Spousal Labor Response to Primary Income: Identification and Heterogeneity*

Yongsung Chang
Bank of Korea & Seoul National U

Elin Halvorsen
Statistics Norway

Marios Karabarbounis
FRB Richmond

March 20, 2025

Based on administrative data from Statistics Norway, we provide a new estimate of the elasticity of spousal labor supply in response to changes in the household's primary worker's income: the so-called "added worker effect." Our estimation uses the sales revenue of the primary worker's employer as an instrument to isolate primary income shifts that are orthogonal to spouse's productivity and persistent over time, both key properties to mitigate the attenuation bias typically present in empirical studies. We find an economically meaningful and statistically significant role for the spousal labor supply, especially among young, hand-to-mouth households. We show that a heterogeneous agent model with dual labor supply calibrated to the Norwegian economy successfully reproduces our empirical estimates. Using our model, we evaluate the welfare implications of various government subsidies designed to mitigate the effects of unexpected income shocks.

JEL Classification: E6, H2, J2
Keywords: Added Worker Effect, Income Risk, Intra-household Insurance

*Emails: yongsung.chang@gmail.com, marios.karabarbounis@rich.frb.org, elin.halvorsen@ssb.no. For helpful comments we thank Serdar Birinci, Richard Blundell, Eric French, Dirk Krueger, Makoto Nakajima, Ananth Seshadri, Gustavo Ventura, and seminar participants at Philadelphia Fed, SED 2023, Norges Bank, IAER, KER 2024, ESAM 2024, and University of New South Wales. Any opinions expressed are those of the authors and do not necessarily reflect those of the Federal Reserve Bank of Richmond, the Federal Reserve System, Statistics Norway, or the Bank of Korea. This research was supported by a grant from the National Research Foundation of Korea, funded by the Korean government (NRF-2023S1A5A2A03083296), and the Norwegian Research Council (Grant 326419).

1 Introduction

Do spouses participate in the labor market when the primary worker in the household loses income? Although it seems natural that families use labor supply to insure against income fluctuations, extensive research has found little evidence in favor of the so-called "added worker effect" (e.g., Lundberg, 1985; Doepke and Tertilt, 2016).

The literature has proposed several explanations for the lack of evidence regarding the insurance role of spousal labor supply. First, estimates of spousal labor supply response may suffer from attenuation bias for at least two reasons: (i) A significant portion of earnings variation reflects predictable or transitory fluctuations that have little impact on household consumption, and consequently, spousal participation (Guvenen, 2007; Blundell, Pistaferri, and Preston, 2008), and (ii) primary worker's income (or productivity) is often correlated with the productivity of the spouse, as is for instance, during a labor market recession. Second, the spousal labor supply response may be small if households have other means to insure such as financial assets, government transfers, or help from relatives (Cullen and Gruber, 2000; Birinci, 2024). Third, there are constraints (for example, childcare or social norms) in adjusting the labor supply of the secondary earner (Halla, Schmieder, and Weber, 2020).

In this paper, based on detailed administrative data from Statistics Norway, we develop an estimation framework to overcome the identification challenges and study the heterogeneity in spousal labor supply responses following a change in primary worker's income.

The Norwegian data provides a unique background to better identify the added worker effect (AWE). First, the data includes a complete description of households' labor income, labor market status, and, more importantly for our analysis, employer information. Based on the matched employer-employee registry we use information on the primary worker's employer (such as sales revenue) as an instrument variable (IV) to isolate primary income shifts that are orthogonal to spouse's productivity and persistent over time, both key properties to reduce the attenuation bias.[1] Second, since households in Norway are obliged to report their wealth to the tax authority, our data includes reliable information on households' assets. In combination with detailed demographic characteristics, we can identify and analyze the response of households with limited means of insurance, or with limited ability to adjust the spousal labor supply, e.g., hand-to-mouth households and young households with children.

We find an economically meaningful and statistically significant role for spousal labor supply, especially among young, hand-to-mouth households. According to our estimation based on the instrument, in response to a 10% decrease in primary income, the participation rate

[1]Similar estimation strategies have been employed to identify the role of precautionary savings and the allocation of financial assets between safe and risky (e.g., Fagereng, Guiso, and Pistaferri, 2018; Chang, Hong, Karabarbounis, Wang, and Zhang, 2022).

of the spouses increases by 1.3 percentage points and earnings by 3.9 percent. Among hand-to-mouth households, spousal participation increases by 2 percentage points and earnings by 16 percent. We also document substantial heterogeneity in responses by age: when primary income decreases by 10 percent, households between ages 30 and 40 increase spousal participation by around 3 percentage points while for households between ages 50 and 60, spousal participation remains about the same. These results rely crucially on our instrumental variable. The ordinary least squares (OLS) estimate does not show a significant response of the spousal labor supply to changes in primary income, consistent with previous findings in the literature.

Next, we construct a quantitative structural model of the household with two goals in mind. First, we test whether a relatively standard model of joint labor supply can replicate the estimated response of spousal labor supply that we find from the data. Second, we use the model to evaluate various forms of government transfer policy against unexpected income fluctuations.

The model is a life cycle model with heterogeneous dual-earner households (e.g., Krueger and Wu, 2021; Guner, Kaygusuz, and Ventura, 2023). Each household consists of a primary worker and a spouse. Households decide every period how much to consume, save, and the labor supply of each member. Labor supply decisions occur along both the participation margin and the amount of hours of work given labor market participation.

We incorporate a rich structure in the income process. First, households have a fixed type (skill) along the life cycle. Second, each earner accumulates human capital while working (Blundell, Costa Dias, Meghir, and Shaw, 2016). Third, earners experience two types of idiosyncratic income shocks, a persistent and a transitory shock. The persistent idiosyncratic shocks are correlated between the two household earners while the transitory shocks are independently distributed across time and earners.

Households self-insure against income fluctuations using labor supply and the household's assets. In addition, a government is providing public insurance in two forms: there is a tax system handing out transfers for low- income households and there is a social security system for retirees. Every period the government sets the tax rate in order to balance its budget.

The model is calibrated to match some salient aspects of the Norwegian economy, such as life cycle profiles of participation, income shares, cross-sectional dispersion of income, dynamics of earnings, transition between employment and non-employment, and the social welfare system. Although not targeted, the model generates a response of spousal labor supply similar to the response in the data. The model can also replicate the discrepancy between the OLS and IV estimates in the data.

Using our quantitative model, we analyze the effectiveness of government subsidies when

households experience an unexpected and temporary income shock. Specifically, we consider four government subsidies that differ by whether the recipients are targeted or not and whether the amount of the subsidy is conditional on employment status or not. All programs require the same amount of additional government expenditure financed by the proportional income tax rate.

We find that the universal and unconditional subsidy is the least effective policy. To fund the policy, the tax rate increases substantially. Moreover, the availability of public insurance crowds out private insurance using the spousal labor supply. Thus, average household consumption and labor supply decline even more than when there is no intervention at all.

The government can improve efficiency in two ways. First, by targeting the subsidy to the most impacted households the government does not need to raise much the proportional tax. Second, by handing out larger transfers to households with dual earners the government mitigates the consumption loss while preserving the incentives of spouses to stay employed. Overall, the policy which targets households that experience a productivity decline *and* is conditional on staying employed in the labor market maximizes both employment and consumption.

We also evaluate these policies based on the welfare of newborn workers. Targeted policies far outperform universal policies not only in terms of efficiency but also welfare. In contrast, conditional policies rely on providing incentives for households to stay in the labor market. Thus, they can improve efficiency but not always welfare.

Our paper contributes to the empirical literature on the AWE. Most empirical studies have lacked the comprehensive data to overcome the identification challenges in estimating the AWE, and thus, find a small or weak spousal labor supply response. For example, based on survey data, Cullen and Gruber (2000) find little evidence of a response in the wife's labor supply to the unemployment of the husband due to the unemployment benefits assistance. Birinci (2024) finds no statistically significant change in spousal earnings following the loss of the husband's job. Casella (2022) emphasizes the role of past labor market experience and finds stronger participation of spouses with work experience during the last five years.

Based on administrative data from Austria, Halla, Schmieder, and Weber (2020) unravel the role of different channels that affect the response of the spousal labor supply using plant closure events of the primary worker. Plant closures generally occur gradually and may well have been foreseen by workers, which may explain why the authors find small spousal labor supply responses (see for example, Stephens, 2002). In addition, the Austrian data do not have linked information on household assets so the authors cannot study heterogeneity in responses by wealth.

Our analysis relies on an IV (sales revenue of the primary worker's employer) that has

proven successful in identifying the income risk that is unanticipated and orthogonal to other household variables (Fagereng, Guiso, and Pistaferri, 2018). In addition, we use detailed information on financial assets and household demographics, which helps us to identify groups that are vulnerable to income risk such as young and hand-to-mouth households – consistent with Bacher, Grübener, and Nord (2024) who find larger spousal responses for young households.

Also consistent with our empirical findings, Autor, Kostol, Mogstad, and Setzler (2019) find a strong negative response of the spousal labor supply to disability insurance receipts in Norway. Albanesi and Prados (2022) find that spouses of husbands with high-income growth during the 1980s and 1990s decreased their participation in the labor market, due to an income effect. Ellieroth (2023) argues that employed spouses may remain employed longer during business cycle downturns due to higher risk of unemployment for the primary worker. Guner, Kulikova, and Valladares-Esteban (2024) show that female participation in the labor market is relatively stable over the business cycle as some women lose their jobs while others join the labor market to compensate for the husband's loss of job.

Our paper also contributes to the development of structural models of consumption and joint labor supply. Blundell, Pistaferri, and Saporta-Eksten (2016) employ a semi-structural model to estimate the degree of consumption insurance and the labor response of working spouses. Their model does not have a participation decision, whereas ours feature participation and thus is more suitable for estimating the AWE. Krueger and Wu (2021) use a model that includes a participation decision for the spouse to evaluate the empirical approach of Blundell, Pistaferri, and Saporta-Eksten (2016) but do not separately analyze the model's predictions for the AWE. Our paper estimates both extensive and intensive margin responses from the data and compares these to a model that includes both margins.

Finally, our paper contributes to the literature that evaluates government policy when household can potentially employ both earners. Our paper follows the categorization of transfers by Guner, Kaygusuz, and Ventura (2020) as means-tested or conditional on household characteristics. Their objective is to analyze permanent changes in the welfare system and more specifically, child transfers. In contrast, our policy objective is to stabilize the economy in the face of income shocks (e.g., McKay and Reis, 2016). Our setup is closer to policies that have been recently implemented such as direct government transfers during the pandemic.

The paper is organized as follows. Section 2 derives an analytic formula for the elasticity of the spousal labor supply to primary income in a static environment. Section 3 provides the empirical analysis. Section 4 sets up the quantitative model and calibrates it to match the salient features of the Norwegian data. Section 5 analyzes various policy experiments based on the quantitative model. Section 6 is the conclusion.

2 The Spousal Labor Supply in a Static Model

We analytically derive the elasticity of the spousal labor supply to primary income in a simple static model. This illustrates the key factors determining the response of the spousal labor supply to primary income shocks. We develop a fully specified quantitative model in Section 4.

Consider a household making decisions about the consumption, c, the labor supply of the primary earner and his/her spouse, h_p and h_s. We assume a unitary model of household labor supply where consumption is fully shared among household members.

Household preferences are given by:

$$U = \frac{c^{1-\sigma}}{1-\sigma} - \frac{h_p^{1+\frac{1}{\gamma_p}}}{1+\frac{1}{\gamma_p}} - \frac{h_s^{1+\frac{1}{\gamma_s}}}{1+\frac{1}{\gamma_s}} \tag{1}$$

where σ is the coefficient of relative risk aversion and γ is the Frisch elasticity of labor supply. We allow for different Frisch elasticity for the primary worker and the spouse. Household's total income consists of primary worker's earnings (wx_ph_p), spouse's earnings (wx_sh_s), and non-labor income (n) where w is the wage rate for the efficiency units of labor and x_p and x_s is productivity of the primary worker and the spouse, respectively. The household's budget constraint is:

$$c = wx_ph_p + wx_sh_s + n.$$

We assume that primary and spouse productivity are correlated through the function $x_s = f(x_ph_p)$.[2] The first order conditions for labor supply are:

$$\frac{1}{\gamma_i}\ln h_i = -\sigma \ln\left[wx_ph_p + wx_sh_s + n\right] + \ln x_i, \text{ with } i = \{p, s\}.$$

Total differentiation of the first-order-condition for the spousal labor supply $i = s$ with respect to (h_s, x_s) gives the Marshallian labor supply elasticity for the spouse:

$$\frac{dh_s}{dx_s}\frac{x_s}{h_s} = \frac{1-\sigma S_s}{1/\gamma_s + \sigma S_s}$$

where $S_s = \frac{wx_sh_s}{wx_ph_p+wx_sh_s+n}$ denotes the spousal earnings share in household's income (Keane, 2011). The primary earnings share S_p can be defined similarly. The Marshallian elasticity is determined by the relative size of the substitution effects and the income effects. With a zero income effect ($\sigma = 0$), the Marshallian elasticity is identical to the Frisch elasticity γ_s (which

[2]We assume that x_s is a function of x_ph_p (instead of x_p) because it is convenient to derive an elasticity with respect to primary income.

captures the substitution effect). The strength of the income effect depends on the relative risk aversion and the income shares of the spouse. When σ is larger (the marginal utility of consumption diminishes rapidly as consumption increases) or when the income share (S_s) is larger, the income effect is also larger: that is, we have a smaller Marshallian elasticity of the labor supply.

Next, we derive the cross-elasticity of spousal labor supply with respect to primary labor income. A total differentiation of the first-order condition for spousal labor supply with respect to $(h_s, x_p h_p)$ gives the following:

$$\frac{dh_s}{d[wx_p h_p]}\frac{[wx_p h_p]}{h_s} = \underbrace{\frac{-\sigma S_p}{\frac{1}{\gamma_s}+\sigma S_s}}_{\text{income effect}} + \underbrace{\frac{dx_s}{d[wx_p h_p]}\frac{wx_p h_p}{x_s}}_{\text{correlation elasticity}} \times \underbrace{\frac{1-\sigma S_s}{\frac{1}{\gamma_s}+\sigma S_s}}_{\text{Marshallian elast.}}$$

The formula highlights why estimates of the spousal response are typically small. Suppose that there is zero correlation between x_s and primary income, the second term disappears. The first term (income effect) is negative: a decline in primary income induces an increase in spousal labor supply. The size of this term depends on σ, S_p, S_s, and γ_s: combination of income effects and Frisch elasticity. With a positive correlation between spousal productivity and primary income, the second term is positive. Its size depends on the strength of correlation and the Marshallian elasticity of labor supply. Since the second term is positive, it weakens the response of the spousal labor supply.

In our view, this issue characterizes the majority of empirical studies that analyze AWE and has not been adequately addressed. In our empirical analysis of Section 3 we provide an IV empirical design to isolate primary labor income shifts orthogonal to spouse's productivity. This allows us to estimate the income effect on the spousal labor supply that is unambiguously negative.

When the correlation is zero the cross-elasticity of spousal labor supply may be small or large, depending on the strength of the income effect. The income effect depends on (i) the primary worker's income share, (ii) the spousal income share, and (iii) the curvature of the utility function. For example, a household with a small non-labor income share (e.g., little financial assets) and a highly productive primary worker is likely to exhibit a large cross-elasticity for the spouse. Therefore, we expect a large heterogeneity in the cross-elasticity of spousal labor supply, which we also examine in our empirical analysis below.

3 Empirical Analysis

3.1 Data

The Norwegian Registry is a comprehensive set of data with detailed information on employment, labor income, and household financial assets. Our empirical analysis is mainly based on five data sets: (1) Income and Wealth Registry Data provides detailed information on income, transfers, and assets; (2) Central Population Register contains individual demographic information (e.g., gender, date of birth, marital status, number of children); (3) National Educational Database has the history and the latest education record for each resident; (4) Employer-Employee Register provides annual information on workers' labor market status (full- and part-time employment, employer ID, beginning and ending dates of job, total payments from each employer, industry, occupation, etc.); (5) Firm Accounts that report companies' profits and balance sheets, and can be linked to workers through the employer ID number.

The Income and Wealth Registry is based on tax records where most information is third-party reported by employers, banks, brokers, insurance companies, and any other financial intermediaries. In addition, there is information from social security registers. Therefore, earned income includes cash salary, capital income, entrepreneurial income, unemployment benefits, taxable benefits, sickness/maternity benefits, and pensions. Since households in Norway are subject to a wealth tax, the data also include detailed information on households' assets.[3]

Empirical studies of spousal labor supply have several shortcomings. Traditionally, the AWE is estimated during business cycle downturns (or expansions) when labor market opportunities are equally depressed (or ample) for both spouses. Second, many studies rely on survey data with limited amount of observations, and hence, little ability to study the heterogeneity in AWE (Casella, 2022; Birinci, 2024). Our data overcome these challenges. First, as we discuss below, with the linked employee-employer registry, we can isolate fluctuations to primary income that are orthogonal to the spousal productivity. In addition, the data are sufficiently rich to analyze the AWE across multiple dimensions: age, wealth, education status, family size, etc.

We start with all individuals in Norway during 2015-2019.[4] We focus on the working stage of the life cycle and thus keep individuals aged 25-60. At the outset, we have 20.6

[3]Traded financial securities are reported at market value. The value of shares in private companies is reported by individuals as well as private companies to the tax authority. The tax authority will then combine the information from the companies' reports with that from individual households and adjust if necessary.

[4]The Employer-Employee Register underwent major changes in 2015 that significantly improved the data quality. Starting in 2015 it has full coverage of all workers and monthly reporting.

million observations, with approximately 2.6 million individuals per year. We further refine the sample using the following criteria. First, we keep married households with spouses also aged 25-60. Second, we drop households where the primary worker is self-employed or works in the public sector. Third, since our instrument relies on idiosyncratic employer performance (i.e., uncorrelated with spousal productivity), we drop households where both members are currently working or have worked in the same company. After our restrictions, we have a sample of 3.1 million observations, about 440,000 individuals or 220,000 households per year.

The primary worker is defined as the household member with the highest lifetime labor income. We classify someone as employed if they are matched to the Employer-Employee Register. This is similar to conditioning a non-zero salary, non-zero workdays, or if a member is matched with at least one firm. For members that are matched with multiple employers in a year, we use the employer (job) which pays the most or is of the longest duration, in that order.

Table 1: Summary Statistics

Variable	Mean	SD	Percentile 25^{th}	50^{th}	75^{th}
Primary worker					
Participation rate	0.94	0.23	1	1	1
Labor income | work (in $)	65,080	177,858	41,637	57,123	78,333
Hours | work (annual)	1,763	238	1,796	1,950	1,950
College education	0.37	0.48	0	0	1
Age	44.8	8	39	46	52
Secondary worker (spouse)					
Participation rate	0.70	0.45	0	1	1
Labor income | work (in $)	38,819	27,161	23,201	37,993	49,617
Hours | work (annual)	1,659	392	1,138	1,772	1,950
College education	0.44	0.50	0	0	1
Age	43.5	9.05	37	44	51
Household					
Hand-to-mouth households	0.10	0.30	0	0	0
Number of children	1.2	1.1	2	2	3

Notes: Labor market and demographic characteristics for primary worker and spouse. Statistics are averaged over 2015–2019.

Table 1 reports the basic summary statistics. The participation rates in the labor market

are high in Norway. In a typical year, on average, 94 percent of primary workers and 70 percent of spouses work. 55 percent of households are dual-working households, 35 percent are a household where only the primary earner is working, 5 percent are households with no worker, and 1 percent with a spouse working. Norwegian primary workers earn 65,080 per year (in US dollars), and those who work dedicate 1,763 hours to the labor market. Norwegian spouses earn 38,819 (in US dollars) and, conditional on participating in the labor market, work 1,659 hours. The average age and college education attainments are similar between primary workers and spouses.

Since consumption and labor supply responses are more sensitive to the amount of financial assets relative to housing assets, we focus on the liquid component of household wealth (Kaplan and Violante, 2014). The financial accounts in our data include bank deposits, bonds, financial securities, mutual funds, stocks and shares, and other financial assets. On average, assets are 1.7 times the average household income. The share of households with wealth less than three times average monthly earnings (our definition of hand-to-mouth) is 10 percent.

Finally, Table 1 reports sources of household income. Primary workers contribute 40 percent of household income, spouses contribute 20 percent, and capital income adds another 10 percent. The government plays a large role in the income of Norwegian households. One component of government transfers is unemployment insurance (UI). For permanent layoffs, it lasts up to two years and replaces, on average, 62.5 percent of the gross income in the last occupation. Furthermore, while employment protection laws make worker firings difficult, it is relatively easy to get workers on sick benefits into disability programs. As a result, in Norway there is a low unemployment rate, but high fractions of workers on sickness or disability benefits. In addition, sick benefits are paid for up to one year and compensate 100 percent of earnings. Although disability programs have complicated rules, we estimate their average compensation to 66 percent of previous earnings. These generous government transfer programs (net of taxes) contribute on average 30 percent of household income.

3.2 Econometric Framework

Let the *true* (log-linear) model of spousal labor supply and household consumption be given by:[5]

$$h_{s,i,t} = \beta_0 + \beta_c c_{i,t} + \beta_\omega \omega_{s,i,t} + u^h_{s,i,t} \tag{2}$$

$$c_{i,t} = \eta_0 + \eta_y y^P_{i,t} + u^c_{i,t}, \tag{3}$$

[5]We allow for individual fixed effects in both labor supply and consumption but suppressed here to save notation. The actual regression also includes controls for various household characteristics such as age, education, number of children, etc. See below (Section 3.3) for the complete list of control variables for the characteristics of the households.

where $h_{s,i,t}$ is the labor supply of the spouse in household i at time t which can be either extensive (participation decision) or intensive margin (log hours worked). Log of household's consumption is $c_{i,t}$ which depends on the persistent component of household labor income, $y^P_{i,t}$ (more discussion on this assumption below). Finally, $\omega_{s,i,t}$ denotes the (log) productivity of the spouse.[6] Here, β_c represents the income elasticity and β_ω represents the substitution elasticity of labor supply. Residual factors in spousal labor supply and household consumption are denoted by $u^h_{s,i,t}$ and $u^c_{i,t}$, respectively.

Let's assume that both primary and secondary earner's log income consists of persistent and transitory components:

$$y_{p,i,t} = y^P_{p,i,t} + y^T_{p,i,t} \text{ and } y_{s,i,t} = y^P_{s,i,t} + y^T_{s,i,t}. \tag{4}$$

The transitory components ($y^T_{p,i,t}$ and $y^T_{s,i,t}$) are assumed to be i.i.d. and uncorrelated with persistent components. However, persistent components may be correlated between members.

As is well understood, consumption responds more strongly to persistent labor income shocks relative to transitory labor income shocks (Blundell, Pistaferri, and Preston, 2008; Kaplan and Violante, 2010). For this reason, we have modeled the consumption equation in (3) as only a function of the persistent component of household labor income, $y^P_{i,t}$.[7] The log of persistent component of household's labor income can be approximated by the weighted sum of the logs of each member's persistent labor earnings:

$$y^P_{i,t} = s_p y^P_{p,i,t} + (1 - s_p) y^P_{s,i,t}, \tag{5}$$

where s_p gives the share of the primary worker's persistent component of labor income with respect to household's persistent income. Using equations (2),(3), and (5) yields:

$$h_{s,i,t} = \bar{\beta}_0 + \beta_c \eta_y s_p y^P_{p,i,t} + \beta_\omega \omega_{s,i,t} + \beta_c \eta_y (1 - s_p) y^P_{s,i,t} + \beta_c u^c_{i,t} + u^h_{s,i,t}. \tag{6}$$

where $\bar{\beta}_0 = \beta_0 + \beta_c \eta_0$.

The typical regression design in empirical studies of spousal labor supply uses some form of the following specification:

$$h_{s,i,t} = \alpha + \beta y_{p,i,t} + e_{s,i,t}. \tag{7}$$

Comparing equations (6) and (7) highlights the identification challenges in estimating the

[6] Spousal productivity is only observed only for working spouses (through wages).

[7] Blundell, Pistaferri, and Preston (2008) find that transitory shocks are nearly perfectly insurable. Using a relative standard Aiyagari (1994) model, Kaplan and Violante (2010) find a higher pass-through of transitory shocks when households are borrowing constrained but still substantially less than the pass-through of persistent shocks.

spousal labor response. First, the error term $e_{s,i,t} = \beta_\omega \omega_{s,i,t} + \beta_c \eta_y s_s y^P_{s,i,t} + \beta_c u^c_{i,t} + u^h_{s,i,t}$. The first two components in the error term, the productivity of the spouse ($\omega_{i,s,t}$), and the persistent component of the spousal income ($y^P_{s,i,t}$), are plausibly correlated with the regressor ($y_{p,i,t}$). This violates the orthogonality condition, and the OLS estimate of β would be biased.

The direction of the bias depends on the relative strengths of the substitution and income effect of spousal labor supply (i.e., the magnitude of the Marshallian elasticity of labor supply, as we illustrated in Section 2). To begin with, it is natural to assume that productivity (or income) of married couples is positively correlated. Thus, the first term is likely to bias β toward zero ($cor(y_{p,i,t}, \omega_{i,s,t}) > 0$ and $\beta_\omega > 0$) and the second term is likely to generate a negative bias in β ($cor(y_{p,i,t}, y^P_{s,i,t}) > 0$ and $\beta_c < 0$). We believe that the bias due to the second term is negligible because the value of $\beta_c \eta_y (1 - s_p)$ is close to zero. For example, the typical value of consumption elasticity in labor supply is $\beta_c = -1$ and that of income elasticity in consumption is $\eta_y = 0.65$ (Blundell, Pistaferri, and Preston, 2008). The average share of spousal income in our sample is $s_s = 0.23$. These values imply $\beta_c \eta_y (1 - s_p) = -0.14$. Given the transitory nature of the spousal labor supply–because secondary workers frequently move into and out of the labor market, the income share of spouses in *persistent* income would be much smaller than 0.23. Thus, the bias due to the second term is almost negligible. Whereas the participation elasticity of the labor supply (which is our main interest) is often estimated to be high and even close to 1: $\beta_\omega = 1$ (e.g., Chang and Kim (2006)).[8] Thus, the bias from the first term is not negligible.

The second issue with the identification of the spousal labor supply is that equation (7) uses $y_{p,i,t} = y^P_{p,i,t} + y^T_{p,i,t}$, i.e., contains persistent and transitory components, while only the persistent component affects consumption. Thus, using the observed income (instead of the persistent income) biases the estimates β toward zero (attenuation bias).

To solve the orthogonality problem, we use information from the primary worker's firm: the sales revenue of the firm where the primary worker is or had been employed ($\Pi_{p,i,t}$). This is a valid instrument because (i) the performance of the primary worker's firm is likely uncorrelated with spousal productivity ($corr(\Pi_{p,i,t}, \omega_{s,i,t}) = 0$) and (ii) it is significantly correlated with the primary worker's earnings ($corr(\Pi_{p,i,t}, y_{p,i,t}) \neq 0$).[9] As it turns out, predicted worker's earnings based on our instrument (firm's sales revenue) is very persistent (with an autocorrelation of 0.95) and thus, it also helps to address the attenuation bias that plagues typical estimates of the AWE. We close the discussion with one caveat. The spousal wage can be used as a proxy for productivity $\omega_{s,i,t}$, in the regression of (6). But wages are observed only when the

[8] Erosa, Fuster, and Kambourov (2016) reports labor supply elasticity of 0.5 for intensive margin.

[9] Our instrument (sales revenue of primary worker's employer) is also orthogonal to the persistent component of spousal income ($cor(\Pi_{p,i,t}, y^P_{s,i,t}) = 0$) and thus eliminates the bias due to the second term in the error term as well.

spouse worked, and our primary interest is estimating the added worker effect (participation decision).

3.3 Two Stage Regression

To eliminate individual fixed effect, consider regressing (7) in first differences:

$$\Delta h_{s,i,t} = \beta \Delta y_{p,i,t} + X'_{i,t}\delta + \upsilon_{s,i,t}. \tag{8}$$

where $X_{i,t}$ controls for household characteristics, which include the number and age of children, change in the log of household's financial assets, age, and education of the primary worker and spouse, dummy variables indicating that the primary worker or the spouse are on sick or parental leave, and finally, dummy variables for time and industry.

The OLS estimate for the elasticity of spousal participation in primary income (β) is close to zero: -0.00 (0.01) where the number in parentheses represents the standard error. This is not surprising given that the literature did not find strong evidence on the added worker effect (e.g., Doepke and Tertilt, 2016). We also estimate (8) using spouse's earnings as a dependent variable. Again, the OLS estimate is also close to zero: 0.009 (0.005).

As discussed in Section 3.2, the OLS estimate of β is likely to be biased towards zero. First, the correlation between primary income and spousal productivity. Second, there is an attenuation bias because the regressor (measured income) contains persistent and transitory components, whereas consumption responds only to persistent income.[10]

We use the firm-side information of the primary worker as an instrumental variable to resolve the issues with the OLS estimation. Our instrument is firm's sales revenue, where the primary worker is employed. The identifying assumption is (i) that an individual worker cannot influence the performance of the firm, and (ii) that the primary worker's employer's sales revenue is not correlated with the productivity of the spouse. At the same time, we require that sales revenue be correlated with the primary worker's earnings, which is very natural. This may not be the case if workers' earnings are set based on aggregate conditions only, such as collective bargaining agreements. In fact, collective bargaining-based wage setting is common in Norway. However, the collective bargaining only sets a wage *norm*, where there is room for wage adjustments based on worker or firm-related idiosyncratic factors (Fagereng, Guiso, and Pistaferri, 2018). We scale the firm's sales revenue (gross revenue minus operating

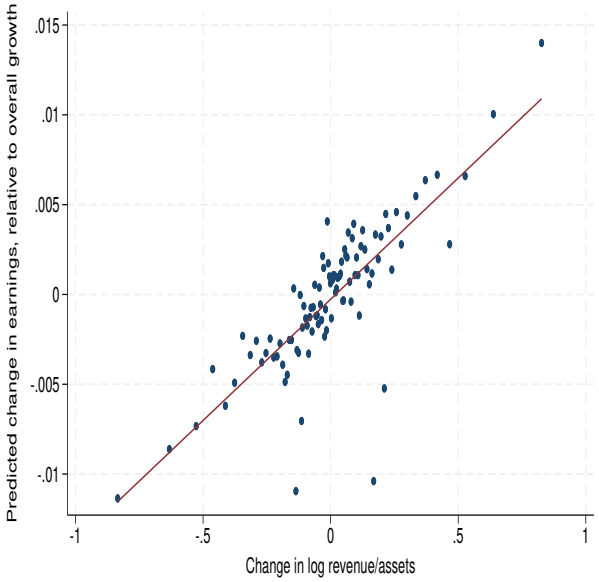
[10]In addition, when households anticipated the changes in (primary) income, they may have already responded by adjusting spousal labor supply and consumption. In fact, a recent literature suggests that a substantial portion of the residual variation in earnings is predictable and reflects individual choices rather than risk (e.g., Primiceri and van Rens, 2009, Guvenen and Smith, 2014). To the extent that income changes are anticipated, the estimated response would be muted.

costs) by the size of the firm's gross assets, denoted by $\Pi_{p,i,t}$. We also experimented with other measures, such as value added, but found that sales revenue works better than value added. The first-stage regression is:

$$\Delta y_{p,i,t} = \phi \Delta \Pi_{p,i,t} + \mathbf{X}'_{i,t}\zeta + \xi_{p,i,t} \tag{9}$$

where ϕ reflects the pass-through coefficient from the firm's revenue growth to the worker's income growth.

Table 2: The First Stage: Pass-through of Sales Revenue onto Worker's Earnings



Dependent Variable = $\Delta y_{p,i,t}$		
	Coefficient	S.E.
$\Delta \Pi_{p,i,t}$ (sales revenue)	0.0135	0.0008
Δ ln household's assets	0.016	0.0007
Number of children	0.004	0.003
Primary, age	−0.0025	0.00009
secondary school	−0.019	0.0008
college	−0.009	0.001
sick leave	−0.039	0.0008
parental leave	−0.072	0.0022

Note: The left panel shows a binscatter plot of changes in firm log revenue-to-assets and changes in primary worker's earnings. The right panel reports estimates of regression (9) for selected control variables.

Table 2 shows the estimated pass through of sales revenue-to-assets onto worker's earnings, the first-stage regression. The left panel of Table 2 is the scatter plot between the primary worker's earnings growth ($y-$axis) and sales revenue growth ($x-$axis). It clearly shows that primary income is highly correlated with sales revenue. The right panel provides some estimates of the first-stage regression (9).[11] The estimated pass-through coefficient $\widehat{\phi} = 0.0135$ with a t ratio of 13. The small pass-through indicates that workers' income is, by and large, insured from the year-to-year fluctuations of sales revenue in Norway. However, some part of year-to-year fluctuations remains uninsurable and constitutes an income effect that can influence the spousal labor supply. Based on our estimates, primary earnings growth decrease

[11]The regression also includes other controls such as time and industry dummies, etc.

Table 3: Spousal Response to Primary Income Shocks: 2SLS

Dependent Variable ($\Delta y_{s,i,t}$) :	Participation (1)	Earnings (2)
$\widehat{\Delta y}_{p,i,t}$ (primary earnings)	-0.13^{***}	-0.37^{**}
	(0.04)	(0.15)
Δ ln household's assets	0.016	0.0007
	(0.04)	(0.15)
Number of children	0.001^{***}	0.004^{***}
	(0.000)	(0.000)
Primary, age	0.00	0.001^{**}
	(0.00)	(0.000)
sick leave	-0.008^{***}	-0.018^{***}
	(0.002)	(0.005)
parental leave	-0.059^{***}	-0.144^{**}
	(0.004)	(0.014)
Spouse, age	0.002^{***}	-0.08^{**}
	(0.00)	(0.000)
sick leave	-0.054^{***}	-0.145^{***}
	(0.001)	(0.005)
parental leave	-0.087^{***}	-0.305^{**}
	(0.004)	(0.014)
Observations	813,035	661,653

Notes: The instrumental variable is the sales revenue/assets of the firm that employed the primary earner. The regression also includes education, year, and industry dummies.

over the life cycle, by 0.25 percent on average. When the primary worker is on sick or parental leave year-to-year earnings growth declines by 3.9 and 7.2 percent, respectively.

The second stage regression, replaces $\Delta y_{p,i,t}$ in Equation (8) with $\widehat{\Delta y}_{p,i,t}$ from (9). Table 3 shows the estimated spousal response. We show estimates for spousal labor market participation and spousal log labor earnings. Earnings are conditional on participating in the labor market in the previous year (that is, changes in the intensive margin only).

The instrumental variable regression generates the response of the spousal labor supply and earnings ($\widehat{\beta} = -0.13$ and -0.37, respectively) that is in line with our theory: the spouse increases the labor supply in response to a decrease in primary income. According to our estimate, in response to a 10 percent decrease in primary income, the probability of the spouse participating in the labor market increases by 1.3 percentage points and the spousal earnings increase by 3.7 percent. Both estimates are statistically significant.

Table 4: Persistence and Variance of Primary Income

	All Sample	Fully employed
ρ_y	0.30	0.95
σ_y	0.34	0.16
Obs.	1,572,483	930,075
$\rho_{\widehat{y}}$	0.94	0.97
$\sigma_{\widehat{y}}$	0.10	0.10
Obs.	1,171,675	735,784

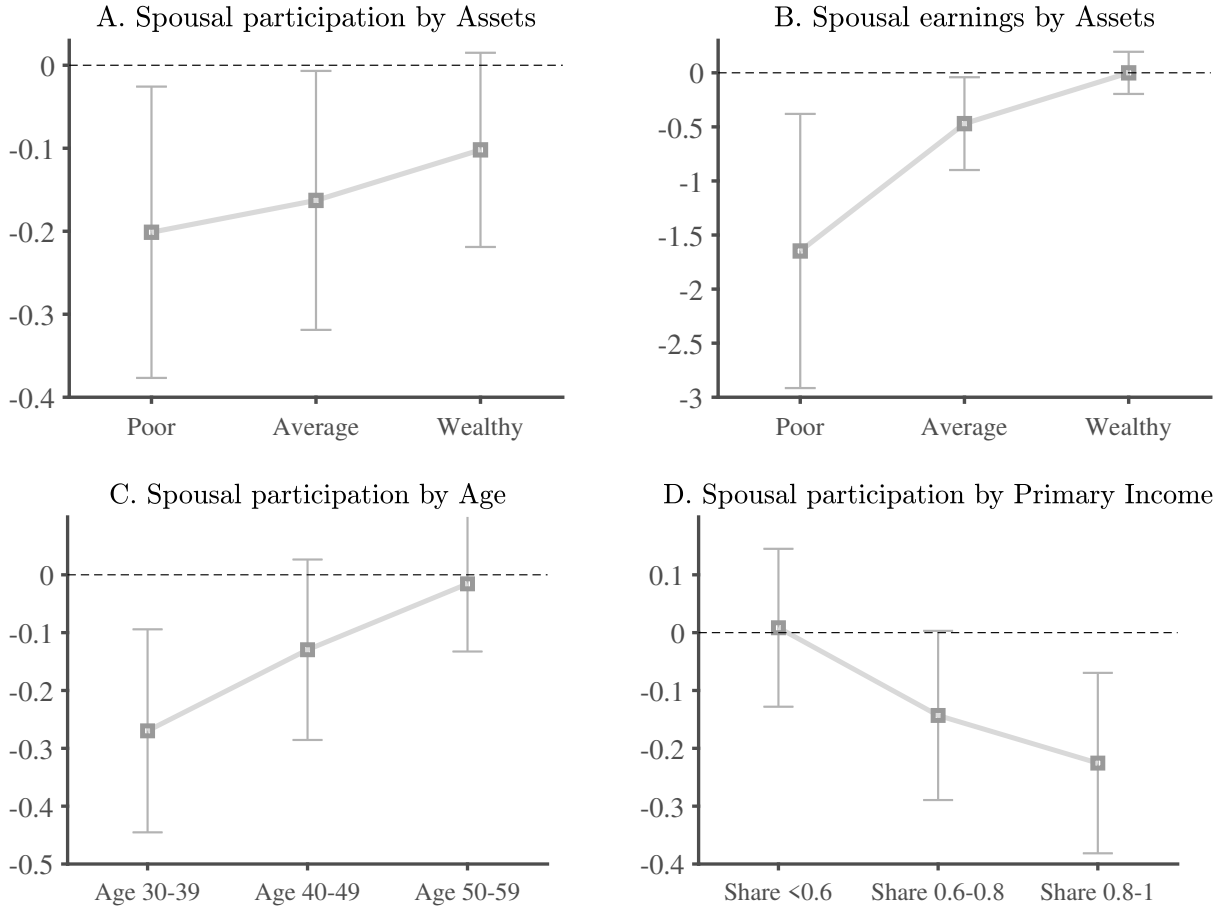
Notes: Estimates from a first-order autoregression of primary income and primary income projected by firm's sales-to-assets. Fully employed includes primary worker that work continuously throughout the year.

In households with a larger number of children, the spouse participates in the labor market and earns slightly more. When the primary worker takes a sick leave, the spousal participation decreases by 0.008 percentage points and earnings decline by 1.8 percent. The spousal participation decline when primary workers take sick leave likely indicates that spouses stay home to take care of a partner in need. When the primary worker takes a parental leave, spousal participation also decreases by 0.05 percentage points. Sickness and parental leave for the spouse reduces the supply of spousal labor which is natural.

As we explain in Section 3.3, the OLS estimate is biased toward zero not only because of correlated productivity, but also because of an attenuation bias: the observed income reflects both persistent and transitory components of income, while only the persistent component generates income effects in the labor supply. Our instrument variable (sales revenue of the employer of the primary worker) is uncorrelated with the productivity of the spouse and, at the same time, highly persistent over time.

To see the latter, we report the time series properties of $y_{p,i,t}$ and $\widehat{y}_{p,i,t}$ (primary income projected by sales revenue). Specifically, we run a first-order auto-regression and estimate the persistence and variance of the process: (ρ_y, σ_y) and $(\rho_{\widehat{y}}, \sigma_{\widehat{y}})$, respectively. Table 4 reports the estimates for the whole sample and a subset of workers that are fully employed during the year. For the entire population, primary income is not highly persistent ($\rho_y = 0.3$) because some workers are often part-time employed or sometimes out of the labor market. Indeed, when we restrict the sample to primary workers who are full-time employed, the persistence

Figure 1: Heterogeneity in Spousal Labor Supply Response



Notes: We run the IV regression separately for each group. We report point estimates and the 95 percent confidence intervals. Upper left and upper right panel reports the estimated of participation and earnings elasticities, respectively, by wealth. Lower left and lower right panel reports the estimated of participation elasticities, by age and the share of primary worker's income, respectively.

increases considerably to 0.95. The projected primary income has much higher persistence: 0.94 for the entire population and 0.97 for full-time employed. At the same time the projected primary income has a lower variance because the instrumental variable method is eliminating noisy, transitory fluctuations.

3.4 Heterogeneity in Spousal Labor Supply Responses

We take advantage of our detailed data and analyze the heterogeneity in the response of spousal labor supply. According to our economic theory in Section 2, the response of spousal labor is stronger for households that do not have sufficient means to insure against primary

income fluctuations. To examine this, we classify the households into three groups based on their financial assets: (i) "wealth-poor" households whose financial assets as of the beginning of the year are less than one month's average earnings in Norway, (ii) 1-6 months' average earnings ("average wealth"), and (iii) more than 6 months' average earnings ("wealthy"). We run regression (8) separately for each group and report the coefficients in the upper panel of Figure 1. The response shows a clear pattern that is consistent with our theory. Looking at the upper-left panel, poor "hand-to-mouth" households (whose financial assets are less than one month's average household income) exhibit a stronger response: $\widehat{\beta} = -0.2$ about 50% larger than the average.

The heterogeneity in spousal earnings response is even more evident (upper right panel). When primary income decreases by 10 percent, wealth-poor households increase their earnings by 16 percent about 4 times more than the average.

The lower left panel of Figure 1 shows the participation responses of the spouse by age. Young households (whose primary earners are 30-39 years old) show a response twice as large as the average of the entire population because they have less wealth on average (Bacher, Grübener, and Nord, 2024). According to our simple formula in Section 2, the elasticity of the spousal labor supply is greater when the primary income share is large. This is verified in the data. The lower right panel of Figure 1 shows that households where primary income is the dominant source of income (larger than 80% of total income) exhibit a response that is almost twice as large as the average.

4 Quantitative Analysis

We build a quantitative model with heterogeneous households and calibrate it to match the salient features of life-cycle profiles of labor supply and earnings with two goals in mind. First, we examine whether the standard model can reproduce the estimated response of the cross elasticity we find in Section 3. Second, the model will be used to evaluate various government policies, with special interest in the effects on the joint labor supply.

4.1 Economic Environment

Demographics We consider a small open economy populated by a continuum of households with a total measure of one. Households are either born as a single or married households and keep their status throughout the life cycle (Guner, Kaygusuz, and Ventura, 2023). Married households, with measure $1 - \pi_m$ consist of two potential earners: a primary worker and a spouse indexed by $i = \{p, s\}$. Similarly to the data, primary workers have higher lifetime labor income (due to higher productivity) than spouses. Although the primary worker is one

year older than the spouse on average in our data, we assume that married individuals have the same age (indexed by j). Each individual enters the labor market at age $j = 1$ and lives until age J. We assume that there is a mandatory retirement age j_R. But some workers may retire before that age if they choose to do so. Starting from age j_R, households receive social security benefits from the government. Finally, we assume that there is an age-dependent probability (s_j) that the household survives.

Preferences Each household maximizes lifetime utility by choosing consumption, c_j, and labor supply of the primary worker and his spouse, h_p and h_s:

$$U = \mathbf{E}\Big[\sum_{j=1}^{J} \beta^{j-1} \mathbf{\Pi}_{k=1}^{j} s_k \ \Big\{\frac{c_j^{1-\sigma}}{1-\sigma} - \psi_p \frac{h_{p,j}^{1+\frac{1}{\gamma}}}{1+\frac{1}{\gamma}} - \psi_s \frac{h_{s,j}^{1+\frac{1}{\gamma}}}{1+\frac{1}{\gamma}}\Big\}\Big] \tag{10}$$

where β is the discount factor, σ the coefficient of relative risk aversion, and γ the Frisch elasticity of labor supply. We assume that disutility from working is different between primary worker and spouse (ψ_p and ψ_s).

Number of Children We assume that the number of children is exogenous and depends on parents' age (Borella, De Nardi, and Yang, 2023). Although the number and age of children are important for household consumption and labor supply, we assume that the cost of children occurs only in monetary form. Each child below age 5 is associated with a cost ζ_5. Thus, the household incurs (child care) expenses $\zeta_5 n_5(j)$ where $n_5(j)$ is the number of children below fiver years old.

Productivity The log productivity (wage) of worker i at age j consists of four components: a permanent fixed effect (skill), α_i, human capital (accumulated over time through labor market experience), $\theta_{i,j}$, persistent productivity shock, $x_{i,j}$, and an i.i.d. shock, $\varepsilon_{i,j}$:

$$\ln w_{i,j} = \ln \alpha_i + \ln \theta_{i,j} + \ln x_{i,j} + \ln \varepsilon_{i,j}.$$

In the beginning of life cycle, the household draws a fixed effect: $\ln \alpha \sim N(0, \sigma_\alpha^2)$, which is common to the couple (i.e. complete positive assortative matching). With permanent skill heterogeneity, government transfers serve not only an insurance role but also a redistributive role. In addition, we allow for the accumulation of human capital through the labor market experience, an important feature when evaluating the effect of fiscal policy on female labor supply (Blundell, Costa Dias, Meghir, and Shaw, 2016). Specifically, human capital $\theta_{i,j}$ increases with labor market experience $l_{i,j}$. The exact functional form of $\theta(l)$ is described in the

following when we calibrate the model.

As we described in Section 3, the response of spousal labor supply depends on the nature of shocks to primary income: persistence and correlation with the productivity of the spouse. To accommodate both features, we assume that idiosyncratic productivity shocks follow a jointly log-normal AR(1) process:

$$\ln x_{i,j} = \rho_x \ln x_{i,j-1} + \nu_{i,j}, \qquad \text{with} \quad \begin{bmatrix} \nu_p \\ \nu_s \end{bmatrix} \sim N\Big(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \sigma_\nu^2 & \sigma_{ps} \\ \sigma_{ps} & \sigma_\nu^2 \end{bmatrix} \Big). \tag{11}$$

where ρ_x represents the persistence and ν_p and ν_s, the innovations to primary worker and spouse, respectively. Innovations have variance σ_ν^2 and covariance σ_{ps}. The transition matrix for this autoregressive process is denoted by $\Gamma_{x,x'}$ where $x = \{x_p, x_s\}$. The initial value of x is drawn from a log-normal distribution with zero mean and variance $\sigma_{x(1)}^2$. Finally, we assume that idiosyncratic productivity shocks ε are independently distributed between time *and* family members with variance σ_ε^2. The probability distribution of these transitory shocks is denoted by $\pi(\varepsilon)$.

Household Labor Supply Households make decisions in both the intensive and extensive margin of labor supply. Let E stand for employment and N non-employment. The employment status of a household at age j is denoted by $e_j = \{EE, EN, NE, NN\}$. We assume that moving from status N to E requires a one-time, age-invariant, utility cost ξ reflecting costs associated with finding and moving to a new job from non-employment. We assume that these costs may be different for p and s. This labor market entry cost serves two purposes. First, it helps match the average levels of the participation rates of the primary worker and the spouse. Second, it influences the frequency of transitions between E and N. Without this cost, the model would predict too frequent transitions between labor market states.

In a model where productivity shock is the main driver of employment decisions, it is mainly wealthy and/or low-productivity workers who are non-employed, whereas the employment rate is weakly correlated with productivity or wealth in the data (Mustre-del Rio, 2015). We introduce an exogenous separation probability λ_i to weaken this correlation.

Private and Public Insurance We assume that the capital market is incomplete (Aiyagari, 1994). The households insure themselves against income risk by accumulating a risk-free asset a that pays a rate of return r (precautionary motive). They also save for retirement (life-cycle motive). There is an exogenous borrowing constraint, $\underline{a}$: $a' \geq \underline{a}$.

The government plays two roles in the model. First, it taxes household income using a tax function $-T + \tau Y$, where T represents lump-sum transfers, τ proportional tax rate, and

Y household income. This simple tax function reflects that in Norway separate tax filing for household members is not available. In addition, the government runs a social security system. When the household retires from the labor market at age j_R, it receives a social security benefit b. We assume for simplicity that all workers receive the same social security benefit, regardless of the history of earnings.[12]

The government spends tax revenue to pay (i) lump sum transfers, (ii) social security benefits, and (iii) government spending G (which does not enter into the household's utility). Thus, the government budget constraint is:

$$\mu_w \tau Y = \mu_w T + (1 - \mu_w) b + G \tag{12}$$

where μ_w is the share of the working age population. The tax parameter τ is adjusted to balance the budget.[13]

Definitions We denote the state of a household with age j and permanent productivity α by:

$$\boldsymbol{z} = (a, x_p, x_s, l_p, l_s, \varepsilon_p, \varepsilon_s, e_{p,j-1}, e_{s,j-1}; j, \alpha)$$

where a denotes asset holdings, x_p, x_s are the individual persistent productivity levels, l_p, l_s are the accumulated human capital stocks, $\varepsilon_p, \varepsilon_s$ are i.i.d. shocks, and $e_{p,j-1}$, $e_{s,j-1}$ denotes the previous employment status for primary worker and spouse, respectively.

Value Functions We present value functions for three types of households: a dual-earner household (e=EE), a household where only the primary worker is working ($e = EN$), and a retiree. The other sets of value functions are provided in the Appendix. To save space, we write the value functions as if there is no exogenous separation.

Dual Earner Household:

$$V^{EE}_{j,\alpha}(\mathbf{z}) = \max_{c,a',h_p,h_s} \left\{ \frac{c_j^{1-\sigma}}{1-\sigma} - \psi_p \frac{h_{p,j}^{1+\frac{1}{\gamma}}}{1+\frac{1}{\gamma}} - \psi_s \frac{h_{s,j}^{1+\frac{1}{\gamma}}}{1+\frac{1}{\gamma}} - \sum_{i=\{p,s\}} \xi_i \mathbb{1}_{\{e_{i,j-1}=N\}} \right.$$
$$\left. + \beta v_{j+1} \sum_{x_p'} \sum_{x_s'} \sum_{\varepsilon_p'} \sum_{\varepsilon_s'} \pi(\varepsilon_p') \pi(\varepsilon_s') \Gamma_{x'|x} \max\left\{ V^{EE}_{j+1,\alpha}(\mathbf{z}'), V^{EN}_{j+1,\alpha}(\mathbf{z}'), V^{NE}_{j+1,\alpha}(\mathbf{z}'), V^{NN}_{j+1,\alpha}(\mathbf{z}') \right\} \right\}$$

[12] As a robustness exercise, we also allow the social security benefit dependent on the human capital level before the retirement year (Guvenen, 2007).

[13] Labor income taxes are not sufficient to fund both transfers, social security benefits and government spending. Thus, we interpret G more broadly to reflect also annual receipts from the Norwegian oil fund that amount to around 20 percent of the government budget.

$$
\begin{aligned}
\textbf{s.t.} \quad & c + a' = (1-\tau)(y_p + y_s) - \zeta_5 n_5(j) - \zeta_{10} n_{10}(j) + T + (1+r)a \\
& a' \geq 0 \\
& y_p = \alpha x_p \theta_{p,j} \varepsilon_p h_p \\
& y_s = \alpha x_s \theta_{s,j} \varepsilon_s h_s \\
& \boldsymbol{z'} = (a', x_p', x_s', l_p + 1, l_s + 1, \varepsilon_p', \varepsilon_s', E, E)
\end{aligned}
$$

Single Earner Household:

$$
\begin{aligned}
V_{j,\alpha}^{EN}(\mathbf{z}) = \max_{c,a',h_p} \; & \left\{ \frac{c_j^{1-\sigma}}{1-\sigma} - \psi_p \frac{h_{p,j}^{1+\frac{1}{\gamma}}}{1+\frac{1}{\gamma}} - \xi_p \mathbb{1}_{\left\{e_{p,j-1}=N\right\}} \right. \\
& \left. + \beta v_{j+1} \sum_{x_p'} \sum_{x_s'} \sum_{\varepsilon_p'} \sum_{\varepsilon_s'} \pi(\varepsilon_p') \pi(\varepsilon_s') \Gamma_{x'|x} \max\left\{ V_{j+1,\alpha}^{EE}(\mathbf{z}'), V_{j+1,\alpha}^{EN}(\mathbf{z}'), V_{j+1,\alpha}^{NE}(\mathbf{z}'), V_{j+1,\alpha}^{NN}(\mathbf{z}') \right\} \right\}
\end{aligned}
$$

$$
\begin{aligned}
\textbf{s.t.} \quad & c + a' = (1-\tau) y_p - \zeta_5 n_5(j) - \zeta_{10} n_{10}(j) + T + (1+r)a \\
& a' \geq 0 \\
& y_p = \alpha x_p \theta_{p,j} \varepsilon_p h_p \\
& \boldsymbol{z'} = (a', x_p', x_s', l_p + 1, l_s, \varepsilon_p', \varepsilon_s', E, N)
\end{aligned}
$$

Retired Household:

$$
V_j^R(a) = \max_{c,a'} \left\{ \frac{c_j^{1-\sigma}}{1-\sigma} + \beta s_{j+1} V_{j+1}^R(a') \right\} \tag{13}
$$

$$
\begin{aligned}
\textbf{s.t.} \quad & c + a' = b + (1+r)a \\
& a' \geq 0.
\end{aligned}
$$

4.2 Calibration

There are five sets of parameters to calibrate: (i) demographics and life-cycle parameters $\{j_R, J, s_j, \pi_m, n_{5,j}, \zeta_5\}$, (ii) preferences parameters $\{\sigma, \gamma, \beta, \psi_i, \xi_i\}$, (iii) asset-market structure $\{r, \underline{a}\}$, (iv) parameters governing the productivity process $\{\alpha, \rho_x, \sigma_\nu^2, \sigma_{ps}, \sigma_\varepsilon, \sigma_{x(1)}^2, \lambda_i, \chi_0^i, \chi_1\}$, and finally, (v) government tax and transfers, government spending, and the social security benefit $\{\tau, T, G, b\}$. We classify the parameters into two sets. In the first set, the parameter values are directly set externally by the data or the existing literature. The remaining parameters are chosen to match some key moments, including the life cycle pattern of labor supply

and earnings of married households in Norway.[14] Table C-1 reports all parameter values.

Externally Set Parameters The model period is one year. Married households are 78 percent of the population. Households are born and enter the labor market at $j = 1$ and live for 80 periods, $J = 80$. This life cycle corresponds to ages 21 to 100. All household members retire exogenously at $j_R = 45$ (age 65) when they start receiving the social security benefit b. The survival probability $\{s_j\}$ and the number of children per married household below the age of 5 is from Statistics Norway (depicted in Figure B-1 in the Appendix). We set ζ_5 based on an annual daycare expense cost equal to \$2,500. The risk aversion σ is set to 2 and the Frisch elasticity of labor supply (in the intensive margin) γ is 0.5, parameters common in the literature.

Our model does not have multiple assets, so we use a single interest rate that captures the overall rates of return. According to estimates from Klovland (2004), the annual rate of return in Norway is 1.4% for risk-free assets and 4.5% for stocks, respectively. According to Chang et al. (2022) the average share of risky assets is 57% in Norway. This gives the average rate of return $r = 3.2\%$. The proportional tax rate, $\tau = 0.32$, is set to match the average tax liability of household income in the data. The social security benefit is set so that it matches the average replacement rate (35%) in the Norwegian social security system.

We estimate a set of income parameters directly from the data of the Norwegian panel. For persistent productivity shocks, we set $\rho_x = 0.955$ (based on our estimates in Table 4) and $\sigma_\nu^2 = 0.025$ (which is close to that in Blundell, Pistaferri, and Saporta-Eksten (2016) based on U.S. panel data). For the variance of transitory shocks, we use $\sigma_\varepsilon^2 = 0.04$ from Blundell, Pistaferri, and Saporta-Eksten (2016). Finally, conditional on education levels, we estimate a correlation between primary and spousal hourly wages equal to 0.20. We use this estimate to calibrate σ_{ps}.

Internally Fitted Parameters We calibrate the discount factor β to match the asset-to-income ratio we observe in the data. We also give each household an initial endowment of assets to match the assets-to-income ration at age 25. The disutility of work ψ_i is set so that on average primary earners work around 34% of their time and spouses 28% of their time. The utility cost of moving from non-employment to employment ξ_i is chosen to match the average participation rates. The exogenous separation rates λ_i are selected by targeting the average transition rate from employment to non-employment.

The variance in initial productivity $\sigma_{x(1)}$ is useful in obtaining a gradually increasing profile of employment rates at the beginning of the life cycle. Without initial dispersion (for example,

[14] For singles, we re-calibrate the discount factor β and the cost ξ to observed asset holdings and participation rates and keep all other parameter values same.

Table 5: Comparison between Model and Data

Statistic	Model	Data	Statistic	Model	Data
Assets/Income	1.80	1.81	Income share, p	0.41	0.49
Participation, p	0.91	0.91	Income share, s	0.30	0.23
Participation, s	0.68	0.68	Gov. transfers share	0.20	0.20
Hours of work, p	0.35	0.34	Share of dual earners	0.60	0.68
Hours of work, s	0.29	0.28	Share of single earners	0.30	0.23
$N \to E$ rate, p	2.9%	2.9%	Debt/Income	0.05	0.05
$N \to E$ rate, s	2.4%	3.2%	Wage growth w_{50}/w_{25}, p	1.46	1.54
Wage discount w_s/w_p	0.77	0.75	Wage growth w_{50}/w_{25}, s	1.45	1.50

Notes: Comparison of the model and the data with respect to income, labor market, and financial statistics.

households starting with the same, median productivity), households behave similarly: either they jump into employment from the beginning or they stay out. We set the dispersion in initial productivity to roughly match the primary employment profile for households between the ages 25-40 in the data.

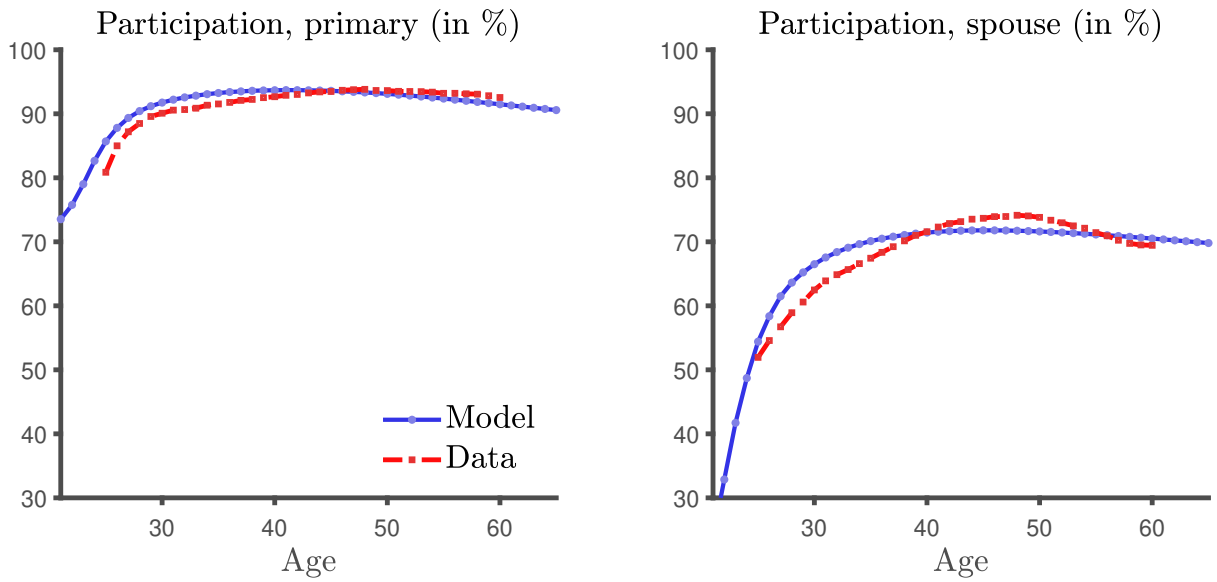
We assume that human capital evolves with labor market experience according to: $\log \theta_{i,j}(l) = \log \chi_0^i + \chi_1^i \log(1 + \frac{1}{q}(l-1))$ where q represents the probability with which human capital increases to a next level. We assume that there are three levels of labor market experience $l = \{1, 2, 3\}$ and thus the probability of accumulation of the experience is $q = \frac{1}{15}$ $(= 3 \times \frac{1}{45})$. At the beginning of the life cycle, all workers have the same level of human capital ($l = 1$). We normalize $\chi_0^p = 1$ and set $\chi_0^s = 0.75$ to reflect an average wage discount for female workers in the data equal to 25 percent. The slope parameter is chosen to match the average lifetime growth in earnings for the primary and spouse (both around 150 percent between ages 25 and 50): $\chi_1 = 0.1$ for both.[15] The size of the lump sum transfers, T, is set to match the average share of government transfers in household income in the data (20 percent).

4.3 Model Fit

Table 5 summarizes some statistics from the model and data. The model reproduces the data very well in terms of the asset-to-income ratio, debt-to-income ratio, average employment rates, and hours worked. The model is also reasonably close to the income shares and the share

[15]Both estimates of the gender-wage gap and the life cycle growth of earnings in the Norwegian data are broadly comparable to estimates for the U.S. labor market (see for example, Heathcote, Storesletten, and Violante, 2010 and Guvenen, Karahan, Ozkan, and Song, 2021).

Figure 2: Life Cycle Profiles of Employment: Model vs. Data



Notes: Average labor market participation by age in the model versus the data. Primary worker participation in the left panel and spousal participation in the right panel.

of households with a dual and single earner (both statistics were not targeted). In addition, in the data, the variance of log-consumption increases by around 4 percentage points between ages 25 and 60. Although not targeted either, the model is able to reproduce closely this fact.

Figure 2 shows the life cycle profiles of labor market participation for primary workers and spouses. In general, the model replicates closely the labor market participation rate in the data. The average participation rate of both primary workers and spouses is matched using the entry costs of participation. In the data, the life cycle participation profile of both primary worker and spouse features a gentle hump shape. The model can generates a similar shape thanks to the interaction of five factors: the entry cost of participation, age-depedent child care costs, initial dispersion in productivity, gradual accumulation of human capital through labor market experience, and the wealth effect from the accumulated assets. Due to the entry cost, low-productivity households delay their labor market entry and wait for a more favorable shock to occur. In addition, estimated child expenses decline after their peak at the age of 30. At the same time, to build both human and financial capital, workers are strongly attached to the labor market in the middle part of the life cycle. Closer to retirement, workers gradually exit the labor market as they accumulate enough financial wealth.

4.4 Cross Elasticity

We would like to examine whether the standard model can replicate the elasticity of the spousal labor supply to the shock of the primary income in the data. In order to do this, we generate a panel data from the model and estimate the cross-elasticity, using the exact same specification (8) used in the empirical analysis:

$$\Delta h_{s,i,t} = \beta \Delta y_{p,i,t} + \mathbf{X}'_{i,t}\delta + v_{s,i,t} \tag{14}$$

where $h_{s,i,t}$ is the outcome variable (participation or log earnings) and $y_{p,i,t}$ is the log of primary income and $\mathbf{X}_{i,t}$ includes the permanent productivity (α), ages, and changes in asset holdings of the household.

In the data, we found the OLS estimate nearly zero due to (i) a positive correlation of productivity between the primary worker and the spouse *and* (ii) attenuation bias caused by transitory components of income. Although we cannot perform an IV estimation in the model (there is no firm variable), we test these effects as follows. We consider three specifications of the model. (1) We allow for both correlated productivity and transitory shocks. The estimation of β based on this model will suffer from both biases. This corresponds to the OLS estimate from the empirical analysis. (2) We shut down the transitory shocks and only allow for correlation of productivity between members. (3) We shut down both the correlation and transitory shocks. This case corresponds to the IV estimation in the empirical analysis.

Table 6: Spousal Labor Supply Response: Model vs. Data

	Data		Model	
	Participation	Earnings	Participation	Earnings
OLS	−0.00	0.009	−0.01	−0.04
IV	−0.13	−0.37	−0.10	−0.33

Notes: OLS in the model reports the estimate of regression 14 using generated data from benchmark model. IV is the estimate from generated data using a model without the correlation of productivity between members and the transitory shocks.

Table 6 compares the IV and OLS estimates between the model and the data. Similarly to the OLS estimate in the data, the model produces a close-to-zero estimate for the spousal response to primary income shocks. But when we shut down both sources of bias, the estimates of $\hat{\beta}$ from the model-generated data are −0.10 and −0.33, respectively, for participation and

earnings, broadly comparable to the IV estimates from the empirical analysis.

Table 7 further decomposes the source of bias using estimates from three specifications of the model. Starting from column (1), when we allow for both a positive correlation of productivity and transitory shocks, the estimates of spousal response are nearly zero for both participation and earnings. In column (2) we shut down transitory shocks and allow a positive correlation of productivity between primary worker and spouse. The estimates for participation and earnings elasticity are $\hat{\beta} = -0.08$ and -0.11, respectively. Thus, correlated productivity explains accounts for a small fraction of the bias in the participation elasticity but accounts for a larger part of the bias in the earnings elasticity.

The reason that correlated productivity does not substantially attenuate the estimate in the participation regression (-0.08 versus -0.10) is that (i) spousal participation mostly responds to big negative primary income changes, and (ii) large correlated income declines are a relatively rare event. For example, in the model, an annual primary income decline of around 10 percent induces around 1 percent of spouses to enter the labor market. In contrast, when both the primary and spousal income decline by 10 percent, nearly no spouse enters the labor market. But since joint income declines are rare, the regression estimate decreases only slightly. At the same time, the attenuation shows more naturally in the earnings regression (-0.11 versus -0.33) because it includes working spouses that adjust more flexibly earnings to any size of shock.

Table 7: Decomposition of Spousal Elasticity

Specification	(1)	(2)	(3)
$corr(x_p, x_s)$	0.2	0.2	0
Transitory Shocks (σ_ε)	0.04	0	0
$\hat{\beta}$ (participation)	-0.01	-0.08	-0.10
$\hat{\beta}$ (earnings)	-0.04	-0.11	-0.33

Notes: Estimates of regression 14 for participation and earnings from three models: (i) Benchmark, (ii) without transitory shocks and (iii) without transitory shocks and correlation between productivity in household members.

Finally, Table 8 reports the estimates of the spousal response by household assets and age for the model without correlated productivity and transitory shocks. The model clearly shows that the effects are stronger among households who have limited ability to insure themselves – consistent with empirical findings. Households with limited financial assets rely more on spousal labor supply to insure against income fluctuations versus households with an ample

financial cushion. The model also produces stronger responses for young households, again, consistent with the empirical results. The age effects are partially explained by younger households having fewer financial assets and partially by a longer horizon of capital accumulation which encourages even more the participation of spouses relative to older households.

Table 8: Heterogeneity in Spousal Participation Response: Model vs. Data

	Model	Data
Household's Wealth		
$<$ 1 month's earnings	-0.25	-0.20
3-6 months'	-0.24	-0.16
$>$ 6 months'	-0.03	-0.10
Household's Age		
30–40	-0.11	-0.26
41–50	-0.04	-0.12
51–60	-0.02	-0.01

Notes: Estimates of regression 14 using model-generated data. The model is without correlation of productivity between members and transitory shocks and is compared to the IV estimates from the data.

4.5 Consumption Elasticity

Blundell, Pistaferri, and Saporta-Eksten (2016) measure the degree of insurance against income shocks by (one minus) the pass through of income shocks to household consumption. We also estimate the pass-through in our model by running the same regression as follows. In fact, this specification is almost identical to the consumption function presented in (3) except that we further distinguish between persistent and transitory productivity:

$$\Delta c_{i,t} = \eta_0 + \eta_x \Delta x_{j,i,t} + \eta_\varepsilon \Delta \varepsilon_{j,i,t} + \nu_{i,t} \tag{15}$$

where $c_{i,t}$ is household consumption, $x_{j,i,t}$ and $\varepsilon_{j,i,t}$ is the persistent and transitory productivity for worker $j = p, s$, respectively. For the primary worker, we obtain the pass through of $\widehat{\eta_x} = 0.23$, i.e., 23 percent of permanent income shocks pass through to consumption and a $\widehat{\eta_\varepsilon} = 0.01$, i.e., transitory shocks are nearly perfectly insured. For spouses, we obtain the pass through of $\widehat{\eta_x} = 0.18$, i.e., 18 percent of permanent income shocks pass through

to consumption and again a nearly zero coefficient for transitory shocks. Using data from the U.S. Blundell, Pistaferri, and Saporta-Eksten (2016) estimate that 35 percent of primary permanent shocks pass through consumption and 19 percent of female permanent shocks pass through consumption.

5 Transfer Policy against Income Shock

What is an effective government transfer policy when a certain part of the economy suffers an income loss? In this section, we introduce an unexpected negative productivity shock to primary workers for μ fraction of households. This corresponds to a shock in an industry where the dominant workforce is male workers (e.g., industries such as construction or manufacturing).[16] An effective policy should support consumption, especially for low-income groups, and at the same time preserve households' willingness to use their own labor supply as a means of self-insurance. To this end, we follow the taxonomy of Guner, Kaygusuz, and Ventura (2020) and divide government transfer programs based on the eligibility of transfer payment (universal or targeted).

1. Universal Unconditional Transfers: The government gives transfers to all households regardless of whether they were hit by the income shock.

2. Universal Conditional Transfers: Transfers to all households, but the amount differs by household employment status in order to encourage labor-market participation (i.e. larger transfers when both members work).

3. Targeted Unconditional Transfers: The government gives the same amount of transfers only to the households hit by the income shock.

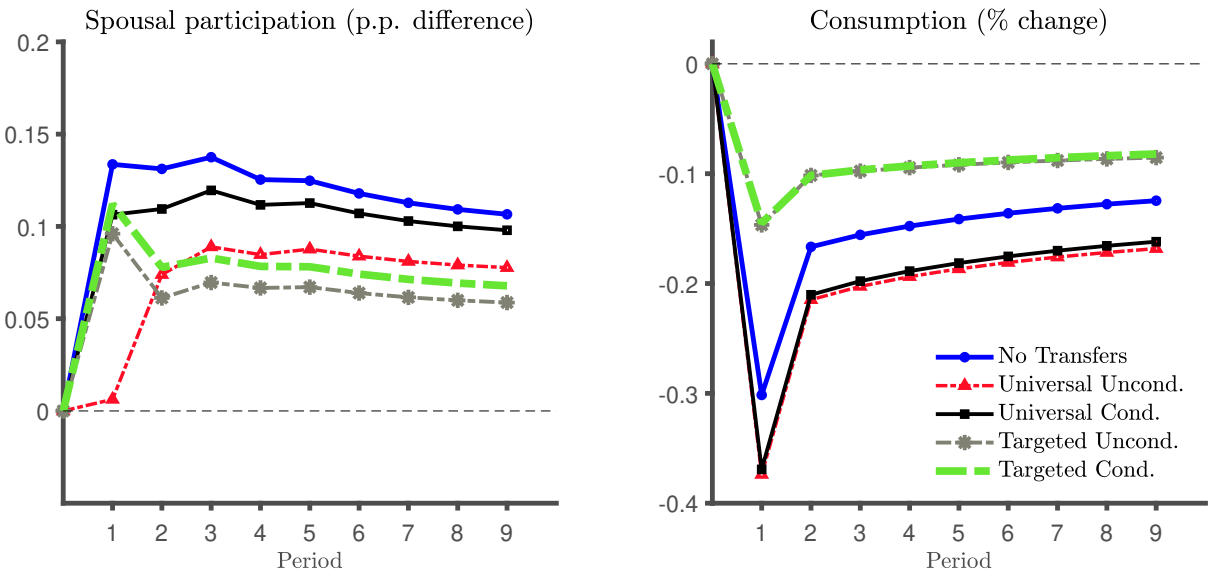
4. Targeted Conditional Transfers: Transfers only to households affected by the income shock. Moreover, the size of the transfer depends on the employment status (i.e. larger transfer for dual earners).

5.1 Productivity Shocks to Primary Workers

Suppose that the primary worker's productivity (x_p), unexpectedly declines by 10 percent for μ share of households in the economy for one year. The government decides to increase the total transfer expenditure (ΔT) by 10 percent, i.e., the increased total transfer payment is the same across four policies. The extra spending is financed by adjusting the labor income tax τ.

[16] We study productivity shocks that affects both members in the next section.

Figure 3: Household Response to A Primary Worker Productivity Decline



Notes: Change in household labor supply and consumption of impacted households from a one-time unexpected shock in primary worker's productivity. We consider the effects when the government does not provide additional transfers and of handing out transfers under four alternative policies.

Under Policy 1, Universal Unconditional Transfers, every household receives 10 percent more transfers. Under Policy 2, Universal Conditional Transfers, in order to encourage labor market participation, dual-earner households (EE) in the economy receive the additional transfer payment twice as much as single-earner households (EN). Under Policy 3, Targeted Unconditional Transfer, only those households affected by the negative productivity shock (μ) receive the extra transfer. Finally, in Policy 4, Targeted Conditional Transfer, while only fraction μ households are eligible, dual income households receive the extra payment twice as much as the rest of the recipients. We also consider a case where the government does not intervene and the transfers level remains as in the steady state ($\Delta T = 0$). As a baseline, we set the share of affected households $\mu = 0.05$.

Figure 3 shows the impulse responses of spousal labor supply and consumption of households that experience a decline in primary productivity under different policy scenarios. Without any government intervention ("No Transfers"), in response to a 10 percent decline in primary productivity, spousal participation increases by 0.14 percent. Participation remains high even after the shock for two reasons. First, due to entry costs, some households do enter immediately to the labor market but wait for a favorable idiosyncratic productivity draw. Second, the impact on spousal participation is long-lived because households decrease their financial assets to mitigate the decline in consumption. Thus, it takes many years for the

participation rate to return to the steady state. This is because the spousal labor supply is operated through income effects. The right panel of Figure 3 shows that without any government intervention, in response to a 10 percent decline in primary productivity, consumption decreases by around 0.30 percent. Consumption falls be a relatively small amount because households can insure against the one-time income shock by using their financial assets and increasing spousal labor supply.

Table 9 shows the initial impact (period 1) for three groups of households which experience the negative productivity shock: the average, poor (bottom half), and wealthy (top half). There is vast heterogeneity in the responses of labor supply, earnings, and consumption based on financial assets. In poor households, primary participation and earnings decrease by -0.14 percentage points and -4.29 percent, respectively. Spousal participation increases by 0.35 percentage points and earnings by 0.82 percent. However, for wealthy households, despite a significant loss in primary earnings (-4.88), the participation of both primary and spousal members hardly moves because the primary worker has high productivity (well above the reservation level) and both household members are well insured by large holdings of assets.

Among the four government programs, the least effective policy is to give universal and unconditional transfers (Policy 1). To finance this program, the government has to increase the income tax rate by 2.8 percentage points in the first year. Even for poor households participation remains broadly same. Given the countercyclical subsidy, households' willingness to self-insure through the labor supply of secondary workers diminishes (Cullen and Gruber, 2000; Birinci, 2024). Furthermore, the higher tax rate further distorts the labor supply decision. As a result, average consumption falls even more than in the case without government intervention.

In order to maintain employment, the government can make the transfer payment conditional on the employment status (e.g., similarly to the Earned Income Tax Credit). Under Policy 2, households receive an extra payment conditional on the employment of the primary worker. The required tax increase is 2.28 percentage points increase (smaller than 2.81 under Policy 1). The policy is successful in preserving the incentives of households to insure using spousal labor supply. On average, spousal participation increases by 0.10 percentage points and it is low-wealth households account for this increase. However, aggregate consumption is only slightly affected and declines by 0.36 percentage points (slightly less than -0.35 under Policy 1).

Under Policy 3, the additional transfers are targeted to impacted households. As a result, the required tax increases by only 0.05 percentage points. As the distortion from the tax increase is small, the aggregate consumption declines by 0.14 percentage points.

Finally, the policy that is targeted and conditional on employment status is the most ef-

Table 9: Effects of Transfers: Negative Productivity Shock to Primary Only

	No Intervention	Policy 1 Universal	Policy 2 Universal Conditional	Policy 3 Targeted	Policy 4 Targeted Conditional
Tax Rate (τ)	+0.05	+2.81	+2.28	+0.05	+0.04
All Households					
Participation, p	−0.07	−0.13	+0.00	−0.09	+0.00
Participation, s	+0.13	+0.00	+0.10	+0.09	+0.11
Earnings, p	−4.43	−6.37	−6.30	−4.65	−4.60
Earnings, s	+0.51	−1.58	−1.47	+0.22	+0.24
Hours, p	−4.36	−6.32	−6.21	−4.67	−4.60
Hours, s	+0.70	−1.44	−1.29	−0.28	−0.31
Consumption	−0.30	−0.37	−0.36	−0.14	−0.14
Poor					
Participation, p	−0.14	−0.28	+0.00	−0.18	+0.00
Participation, s	+0.35	+0.01	+0.28	+0.25	+0.29
Earnings, p	−4.29	−6.27	−6.11	−4.69	−4.58
Earnings, s	+0.82	−1.39	−1.17	−0.32	−0.37
Consumption	−0.44	−0.53	−0.51	−0.16	−0.16
Wealthy					
Participation, p	−0.00	−0.00	−0.00	−0.00	−0.00
Participation, s	+0.00	−0.00	−0.00	+0.00	+0.00
Earnings, p	−4.53	−6.44	−6.44	−4.62	−4.62
Earnings, s	+0.22	−1.75	−1.75	+0.12	+0.13
Consumption	−0.18	−0.24	−0.24	−0.12	−0.12

Note: Effect of a negative productivity shock on selected variables without government intervention and under four transfer policies. Poor and wealthy households are above and below the median wealth at the time of the shock, respectively. All variables are expressed in percentage change from steady state value. Participation and the tax rate are reported as a change in percentage points.

fective one among the four programs we considered. The required tax is only 0.04 percentage points. The primary participation does not decline in spite of the shock. Spousal participation rises almost as much as in the case without additional transfers at all. Therefore, this policy can both support consumption and preserve labor supply incentives. The aggregate consumption falls by 0.14 as in Policy 3.

We report the welfare associated with each transfer scheme. Welfare is in terms of con-

sumption equivalent. We ask how much consumption would a newborn worker (age 21) be willing to sacrifice to be born in an economy where the shock has disappeared and the economy has returned to its steady state instead of being born in the economy that has just experienced the productivity shock. When the government does not intervene (no transfers), impacted newborn households have a welfare loss equal to 0.10 percent. Due to a higher tax rate to balance the government budget (equal to a 0.05 percentage point rise), non-impacted households also face a welfare loss of 0.01 percent.

When we give universal and unconditional transfers the impacted newborn households have a welfare loss equal to 0.08 percent. The government intervention helps reduce the welfare loss but the reduction is not large. This happens because the large increase in tax rates by 2.8 percentage points reduce household income. The non-impacted new born households receive transfers but at the same time pay a higher tax rate. Thus, they experience a very small positive welfare gain equal to 0.001 percent.

With universal and conditional transfers the impacted newborn households have a welfare loss also equal to 0.10 percent and the non-impacted new born households a small positive welfare gain equal to of 0.004 percent. Thus, under our parameterization, universal conditional transfers are not preferable universal unconditional in terms of welfare even though they mitigate slightly more consumption loss and promote more employment. The benefits of smaller decline in consumption are canceled out by the cost of working more.

With targeted and unconditional transfers the impacted newborn households have a reduced welfare loss equal to 0.01 percent. Transfers are more concentrated in this case; only impacted households receive the transfer so the tax rate rises only slightly by 0.05 percentage points. The non-impacted new born only pay for the higher tax rate without collecting the transfer so their welfare is lower by 0.003 percent.

When transfers are targeted and conditional the impacted newborn households have a reduced welfare loss equal to 0.03 percent. Once more, although consumption declines less in this case, households are also working more to sustain the larger consumption level. Thus, the impacted newborn is not as well off as in the case with targeted and unconditional transfers. The same is true for non-impacted newborns whose welfare is lower by 0.01 percent, i.e., slightly less than in the targeted and unconditional policy.

In general, we find that targeted policies far outperform universal policies in terms of efficiency gains and welfare. Universal transfer policies require such a large increase in tax rates that households may experience even lower consumption levels compared to a no-transfer policy. Conditional policies rely on providing incentives for households to stay in the labor market. We find that they slightly improve efficiency gains, but not always welfare.

5.2 Income Shocks to Both Earners

We consider a shock that affects both members: for a $\mu = 0.05$ fraction of households, productivity (x_p and x_s) declines by 10 percent for one year. As before, the government increases the total transfers payment by 10 percent and adjusts the income tax rate τ to finance the additional expenditure. In this case both primary and spousal participation, earnings, and hours decline. The decline in spousal participation – especially among poor households – is larger because the average productivity of secondary workers is low (thus closer to the reservation level). Movements in the participation rate arise from less wealthy households. Once more, Policy 4 (targeted conditional transfers) is the most effective policy: the average consumption (which is a good proxy for welfare) falls the least (-0.29 percentage points). The most ineffective one is Policy 1 (universal unconditional transfers) where average consumption declines by -0.42 (even higher than the case without intervention). This exercise shows how correlated income shocks can lead to a muted spousal labor supply in response to primary income. In single-earner households the spouse does not join the labor market while in dual earner households the spouse may well leave the market due to the negative productivity shock.

6 Conclusion

This paper examines the role of spousal labor supply as an insurance mechanism against primary income fluctuations. Using administrative data from Norway, we identify primary income shocks that are unanticipated, persistent, and uncorrelated with the spousal's productivity and estimate the spousal participation and earnings response. Our findings highlight a significant added worker effect, particularly for households with limited insurance options. We develop a structural model calibrated the Norwegian labor market and welfare system and confirm that it replicates the spousal labor supply responses we document in the data. We use the model to quantify the heterogeneous effects of various transfer programs and evaluate their effectiveness to jointly support consumption and promote labor supply.

References

Aiyagari, S. R. (1994). "Uninsured idiosyncratic risk and aggregate saving". *Quarterly Journal of Economics*, *109*(3), 659-684.

Albanesi, S., and Prados, M. J. (2022). Slowing women's labor force participation: The role of income inequality. *Working Paper*.

Autor, D., Kostol, A., Mogstad, M., and Setzler, B. (2019). Disability Benefits, Consumption Insurance, and Household Labor Supply. *American Economic Review*, *109*(7), 2613–2654.

Bacher, A., Grübener, P., and Nord, L. (2024). Joint search over the life cycle. *Journal of Monetary Economics*, 103696.

Birinci, S. (2024). *Spousal Labor Supply Response to Job Displacement and Implications for Optimal Transfers* (Working Paper). Federal Reserve Bank of St. Louis.

Blundell, R., Costa Dias, M., Meghir, C., and Shaw, J. (2016). Female Labor Supply, Human Capital Accumulation, and Welfare Reform. *Econometrica*, *84*(5), 1705–1753.

Blundell, R., Pistaferri, L., and Preston, I. (2008). Consumption Inequality and Partial Insurance. *American Economic Review*, *98*(5), 1887-1921.

Blundell, R., Pistaferri, L., and Saporta-Eksten, I. (2016). Consumption Inequality and Family Labor Supply. *American Economic Review*, *106*(2), 387-435.

Borella, M., De Nardi, M., and Yang, F. (2023). Are Marriage-Related Taxes and Social Security Benefits Holding Back Female Labour Supply? *Review of Economic Studies*, *90*(1), 102-131.

Casella, S. (2022). *Women's Labor Force Participation and the Business Cycle* (Working Paper).

Chang, Y., Hong, J., Karabarbounis, M., Wang, Y., and Zhang, T. (2022). Income Volatility and Portfolio Choices. *Review of Economic Dynamics*, *44*, 65–90.

Chang, Y., and Kim, S.-B. (2006). From individual to aggregate labor supply: A quantitative analysis based on a heterogeneous agent macroeconomy. *International Economic Review*, *47*(1), 1–27.

Cullen, J., and Gruber, J. (2000). Does unemployment insurance crowd out spousal labor supply? *Journal of Labor Econoics*, *18*(3), 546–572.

Doepke, M., and Tertilt, M. (2016). Families in macroeconomics. *Handbook of Macroeconomics*, *2*.

Ellieroth, K. (2023). Spousal insurance, precautionary labor supply, and the business cycle. *Working Paper*.

Erosa, A., Fuster, L., and Kambourov, G. (2016). Towards a Micro-Founded Theory of Aggregate Labour Supply. *Review of Economic Studies*, *83*(3), 1001–1039.

Fagereng, A., Guiso, L., and Pistaferri, L. (2018). Portfolio choices, firm shocks, and uninsurable wage risk. *Review of Economic Studies*, *85*(1), 396-436.

Guner, N., Kaygusuz, R., and Ventura, G. (2020). Child Related Transfers, Household Labor Supply, and Welfare. *Review of Economic Studies*, *87*(5), 2290–2321.

Guner, N., Kaygusuz, R., and Ventura, G. (2023). Rethinking the Welfare State. *Econometrica*, *91*(6), 2261–2294.

Guner, N., Kulikova, Y., and Valladares-Esteban, A. (2024). *Does the Added Worker Effect Matter?* (Working Paper). HCEO, The University of Chicago.

Guvenen, F. (2007). Learning your earning: Are labor income shocks really very persistent? *American Economic Review*, *97*(3), 687-712.

Guvenen, F., Karahan, F., Ozkan, S., and Song, J. (2021). What do data on millions of U.S. workers reveal about life-cycle earnings risk. *Econometrica*, *89*(5), 2015–2557.

Guvenen, F., and Smith, T. (2014). Inferring labor income risk and partial insurance from economic choices. *Econometrica*, *82*(6), 2085-2129.

Halla, M., Schmieder, J., and Weber, A. (2020). Job displacement, family dynamics, and spousal labor supply. *American Economic Journal: Applied Economics*, *12*(4), 253-287.

Heathcote, J., Storesletten, K., and Violante, G. L. (2010). "The macroeconomic implications of rising wage inequality in the united states". *Journal of Political Economy*, *118*(4), 681–722.

Kaplan, G., and Violante, G. (2010). How Much Consumption Insurance Beyond Self-Insurance? *American Economic Journal: Macroeconomics*, *2*(4), 53-87.

Kaplan, G., and Violante, G. L. (2014). "A model of the consumption responses to fiscal stimulus payments". *Econometrica*, *82*(4), 1199-1239.

Keane, M. (2011). Labor supply and taxes: A survey. *Journal of Economic Literature*, *49*(4), 961-1075.

Klovland, J. (2004). Bond markets and bond yields in norway 1819-2003. In O. Eitrheim, J. Klovland, and J. Qvigstad (Eds.), *Historical monetary statistics for norway 1819-2003.* Norges Bank.

Krueger, D., and Wu, C. (2021). Consumption insurance against wage risk: Family labor supply and optimal progressive income taxation. *American Economic Journal: Macroeconomics*, *13*(1), 79-113.

Lundberg, S. (1985). The added worker effect. *Journal of Labor Economics*, *3*(1), 11–37.

McKay, A., and Reis, R. (2016). The Role of Automatic Stabilizers in the U.S. Business Cycle. *Econometrica*, *84*(1), 141–194.

Mustre-del Rio, J. (2015). Wealth and labor supply heterogeneity. *Review of Economic Dynamics*, *18*(3), 619-634.

Primiceri, G., and van Rens, T. (2009). Heterogeneous life-cycle profiles, income risk and consumption inequality. *Journal of Monetary Economics*, *56*(1), 20-39.

Stephens, M. (2002). Worker Displacement and the Added Worker Effect. *Journal of Labor Economics*, *20*(3), 406-422.

Appendix: For Online Publication

A Static model of household labor supply

In this appendix we solve for the household's static model of labor supply. Assume there is a household with the following separable preferences.

$$U = \frac{c^{1-\sigma}}{1-\sigma} - \frac{h_p^{1+\frac{1}{\gamma}}}{1+\frac{1}{\gamma}} - \frac{h_s^{1+\frac{1}{\gamma}}}{1+\frac{1}{\gamma}}$$

where c is household consumption, h_p is hours by the primary worker, and h_s is the hours by the spouse. The household consumption is equal to the primary worker's earnings, the spouse's earnings and non labor income N.

$$c = wx_p h_p + wx_s h_s + n$$

It is convenient to define the labor share of primary/secondary income:

$$S_i = \frac{x_i h_i}{x_p h_p + x_s h_s + N} \quad \text{with } i = \{p, s\}$$

The first-order condition for spousal labor supply is:

$$\frac{1}{\gamma} \ln h_s = -\sigma \ln \left[wx_p h_p + wx_s h_s + N \right] + \ln x_s$$

By total differentiation of the first-order condition with respect to (h_s,x_s) we get:

$$\frac{1}{\gamma}\frac{1}{h_s} dh_s = -\sigma \frac{x_s}{x_p h_p + x_s h_s + N} dh_s - \sigma \frac{h_s}{x_p h_p + x_s h_s + N} dx_s + \frac{1}{x_s} dx_s$$

After some slight algebra manipulation we get

$$\frac{dh_s}{dx_s}\frac{x_s}{h_s}\left[\frac{1}{\gamma}\frac{1}{h_s} + \sigma \frac{x_s}{x_p h_p + x_s h_s + N}\right] = -\sigma \frac{x_s}{x_p h_p + x_s h_s + N} + \frac{1}{h_s}$$

and by multiplying by h_s in both sides we derive:

$$\frac{dh_s}{dx_s}\frac{x_s}{h_s}\left[\frac{1}{\gamma} + \sigma \underbrace{\frac{x_s h_s}{x_p h_p + x_s h_s + N}}_{S_s}\right] = -\sigma \underbrace{\frac{x_s h_s}{x_p h_p + x_s h_s + N}}_{S_s} + 1$$

which takes us to the Marshallian labor supply elasticity:

$$\frac{dh_s}{dx_s}\frac{x_s}{h_s} = \frac{1 - \sigma S_s}{1/\gamma + \sigma S_s}$$

To estimate the elasticity of spousal labor supply with respect to primary income we take the total differentiation of the first order condition with respect to $(h_s, x_p h_p)$:

$$\frac{1}{\gamma}\frac{1}{h_s}dh_s = -\sigma\frac{x_s}{x_p h_p + x_s h_s + N}dh_s - \sigma\frac{1}{x_p h_p + x_s h_s + N}d[x_p h_p]$$

After some algebra this takes us to:

$$\frac{1}{\gamma}\frac{dh_s}{d[x_p h_p]}\frac{[x_p h_p]}{h_s} = -\sigma\underbrace{\frac{x_s h_s}{x_p h_p + x_s h_s + N}}_{S_s}\frac{dh_s}{d[x_p h_p]}\frac{[x_p h_p]}{h_s} - \sigma\underbrace{\frac{[x_p h_p]}{x_p h_p + x_s h_s + N}}_{S_p}$$

and

$$\frac{dh_s}{d[x_p h_p]}\frac{[x_p h_p]}{h_s}\left[\frac{1}{\gamma} + \sigma S_s\right] = -\sigma S_p.$$

The spousal labor supply elasticity is thus:

$$\{\frac{dh_s}{d[x_p h_p]}\frac{[x_p h_p]}{h_s} = \frac{-\sigma S_p}{\frac{1}{\gamma} + \sigma S_s}$$

We now consider the case where the primary earnings $wx_p h_p$ is correlated with the spouse's productivity. Specifically, we assume that $x_s = f(wx_p h_p)$. The first-order condition of spousal labor supply remains the same. But the total differentiation of the first order condition with respect to $(h_s, x_p h_p)$ gives the following expression

$$\begin{aligned}\frac{1}{\gamma}\frac{1}{h_s}dh_s = &-\sigma\frac{x_s}{x_p h_p + x_s h_s + N}dh_s - \sigma\frac{1}{x_p h_p + x_s h_s + N}d[x_p h_p] \\ &-\sigma\frac{h_s}{x_p h_p + x_s h_s + N}f' d[x_p h_p] + \frac{1}{x_s}f' d[x_p h_p]\end{aligned}$$

After some algebra we can write this as:

$$\frac{dh_s}{d[x_p h_p]}\frac{[x_p h_p]}{h_s}\left[\frac{1}{\gamma} + \sigma S_s\right] = -\sigma S_p - \sigma\frac{h_s x_s}{x_p h_p + x_s h_s + N}f'\frac{[x_p h_p]}{x_s} + \frac{x_p h_p}{x_s}f'$$

or equivalently:

$$\frac{dh_s}{d[x_p h_p]}\frac{[x_p h_p]}{h_s}\left[\frac{1}{\gamma} + \sigma S_s\right] = -\sigma S_p + [1 - \sigma S_s]\frac{x_p h_p}{x_s}f'$$

Note that $\frac{x_p h_p}{x_s}f'$ is the elasticity of spousal productivity wrt to primary earnings, so:

$$\frac{dh_s}{d[x_p h_p]}\frac{[x_p h_p]}{h_s} = \underbrace{\frac{-\sigma S_p}{\frac{1}{\gamma} + \sigma S_s}}_{\text{income elasticity}} + \varepsilon_{x_s, x_p h_p} \times \underbrace{\frac{1 - \sigma S_s}{\frac{1}{\gamma} + \sigma S_s}}_{\text{Marshallian elast.}}$$

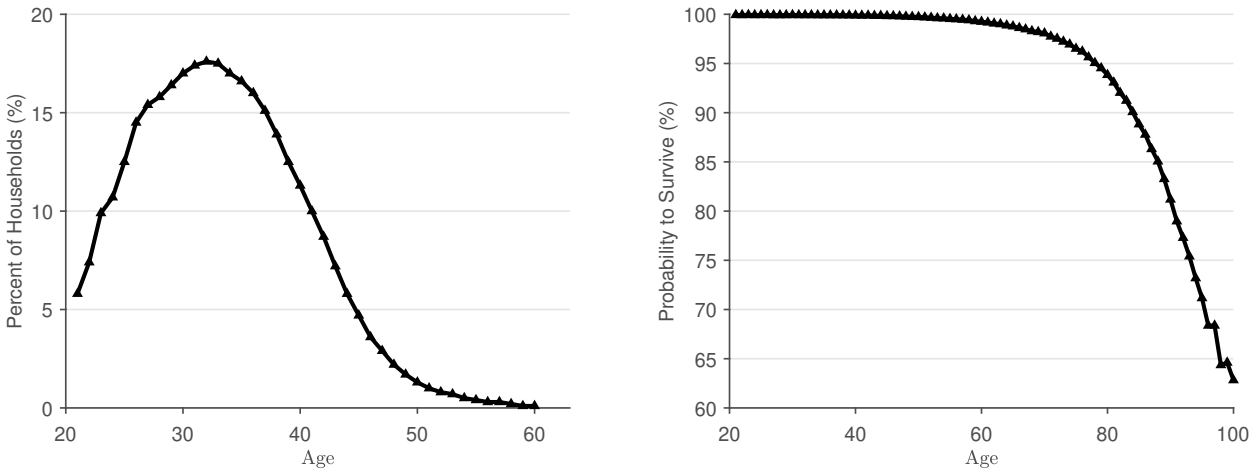
which is the equation we have in the main text.

B Additional Empirical Analysis

B.1 Number of Kids by Age

Our model includes age-dependent expenses for child care. To calibrate these expenses we need to measure the number of young children by age group. In the administrative data we observe the age of the youngest child of each family. Based on this measure we compute the share of households that have a child younger than the age of 5. Figure B-1 shows the percentage of married households that have a child younger than 5 years old in Norway. The x-axis is represents the age of the head of the household. The share is hump-shaped. At the peak, which occurs at the age of 32, 17.6 percent of Norwegian households have a child younger than the age of 5. After that age the share is monotonically declining.

Figure B-1: Share of Households with a Child < 5 Years Old and Survival Probability



Notes: The left panel shows the share of households with a child less than age of 5 by household age. The right panel shows the survival probability by age.

C Additional Quantitative Analysis

C.1 Calibration

For every parameter in the model we describe its value and how we choose the value. In Table C-1 describes one set of parameters is set externally while the other set is internally calibrated to match salient features of the Norwegian economy.

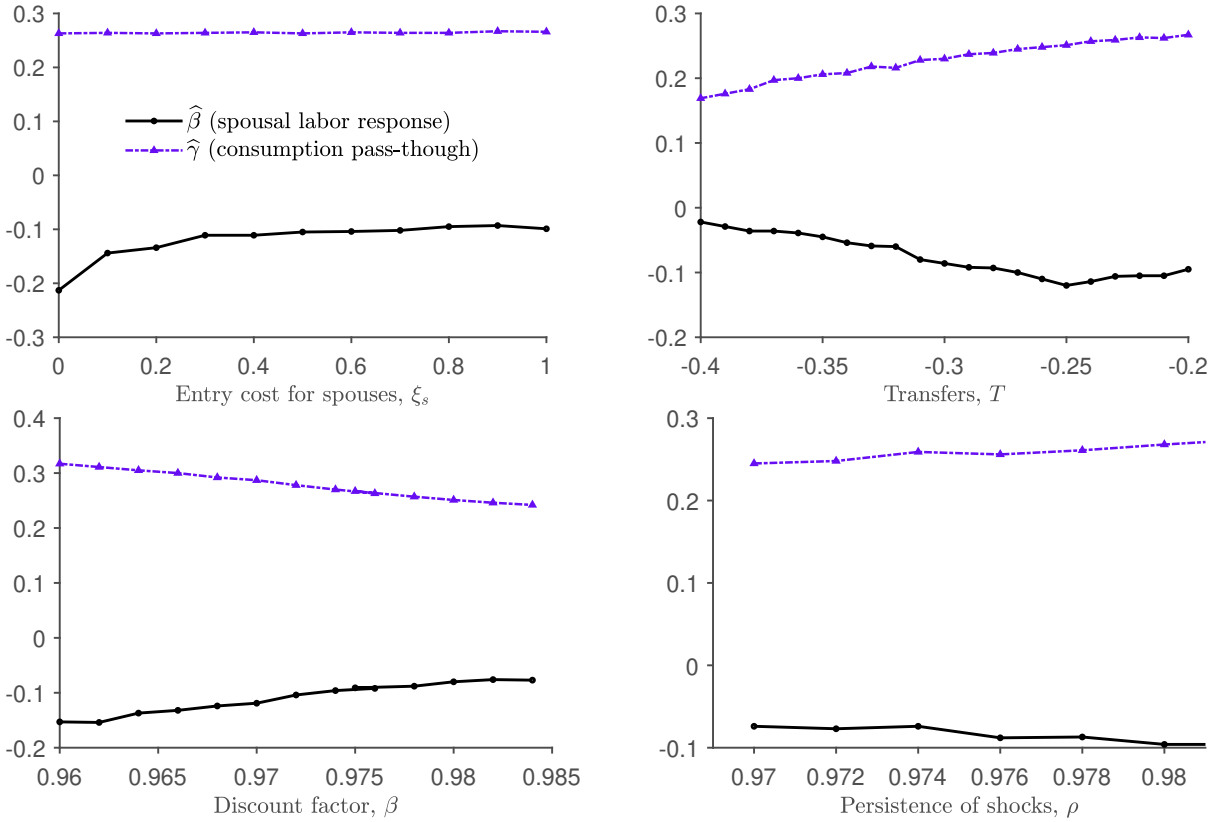
Table C-1: Parameter Values

Description	Variable	Value	Reference / Target
Externally Set Parameters			
Life cycle	J	80	
Retirement age	j_R	45	
Survival probabilities	$\{s_j\}$	–	Statistics Norway
Number of children w/ age < 5	$\{n_{5,j}\}$	–	Statistics Norway
Daycare cost	$\{\zeta_5\}$	\$2,500	Statistics Norway
Risk aversion	σ	2	–
Frisch elasticity of labor	γ	0.5	–
Interest rate	r	3.2%	Klovland (2004)
Persistence of shocks	ρ_x	0.955	Table 4
Var of innovation to persistent shocks	σ^2_ν	0.025	Table 4
Variance of transitory shocks	σ^2_ε	0.04	Blundell et. al (2016)
Proportional tax rate	τ	0.32	Statistics Norway
Social security benefit	b	0.40	Replacement rate
Internally Fitted Parameters			
Discount factor	β	0.967	Average asset-to-income
Separation rate, primary	λ_p	1%	Average $N \to E$, primary
Separation rate, spouse	λ_s	0.5%	Average $N \to E$, spouse
Initial human capital, spouse	χ^s_0	0.75	Gender-wage premium
Human capital growth	χ_1	0.10	Wages growth (25-60)
Covariance of persistent shocks	σ_{ps}	9.0	Correlation in productivity
Variance of initial productivity	$\sigma^2_{x(1)}$	0.10	Average participation (25-40)
Entry cost, primary	ξ_p	4.2	Average participation, primary
Entry cost, spouse	ξ_s	1.6	Average participation, spouse
Borrowing limit	$\underline{a}$	0.05	Average credit card debt
Disutility from work, primary	ψ_p	9.0	Average hours, primary
Disutility from work, spouse	ψ_s	8.7	Average hours, spouse
Government transfers	T	0.16	Transfers-income ratio

C.2 Robustness

We examine how two elasticities (spousal labor supply, $\hat{\beta}$ and consumption pass-through, $\hat{\eta_x}$) change with respect to four parameters: labor market entry cost, the size of government transfer, time discount factor, and persistence of productivity. In all exercises, we keep all the other parameters unchanged.

Figure C-2: Comparative Statics



Halla, Schmieder, and Weber (2020) argues that social gender norms make it harder for women to join the workforce. We vary the entry cost of labor market participation for spouses ξ_s in the upper left panel of Figure C-2. With no entry costs ($\xi_s = 0$), $\hat{\beta} = -0.21$ is almost twice as large as in the benchmark case. As the entry cost increases to $\xi_s = 1$, the participation elasticity decreases by half to $\hat{\beta} = -0.12$ The consumption elasticity with respect to persistent productivity η_x decreases slightly but remains more or less the same.

According to Cullen and Gruber (2000) and Birinci (2024), government transfers provide additional insurance to households and thus mitigate the spousal labor response to primary income shocks. The upper right panel of Figure C-2 shows that when T by 40% (compared to the reference), β approaches zero and η_x decreases to 0.16.

Next, we consider changing the time discount factor in the lower left panel of Figure C-2. As households become more patient, both elasticities increase. In the last panel, we increase the persistence of productivity shocks. As the shocks become persistent, both elasticities increase.

C.3 Frisch Elasticity of Labor Supply

While our primary interest is the added worker effect (the response of spousal labor in the extensive margin), we also estimate the cross-elasticity in the intensive margin in the model. First, we estimate the measure of the Frisch elasticity of labor supply for the primary earner and the spouse. We simulate a one-time, unanticipated decline in primary and spousal wage by 10% and calculate the percentage change in household labor supply. Since the wage change is unanticipated and lasts for one period (i.e., wealth remains constant), the estimated elasticities can be interpreted as Frisch elasticities of labor supply. Frisch elasticities are more informative about the household responses to temporary changes in income shocks or policy like the ones we analyze in Section 6. Table C-2 reports the Frisch labor supply elasticity to the changes in the own wage and to the changes in the partner's wage (cross-elasticity).

Table C-2: Frisch Elasticities of Labor Supply

Labor supply	Primary earner		Spouse	
	own wage	cross wage	own wage	cross wage
Age 25	0.49	–0.02	0.57	–0.07
Age 30	0.45	–0.02	0.48	–0.04
Age 40	0.43	–0.03	0.44	–0.05
Age 50	0.45	–0.01	0.46	–0.03
Aggregate	0.46	–0.01	0.50	–0.03

Notes: Spousal labor and earnings responses to primary income shocks in model and the data. Model estimates are based on same regression as in the data. We report estimates by wealth as a share of earnings and age.

As shown in Table C-2, the primary Frisch elasticity of labor supply is on average 0.46 and the spousal is on average 0.50. Average estimates depend heavily on our choice of the intertemporal elasticity of the labor supply $\gamma = 0.5$. The heterogeneity in the elasticity of labor supply arises endogenously in the model through differences to participate in the labor market. Young households that possess some level of wealth and have a long time horizon are willing to move temporarily to non-employment. As households reach their middle age, the elasticities decline since households need to build up life cycle savings. Elasticities increase closer to retirement when savings are more abundant and returns to human capital are lower. Table C-2 also reports the cross-Frisch elasticity of labor supply for the primary and the spouse. The cross-Frisch elasticity for the spouse (with respect to a decline in primary earnings) is -0.03 and is larger at younger ages. The cross-Frisch elasticity

for the primary (with respect to a decline in spousal earnings) is uniformly small across ages, since spousal income is typically a small share of household income.