l(t)$ denotes leisure, and $\bar{l}$ denotes the time endowment at each instant. Labor supply is defined as the time residual of leisure, and $w(t)$ denotes the labor payoff function. Any savings are appreciated by a constant, risk-free interest rate r and flow into the individual's asset account $k(t)$. The utility function u increases in each argument and is strictly concave. This problem can be written in terms of a Hamiltonian function:

$$H(t) = u(c(t),l(t)) + \mu(t)[rk(t) + (\bar{l} - l(t))w(t) - c(t)], \tag{2}$$

where $\mu(t)$ is the co-state variable that indicates the marginal utility of wealth.

Following the maximum principle, the first order conditions satisfy:

$$\frac{\partial H(t)}{\partial c(t)} = \frac{\partial H(t)}{\partial l(t)} = 0. \tag{3}$$

This model provides a simple framework for analyzing the intensive margin labor–leisure trade-off. However, this problem specification does not guarantee that the solution for leisure satisfies the time endowment constraint. By abstracting from such a constraint, the solution for the intensity of labor supply and leisure may not be feasible with $l(t) > \bar{l}$ for some time. This becomes an obvious limitation when studying retirement, which constitutes a salient feature of the life-cycle experience. To analyze such a corner solution, the model needs to be slightly modified such that the constraint $l(t) \leq \bar{l}$ is always satisfied.

One approach to explicitly integrating such an inequality constraint into the optimal control problem is by augmenting the Hamiltonian function in Equation (2) with the restriction $l(t) \leq \bar{l}$, which introduces a marginal penalty $\lambda(t)$ to the utility for any infeasible leisure values $l(t) > \bar{l}$ (Gahramanov et al., 2013).

$$H(t) = u(c(t), l(t)) + \mu(t)[rk(t) + (\bar{l} - l(t))w(t) - c(t)] - \lambda(t)[l(t) - \bar{l}]. \tag{4}$$

While this approach explicitly accounts for the time constraint, it is difficult to obtain closed-form solutions for $\lambda(t)$ and thus solve analytically for $l(t)$. More specifically, it can be difficult to algebraically pinpoint if and when the inequality constraint actually binds. It may be the case that an individual enters and exits the labor market sporadically throughout the life cycle, which would imply that the time endowment constraint could be binding on multiple occasions.

To obtain an analytical solution, the researcher can make a conjecture about the structure of the solution for leisure. This is only a viable option when the trade-off between labor supply and leisure is represented by a well-behaved mathematical relationship. Assuming that the wage function, the cost of leisure, is hump-shaped, it is expected that retirement will follow a smooth reduction in hours as skill depreciation dominates returns to experience. This would then imply that $\lambda(t)$ starts to bind at the age of retirement as the result of a predictable path of leisure intensity.

While such a solution is intuitive for a self-financing agent in a friction-free environment, the presence of pension systems or fixed costs of labor supply may modify incentives for labor supply in such a way that optimal retirement follows from a discrete drop in hours worked. Under such scenarios, it becomes more difficult to predict when $\lambda(t)$ starts to bind. One way to circumvent this analytical lacuna is by modeling the timing of retirement as a decision variable in a two-stage optimal control problem. This implies a modification of the optimal

control problem as follows:

$$\max_{c(t),h(t),R} \int_0^R u(c(t), l(t))dt + \int_R^T u(c(t), \bar{l})dt$$

subject to:

$$\dot{k} = \begin{cases} rk(t) + (\bar{l} - l(t))w(t) - c(t) & \text{for } t \in [0, R), \\ rk(t) - c(t) & \text{for } t \in [R, T]; \end{cases} \tag{5}$$

$$k(0) = k_0,$$

$$k(T) = k_T.$$

In this modeling framework, retirement age R is the terminus of the first phase and the start of the second phase, which coincides with the discrete change to the state dynamics $\dot{k}$. Conditional on the discrete change of the state dynamics at time $t = R$, it is possible to specify two separate Hamiltonian functions, as follows:

$$H_1(t) = u(c(t), l(t)) + \mu_1(t)[rk(t) + (\bar{l} - l(t))w(t) - c(t)], \tag{6}$$

$$H_2(t) = u(c(t), \bar{l}) + \mu_2(t)[rk(t) - c(t)]. \tag{7}$$

What remains to be solved is the optimal switching point, which, in this case, coincides with the retirement date. Following Kamien and Schwartz (2012), the optimal switching point of a two-stage optimal control problem satisfies the following conditions:

$$\mu_1(R^*) = \mu_2(R^*), \tag{8}$$

$$H_1(R^*) = H_2(R^*). \tag{9}$$

So far, the model only considers the economic behavior of a self-financing agent. This is convenient from a modeling perspective since it means that controls for each subdomain can be solved independently. However, suppose that the agent participates in an earnings-based pension system, in which individual benefits are based on life-cycle contributions. Under such a specification, the choice of optimal leisure for the first subdomain $t \in [0, R)$ will have a direct effect on the state dynamics of the second subdomain $t \in [R, T]$. Without making any structural assumption about the contribution–benefit formula, let τ correspond to a proportional contribution rate and b to a pension benefit annuity that is partially determined by $l(t)$. To make a distinction between the different time regimes, I introduce time variable s as a continuation of t beyond the age of retirement. Following this modification, the optimal control problem can be

expressed as follows:

$$\begin{aligned}
&\max_{c(t),c(s),h(t),R} \int_0^R u(c(t), l(t))dt + \int_R^T u(c(s), \bar{l})ds \\
&\textit{subject to:} \\
&\dot{k} = \begin{cases} rk(t) + (\bar{l} - l(t))w(t)(1-\tau) - c(t) & \text{for } t \in [0, R) \\ rk(s) + b(l(t), s | t \in [0, R)) - c(s) & \text{for } s \in [R, T] \end{cases} \\
&k(0) = k_0 \\
&k(T) = k_T.
\end{aligned} \tag{10}$$

By introducing the pension system, the Hamiltonian function in Equation (7) needs to be modified as follows:

$$H_2(s) = u(c(s), \bar{l}) + \mu_2(s)[rk(s) + b(l(t), s | t \in [0, R)) - c(s)]. \tag{11}$$

Following the maximum principle, the welfare effect of a marginal change in a control variable has to be zero in optimum. In a scenario with an earnings-based public pension system, the labor–leisure trade-off bears a resemblance to a delayed response problem, as the total effect is partially realized as a contemporaneous effect and partially as a lagged effect. By devoting more time $t \in [0, R)$ to leisure, the agent experiences a contemporaneous effect on utility and net earnings and a future effect on pension income for $s \in [R, T]$. Given the specification of retirement benefits as annuities, the total effect of increasing leisure at any time t becomes:

$$\frac{\partial H_1(t)}{\partial l(t)} + \int_R^T \frac{\partial H_2(s)}{\partial l(t)} ds = 0. \tag{12}$$

Equation (12) characterizes the modified first order condition for optimal leisure in the presence of an earnings-based pension system. This result can be obtained in a more formal way by stating the maximization problem in Equation (10) in terms of a Lagrangean function:

$$\begin{aligned}
\max_{c(t),c(s),h(t),R} L = &\int_0^R u(c(t), l(t))dt + \int_R^T u(c(s), \bar{l})ds \\
+\mu_0 \Bigg\{ &\int_0^R (\bar{l} - l(t))w(t)(1-\tau)e^{-rt}dt + \int_R^T b(l(t), s | t \in [0, R))e^{-rs}ds \\
&- \int_0^T c(t)e^{-rt} \Bigg\}.
\end{aligned} \tag{13}$$

The first order condition characterizing the optimal choice of leisure at any time

becomes:

$$\begin{aligned}\frac{\partial L}{\partial l(t)} = \frac{\partial u(c(t), l(t))}{\partial l(t)} - \mu_0 w(t)(1-\tau)e^{-rt} \\ +\mu_0 \int_R^T \frac{\partial b(l(t), s|t \in [0, R))}{\partial l(t)} e^{-rs} ds = 0.\end{aligned} \tag{14}$$

The first two additive terms on the RHS of Equation (14) correspond to the partial derivative of the Hamiltonian function in Equation (6) with respect to leisure and the terms within the integral to the partial derivative of the Hamiltonian function in Equation (11) with respect to leisure at time $t \in [0, R)$.

In the context of this thesis, I recognize that the continuous-time approach has at least two advantages over discrete-time modeling. First, a high frequency time variable allows the model to capture marginal adjustments on the extensive margins in a straightforward manner. Second, discrete-time OLG models with more than three periods quickly become analytically intractable. Ultimately, the modeling contributions of this thesis provide a mathematically convenient approach to simultaneously studying the marginal effects of changes in tax and retirement policy on incentives for both intensive and extensive margins of labor supply.

2 Summary of the Papers

Paper [I]: Earnings Dynamics in Sweden: The Recent Evolution of Permanent Inequality and Earnings Volatility

This paper studies the evolution of permanent and transitory earnings inequality in Sweden for the years 2002–2015. Consistent with the PIH, we model earnings in terms of their permanent and transitory components. This allows us to elucidate how the structure of inequality has evolved over the course of the financial crisis of 2008–2009 and the lingering Euro debt crisis. Novel for studies based on Swedish data, we examine both male and female earnings inequality. Lastly, we use the earnings model to simulate pension entitlements based on the income pension pillar and thus link earnings inequality to lifetime inequality.

We draw inferences from the variance–covariance structure of earnings observed in Swedish tax data obtained from the ASTRID database. The theoretical moments predicted by the model are matched with the empirical autocovariance moments using minimum distance estimation.

Contesting the findings of previous studies on Swedish data, we find that the observed earnings variance over the life cycle is not consistent with the predictions of either of the established model specifications within the RIP or HIP dichotomy.

More specifically, we observe non-increasing variances in earnings over the life cycle, while both the HIP and RIP models predict increasing variances. We instead settle for a parsimonious model that is serially correlated by an ARMA(1,1) process, similar to the specification in Lillard and Willis (1978). The persistence of earnings shocks is found to be of small to moderate magnitude, with the AR parameter within the interval of [0.335, 0.411] for men and [0.344, 0.375] for women.

Studying the earnings of both men and women, we conclude that the financial crisis mainly affected the earnings variability of low-educated men by making their earnings more volatile. No significant effect on the volatility of female earnings is found. This result is consistent with the argument that the financial crisis mainly affected private firms in the export sector, which are typically dominated by male workers.

Simulations show that the variability of income pension entitlements is larger among individuals with less than a high school education than those with higher educational attainment. This observation can likely be attributed to the fact that contributions to the income pension system are subject to an upper limit, which truncates contribution amounts of higher income groups.

Paper [II]: Age-Targeted Income Taxation, Labor Supply, and Retirement

This paper models and studies the effect on life-cycle labor supply and savings when introducing a favorable marginal tax treatment in the form of earned income tax credits for older workers. Such a policy was introduced in Sweden in 2007 as part of a general tax reform aimed at increasing overall labor market participation in the economy. By increasing the net returns to labor supply, retirement becomes more expensive in terms of foregone labor income, which is expected to promote a delayed labor market exit.

The model focuses on the life cycle of a representative agent. The agent derives utility from the consumption of non-durable goods and leisure throughout their lifetime. To allow for the possibility of returning to the labor market post-treatment, retirement is only treated as an absorbing state if withdrawal takes place after receiving the treatment.

I consider a continuous-time model framework, allowing me to model retirement behavior in a tractable fashion. The life cycle is modeled as a three-stage optimal control problem following an augmentation of the model in Heckman (1974). The first stage begins at labor market entry and ends at tax treatment. The second stage is initiated at the age of tax treatment and ends at retirement. Lastly, the third stage corresponds to the retirement phase. The problem is solved as a generalized salvage value problem.

The utility function is specified as a Cobb–Douglas function to allow for complementarity between the consumption of non-durable goods and leisure. This assumption allows the model output to reproduce a co-movement between consumption and earnings, as commonly observed from the data.

I calibrate the model to Swedish earnings and mortality data and then simulate the effects of the tax treatment. Two main scenarios are considered: one with a self-financing agent and one in which the agent enrolls in a defined contribution public pension system. The main simulations concern the tax treatment at ages 60 or 65. At these ages, the depreciation of marginal productivity will have reached different values, which, in turn, will have quantitative implications.

The results of the numerical simulations are intuitive: the agent realizes that the tax change implies that leisure becomes relatively more expensive post-treatment and will therefore allocate more time to leisure during the primary working life. The intertemporal substitution effect on savings behavior is evident when considering a more impatient agent since it puts more utility weight on leisure activities and consumption at earlier stages of the economic life. Tax treatment then coincides with a period far into the future to which the agent assigns less weight in terms of disutility from work compared with a more patient agent. Therefore, the agent accumulates a larger debt and compensates by working relatively more when older.

Paper [III]: Public Pension Reform and the Equity-Efficiency Trade-off

This paper studies how the structure of the contribution–benefit formula of public pension systems affect human capital formation, labor supply, and savings behavior over the life cycle. It builds on two main insights from previous literature. The first follows from Sommacal (2006), who shows that labor supply effects are important when studying the redistributive performance of social security. In particular, the reallocation of pension contributions toward a flat benefit component introduces larger labor supply disincentives for low-income households relative to high-income households. The increased dispersion in hours worked increases earnings inequality and may offset the reduced pension inequality such that lifetime inequality is unchanged. The second follows from Jacobs (2009) who shows that labor distortions induced by labor income taxation are predicted to be substantially larger in a model with endogenous extensive margins of labor supply compared to models abstracting from such margins.

To integrate the extensive margin decisions without completely abandoning analytical tractability, I consider a continuous-time model that follows the baseline framework in Jacobs (2009), to which I add heterogeneous agents with regards to earnings ability and an endogenous pension system with a redistributive and notional defined contribution earnings-based pillar.

The contribution rate to the public pension system is levied on earnings, and benefit eligibility is conditional on full-time retirement. The lifetime of each individual is decomposed into three distinct phases: tertiary education, working life, and retirement.

The paper also contributes by introducing a way to solve for an earnings-based pension system in an optimal control framework as a delayed response problem. The effect of a small change in the intensive margin labor supply at each instant is partially recognized contemporaneously through the instantaneous labor–leisure trade-off and partially in a lagged fashion as it affects realized pension benefits.

The main findings are as follows. First, a redistributive pension system may increase lifetime inequality following a larger dispersion in labor market participation. This result is not consistent with the idea of an equity-efficiency trade-off. If individuals are impatient enough, and the public pension system is sufficiently extensive, a reform from a flat-rate type pension system toward an earnings-based scheme may, in fact, promote both efficiency and equity. Second, if the NDC scheme is return-dominated by private savings, the labor supply becomes relatively more attractive later in life compared to earlier in life, as the foregone compound interest is smaller for any pension contributions made when the individual is older. The implicit tax treatment of the contribution rate on the returns to labor supply is thus larger for younger workers than for older workers. Third, such tax treatment lowers the opportunity cost of tertiary education and promotes a delayed labor market entry.

Paper [IV]: Implications of Pension Illiteracy for Labor Supply and Redistribution

Extensive evidence suggests that many individuals fail to understand how their economic behavior interacts with important features of the pension system. The Swedish Pensions Agency claims that a common misconception among individuals enrolled in the public pension system is that they cannot affect their future retirement benefits. This suggests that many individuals do not comprehend the benefit formula.

This study explores the effects on aggregate labor supply and the performance of social security in terms of redistribution when some agents fail to rationalize the financial incentives implicit in the benefit formula of the pension system.

Reforming public pension systems from flat-rate to earnings-based is expected to increase labor market efficiency. This follows an argument of lower implicit income taxation: when contributions are realized as future benefits, the public pension system introduces fewer disincentives for working. However, this result is conditional on individuals taking this correlation into account when planning their labor supply behavior on both the intensive and extensive margins.

To explore the implications of pension illiteracy, I introduce two types of myopic individuals and a canonical rational individual (lifecycler). Type 1 myopes acknowledge that retirement timing affects the annuitization of benefits but fail to recognize how contributions and benefits correlate. Type 2 myopes treat retirement annuities as strictly exogenous. One can think of these individuals as receiving a statement of pension benefits at different moments in time from the pension agency.

The analysis is based on a continuous-time OLG model. In addition to heterogeneity in optimization sophistication, individuals differ in terms of productivity to provide a rationale for redistribution. Individuals enroll in a public pension system with a redistributive (Beveridgean) and earnings-based (Bismarckian) pillar, and the relative weights given to the pillars determine the effective taxation induced by the public pension system.

The study presents three main conclusions: First, pension illiteracy can reduce aggregate labor supply. Second, pension illiteracy can increase both earnings inequality and lifetime inequality. Third, because of pension illiteracy, a Beveridgean, flat benefit scheme can yield higher lifetime income inequality relative to a Bismarckian, earnings-based scheme. These results suggest that if pension illiteracy is common, the equity–efficiency trade-off between implementing a redistributive or earnings-based pension system may be equivocal. This is potentially distressing for policymakers who aim to either increase labor market participation by increasing the correlation between earnings and benefits or reduce economic inequality by implementing a redistributive pension system.

3 Future Research

Understanding how individual behavior interacts with tax and retirement policies, and, in turn, the structure of inequality over the life cycle is fundamental for making thorough assessments of the effectiveness of public pension policies. The findings of the papers in this thesis suggest many avenues for future research.

Since the findings in the first paper reject the HIP model, it is not obvious how differences in human capital formation contribute to permanent earnings inequality in Sweden. One suggestion for future studies is to examine this relationship more explicitly with a model that can incorporate changes in centralized wage bargaining and other institutional policies, which are expected to limit the wage premium of education.

There are at least two general directions for extending the analyses in papers [II]–[IV]. First, accounting for general equilibrium effects through endogenous factor prices would make it possible to analyze the effects of changes in factor prices on life-cycle behavior. For example, because the introduction of a public pension system is predicted to crowd out private savings and raise interest rates, it is also expected to lower the present-value lifetime returns of educational attainment. The current model environment does not account for such feedback effects and could thus potentially risk overlooking an important long-run effect of public pension reform. Second, the models could arguably be extended beyond the canonical framework by incorporating elements of risk. In the dimension of savings, the addition of portfolio choice and risky investments could provide a more realistic framework for analyzing the overall performance of a fully funded scheme relative to unfunded alternatives. It is also reasonable to consider the effects of labor market imperfections. In particular, it can be difficult to justify the assumption that European labor markets are perfectly competitive and friction free.

In terms of additional retirement policies, it would be intriguing to compare the carrot policy of favorable earned income tax treatment as described in paper [II] to a stick policy, such as pension penalties for early retirement, and research whether the two will have different effects on economic aggregates and the government budget via general equilibrium effects. Conclusions from such a study can guide policymakers who aim to promote the fiscal sustainability of public pension systems in the long run.

Regarding paper [IV], a natural extension would be to study how pension illiteracy may interact with other behavioral failures, such as time inconsistency. The model framework could also be used to conduct cost–benefit analyses of launching pension-related information campaigns. The pension agency may increase the efficiency of public pension systems by making information more accessible and less costly for the public.

A major development would be to integrate horizontal constraints (such as credit rations) into the optimal control framework in an analytically tractable

way. This would provide more realism, as borrowing against illiquid pension entitlements is rarely allowed. This constrains individuals to borrow too much when they are young, which could have substantial implications for impatient individuals. While credit rations have been integrated into optimal control theory, as exemplified by Park (2006), closed-form solutions are only obtained when the structure of the solution is analytically predictable, such as in a canonical life-cycle consumption model with inelastic labor supply.

References

Alpert, A. and Powell, D. (2013). Estimating intensive and extensive tax responsiveness: Do older workers respond to income taxes?

Baker, M. and Solon, G. (2003). Earnings dynamics and inequality among canadian men, 1976–1992: Evidence from longitudinal income tax records. *Journal of Labor Economics*, 21(2):289–321.

Biggs, A. (2012). A payroll tax cut could help social security. *Wall Street Journal*, 25.

Castanheira, M. and Galasso, V. (2011). Which reforms for a fair and sustainable pension system? *Reflets et perspectives de la vie économique*, 50(3):187–202.

Cesaratto, S. (2006). Transition to fully funded pension schemes: a non-orthodox criticism. *Cambridge Journal of Economics*, 30(1):33–48.

Cigno, A. (2008). Is there a social security tax wedge? *Labour Economics*, 15(1):68–77.

Doris, A., O'Neill, D., and Sweetman, O. (2010). Aggregate earnings inequality in europe: Permanent differences or transitory fluctuations?

Doris, A., O'Neill, D., and Sweetman, O. (2011). Gmm estimation of the covariance structure of longitudinal data on earnings. *The Stata Journal*, 11(3):439–459.

Feldstein, M. (1974). Social security, induced retirement, and aggregate capital accumulation. *Journal of political economy*, 82(5):905–926.

Friedman, M. (1957). The permanent income hypothesis. In *A theory of the consumption function*, pages 20–37. Princeton University Press.

Friedman, M. and Kuznets, S. (1945). Income from independent professional. *New York*.

Gahramanov, E., Tang, X., et al. (2013). Solving for the retirement age in a continuous' time model with endogenous labor supply. Technical report.

Gustavsson, M. (2007). The 1990s rise in swedish earnings inequality–persistent or transitory? *Applied Economics*, 39(1):25–30.

Gustavsson, M. (2008). A new picture of swedish earnings inequality: Persistent and transitory components, 1960–1990. *Review of Income and Wealth*, 54(3):324–349.

Hachon, C. (2008). Do redistributive pension systems increase inequalities and welfare?

Hause, J. C. (1977). The covariance structure of earnings and the on-the-job training hypothesis. In *Annals of Economic and Social Measurement, Volume 6, number 4*, pages 335–365. NBER.

Heckman, J. (1974). Life cycle consumption and labor supply: An explanation of the relationship between income and consumption over the life cycle. *The American Economic Review*, pages 188–194.

Holzmann, R., Hinz, R. P., and Dorfman, M. (2008). Pension systems and reform conceptual framework. *World Bank Discussion Paper*, 824.

Jacobs, B. (2009). Is prescott right? welfare state policies and the incentives to work, learn, and retire. *International Tax and Public Finance*, 16(2):253–280.

Kamien, M. I. and Schwartz, N. L. (2012). *Dynamic optimization: the calculus of variations and optimal control in economics and management*. Courier Corporation.

Kässi, O. (2014). Earnings dynamics of men and women in finland: permanent inequality versus earnings instability. *Empirical Economics*, 46(2):451–477.

Lacomba, J. A. and Lagos, F. (2006). Population aging and legal retirement age. *Journal of Population Economics*, 19(3):507–519.

Laitner, J. and Silverman, D. (2012). Consumption, retirement and social security: Evaluating the efficiency of reform that encourages longer careers. *Journal of public economics*, 96(7-8):615–634.

Laun, L. (2017). The effect of age-targeted tax credits on labor force participation of older workers. *Journal of Public Economics*, 152:102–118.

Lillard, L. A. and Weiss, Y. (1979). Components of variation in panel earnings data: American scientists 1960-70. *Econometrica: Journal of the Econometric Society*, pages 437–454.

Lillard, L. A. and Willis, R. J. (1976). Dynamic aspects of earnings mobility. Technical report, National Bureau of Economic Research.

Moffitt, R. A. and Gottschalk, P. (1995). *Trends in the Autocovariance Structure of Earnings in the US, 1969-1987*. Johns Hopkins University, Department of Economics.

Park, M.-H. (2006). An analytical solution to the inverse consumption function with liquidity constraints. *Economics Letters*, 92(3):389–394.

Sommacal, A. (2006). Pension systems and intragenenerational redistribution when labor supply is endogenous. *Oxford Economic Papers*, 58(3):379–406.

Sundén, A. et al. (2000). How will sweden's new pension system work? Technical report.

Williamson, J. B. and Williams, M. (2003). The notional defined contribution model: An assessment of the strengths and limitations of a new approach to the provision of old age security. *Boston College Center for Retirement Research Working Paper*, (2003-18).

World Bank (1994). *Averting the old age crisis: Policies to protect the old and promote growth. Summary*. The World Bank.