The Transmission of Income Shocks Inside the Household

Guillermo Carranza* Jin Cao

Yale University Norges Bank

November 10, 2026

Please click here for the most recent draft.

Abstract

The transmission of income shocks to consumption is typically estimated at the household level because within-household allocations are difficult to observe. Yet household-level estimates can conceal differences in individual exposure when spouses share risk imperfectly. Using spouse-level Norwegian administrative expenditure data and plausibly exogenous firm-driven income variation, we estimate the pass-through elasticity of spousal income shocks to household, personal, and joint expenditures. Household consumption responds more to women's permanent shocks than to men's, indicating imperfect risk-sharing. These income shocks pass strongly into the affected spouse's personal consumption, with substantially higher pass-through for women than for men. At the same time, the effect of one spouse's shock on the other's personal consumption is very small. Joint expenditures, including groceries and utilities, exhibit greater pooling across spouses' shocks. These responses are difficult to reconcile with standard models in which spouses pool all resources, and survey data we collect show financial separation is widespread among Norwegian couples. We develop a limited-commitment model with separate budgets that can reproduce our findings. Relative to full risk-sharing, our model implies lower insurance gains from marriage and higher welfare gains from public insurance.

***Job Market Paper.** Corresponding author: guillermo.carranzajordan@yale.edu. We thank Joseph Altonji, Orazio Attanasio, Cormac O'Dea, and Seth Zimmerman for invaluable feedback and support. We also thank Jaime Arellano-Bover, José Casco, Peter Ganong, Sevin Kaytan, Costas Meghir, Pascal Noel, Marius Ring, Matthew Schwartzmann, and Eric Solomon for comments, as well as seminar participants at Norges Bank, Banco de México and the Yale Labor/Public workshop. We thank the Cowles Foundation for Research in Economics for financial support. The Norwegian survey received a Yale IRB exemption (Protocol ID 2000042886). All errors are our own.

1 Introduction

The ability to maintain consumption when income fluctuates determines how well people are insured against income risk. Measuring the magnitude of this insurance is central to understanding the welfare costs of income volatility, the effectiveness of private insurance, and the adequacy of public insurance. Yet nearly all of the existing literature approaches this question from a household perspective, measuring how household consumption responds to changes in income (Attanasio and Davis, 1996; Blundell et al., 2008; Guvenen and Smith, 2014; Ganong et al., 2025). While these estimates characterize household-level insurance, they may conceal differences in each spouse's exposure to income risk when spouses share risk imperfectly. Thus, two households may appear equally insured while providing different levels of protection to each spouse.

Past empirical research shows that intrahousehold consumption allocations depend on how income is distributed across spouses (Browning et al., 1994; Attanasio and Lechene, 2002; Chiappori and Meghir, 2015) and that spouses often do not achieve perfect risk sharing (Duflo and Udry, 2004; Robinson, 2012). These findings suggest that the allocation of consumption insurance within the household is a first-order question for understanding the welfare impact of income risk. However, tracking how consumption is allocated within households over time is difficult: existing evidence relies almost exclusively on small-scale surveys of individual consumption. Longitudinal administrative expenditure data linking individual consumption to plausibly exogenous income shocks are especially scarce. As a result, our understanding of how households transmit and share income risk remains incomplete.

This paper measures the transmission of risk within households using unique Norwegian administrative panel data on individual-level spending. We ask three related questions. First, when one spouse experiences an income shock, how much does aggregate household consumption respond? Second, looking within the household, how does this response differ between personal expenditures, which primarily benefit one spouse, and joint expenditures, which are shared by the household? Third, what models of household risk sharing can rationalize the resulting patterns? Our main finding is that the provision of insurance differs substantially between joint and personal expenditures. Couples pool income shocks through joint expenditures, while personal consumption remains closely tied to the income of the spouse who consumes it. This incomplete risk sharing is particularly pronounced for women: women's personal consumption responds considerably more to their own income variation than men's personal consumption responds to theirs.

Our main administrative transaction data come from Norway's domestic payment network, cover about 95% of taxpayers and roughly 80% of electronic payments, and can be

linked to administrative registers on income, employment, wealth, and family structure. This national, spouse-linked individual-level dataset allows us to construct measures of household, personal, and joint consumption at a scale and with coverage typically unavailable in surveys or other transaction datasets. Norway is also a particularly useful setting to study household dynamics using expenditure data. Only about one in ten point-of-sale payments is made in cash, and Norway's cash share is one of the lowest in advanced economies (Norges Bank, 2019).

Identifying causal responses to income shocks requires addressing both the persistence and endogeneity of income changes. Income shocks can be permanent or transitory, and consumption should respond differently depending on their persistence (Hall, 1978). At the same time, income and consumption may co-move because of other shocks or reverse causality, making plausibly exogenous variation in income essential. To address these challenges, we first build on Ganong et al. (2025) and instrument each spouse's income growth with firm-driven income shocks, proxied by changes in co-workers' income. We then adapt the decomposition in Blundell et al. (2008) (BPP) to these firm-level shocks and a two-earner household, allowing us to separate permanent from transitory innovations. The approach exploits differences in the autocovariance patterns of permanent and transitory income shocks to identify their effects on household, personal, and joint consumption. We refer to this approach as BPP-IV estimation.

Our analysis yields four main empirical results. First, the source of an income shock matters even for aggregate household consumption responses. Throughout the analysis, we compare equally sized percentage changes in aggregate household income driven by shocks to women's versus men's income. A permanent income change that reduces household income by 10% lowers nondurable consumption by 5.1% if it originates from women's income, compared with 2.9% if it originates from men's income. These responses correspond to pass-through elasticities—the percentage change in consumption associated with a 1% change in household income—of 0.51 and 0.29, respectively. This dependence on the source of the shock is inconsistent with perfect intrahousehold risk sharing.

Second, permanent shocks to a spouse's own income pass strongly into that spouse's personal consumption. We estimate own-income pass-through to personal consumption of 1.44 for women and 0.88 for men, approximately 64% larger for women. These findings point to substantial differences in the transmission of own-income shocks to personal consumption between women and men.

Third, strikingly, cross-spouse pass-through to personal consumption is small. The pass-through of men's permanent income shocks to women's personal consumption is 0.10, while the pass-through of women's permanent income shocks to men's personal consumption is

0.07 and statistically indistinguishable from zero. This large difference between own- and cross-spouse responses points to very limited pooling of personal consumption. Instead, personal consumption remains strongly connected to own permanent income and responds negligibly to the spouse's income.

Fourth, joint expenditures, including utilities, groceries, and other shared goods, respond similarly to permanent income shocks from either spouse. The pass-through of joint expenditures shows much smaller differences across spouses than those of either spouse's personal consumption. Overall, our findings indicate that households appear to pool resources for shared goods while leaving personal consumption exposed to individual income risk, especially for women.

Our personal-consumption results hold after accounting for potential mismeasurement of consumption across household members. We construct an imputed measure of personal consumption that is less sensitive to expenditure spillovers across spouses, using clothing expenditures and pre-union spending relationships. This approach builds on past literature that uses clothing as an assignable good to identify individual resource shares within households (Dunbar et al., 2013; Lechene et al., 2022). We show that clothing expenditure changes little around the start of a partnership, supporting its use as an assignable measure during cohabitation or marriage. Using clothing-imputed consumption, we find the same cross-spousal patterns: own-income pass-through is large and larger for women, while cross-spouse pass-through is small. We also find similar results for personal consumption pass-through among couples without children, reducing concerns that our findings are driven by expenditures on children.

In the second half of the paper, we ask which household models can explain our pass-through results. Our findings help distinguish between two mechanisms that can generate imperfect insurance across spouses but differ in their ability to explain the small cross-spouse responses in personal consumption. The first mechanism is limited commitment. When spouses cannot commit to future allocations, an income innovation can alter a spouse's outside option and trigger renegotiation of consumption allocations (Mazzocco, 2007; Voena, 2015; Lise and Yamada, 2019). A favorable shock can then shift personal consumption toward the spouse receiving it. The second mechanism is budget separation, drawing on the literature on "separate spheres" in household decision-making (Lundberg and Pollak, 1993; Duflo and Udry, 2004). If spouses retain separate budgets rather than a single joint one, additional income for one spouse expands the resources available for their own consumption, but these resources may not be available to the other spouse.

We introduce these frictions sequentially and discipline the models using the estimated pass-through coefficients on personal and joint expenditures. The standard limited-commitment

model with a joint budget, the workhorse model in the dynamic consumption literature, can generate larger personal consumption responses to own than to spousal income shocks. Quantitatively, however, it generates cross-spousal responses that are too large. In its best-fitting parameterization, a male income shock produces a pass-through of 0.77 to women's personal spending, while a female shock produces a pass-through of 0.49 to men's personal spending, several times larger than in the data. Thus, limited commitment alone struggles to rationalize the near-disconnect between one spouse's income and the other's personal consumption that we observe.

We therefore consider budget separation, a mechanism that has received less attention in economic models of household risk sharing than limited commitment. The sociological literature, however, has documented substantial heterogeneity in how couples organize their finances, including systems of full pooling, partial pooling, and separate finances (Pahl, 1995; Vogler et al., 2008). Motivated by these arrangements, we develop a limited-commitment model that allows spouses to maintain individual budgets while sharing expenses on joint consumption in proportion to their incomes. Under separate budgets, personal consumption remains tied to own resources, while joint consumption provides a channel for risk sharing.

We provide direct evidence supporting the budget-separation mechanism using a new survey of Norwegian households that we conducted in August 2026.[1] Only 28% of respondents report fully pooling all income, while 51% report meaningful financial separation from their spouse. Moreover, 87% report either having no access to their partner's personal account or never using it for personal consumption, and around half agree that "the money I earn feels like mine." Turning to how couples share expenses, respondents report contributing to joint expenditures, such as housing and utilities, broadly in proportion to their incomes, while the majority report paying entirely for their own personal consumption. Together, these facts point to a household financial structure in which cooperation over joint expenditures coexists with meaningful control over personal resources.

We find that a limited-commitment model with separate budgets fits our pass-through estimates much more closely, reproducing the small cross-spouse responses in personal consumption and the more similar responses of joint expenditures to shocks from either spouse. Our results have implications for the value of intrahousehold and public insurance against income risk. We begin by calculating the welfare gain from sharing risk across spouses. In our estimated model, the consumption-equivalent insurance gain from marriage relative to autarky is 3.2% of lifetime consumption for women and 3.6% for men, about 14% higher for men than for women. Relative to a pooled full-commitment benchmark, which represents the full potential insurance surplus under perfect risk sharing, limited commitment and separate

[1] We report selected survey questionnaire items in Appendix D.1.

budgets eliminate about 40% of the insurance gains for women and 24% for men. These risk-sharing frictions therefore not only reduce the insurance provided by marriage relative to full risk sharing, but also distribute that insurance unevenly across spouses.

We next examine how risk-sharing frictions affect the value of public insurance. Public insurance can crowd out the provision of within-household insurance, but if the latter is limited, the welfare gains from public insurance can be larger. We illustrate this point by measuring the welfare gains of a individual-level policy that supplements income below the 20th percentile, financed by a tax on income above the 80th percentile. Under full risk sharing, the median gain is 0.14% of lifetime consumption for both spouses. However, under our model, the gains are 0.47% for women and 0.46% for men, more than three times as large, reflecting the greater need for insurance due to limited commitment and budget separation.

What we learn from the Norwegian setting is that how much private insurance the household provides and, therefore, how much public insurance is worth are ultimately empirical questions. Risk sharing can differ across countries, marital institutions, and household types. Public insurance should be designed based on empirical evidence of how spouses share income risk, rather than assuming that household transfers protect both spouses equally.

Contributions to the literature: First, we provide the first large-scale administrative evidence on how income risk is transmitted to consumption within the household. Most evidence on consumption insurance relies on household-level consumption measures, which may mask differences in individual exposure to income risk when spouses share risk imperfectly (Attanasio and Davis, 1996; Blundell et al., 2008; Kaplan and Violante, 2010; Guvenen and Smith, 2014; Arellano et al., 2017; Eika et al., 2020; Ganong et al., 2025).[2] Previous studies have documented important departures from perfect intrahousehold risk sharing, showing that consumption and nutritional allocations depend on which household member experiences income shocks (Dubois and Ligon, 2009; Duflo and Udry, 2004).

Direct evidence on insurance of private consumption, however, remains limited. Using the Japanese Panel Survey of Consumers, Lise and Yamada (2019) find that wage innovations affect spouses' personal consumption allocations, which they argue is consistent with limited commitment.[3] More recently, Bellue et al. (2026) use Dutch LISS survey data to show that imperfect information about spouses' personal consumption is associated with incomplete risk sharing. Evidence exploiting plausibly exogenous variation in income remains

[2]A related literature documents how spouses adjust their labor supply in response to their partners' earnings shocks, stabilizing household income but not revealing how the additional resources are ultimately allocated to consumption between spouses (Lundberg, 1985; Cullen and Gruber, 2000; Stephens, 2002; Blundell et al., 2016).

[3]Using the same dataset, Fukuda (2024) examine how husbands' involuntary job losses affect both spouses' private consumption, but do not explore how consumption insurance differs following wives' job losses.

particularly scarce. One exception is Robinson (2012), who uses randomized income transfers to married couples in Kenya to examine how transfers received by individual spouses affect private consumption. Compared to this literature, we leverage a large-scale administrative dataset that allows us to exploit plausibly exogenous variation in year-to-year income fluctuations to study how spouses share income risk on personal and joint consumption.

Second, we contribute to the literature on intrahousehold inequality by highlighting disparities in consumption insurance. As emphasized by Chiappori and Meghir (2015), household-level consumption measures can mask substantial differences in individual welfare. A substantial literature documents inequality in consumption levels and resource shares, with implications for individual poverty (Lise and Seitz, 2011; Dunbar et al., 2013; Calvi, 2020; Lechene et al., 2022; Bargain et al., 2022; Calvi et al., 2023; Arduini, 2026). Related work documents disparities in leisure (Beblo and Robledo, 2008; Barigozzi et al., 2023) and the burden of home production (Bertrand et al., 2015; Hancock et al., 2025). We highlight the provision of insurance as an additional dimension of intrahousehold inequality and show that even when average consumption allocations are relatively equal, as in our data, spouses may differ substantially in how much they benefit from sharing income risk.

Third, we show how empirical evidence on intrahousehold risk sharing can help distinguish between competing models of household behavior. Workhorse models of dynamic household behavior typically combine a common household budget with either full or limited commitment (Mazzocco, 2007; Voena, 2015; Lise and Yamada, 2019; Chiappori and Mazzocco, 2017). Past work has focused on distinguishing between commitment regimes by testing how shocks and their histories affect household allocations, labor supply, and bargaining weights (Mazzocco, 2007; Blau and Goodstein, 2016; Lise and Yamada, 2019; Theloudis et al., 2025). We show that standard full- and limited-commitment models struggle to match the cross-spousal consumption responses we estimate. To account for these patterns, we introduce separate budgets into a limited-commitment model, improving its ability to reproduce the observed degree of risk sharing. Our results highlight the importance of budget separation as an additional friction in intrahousehold risk sharing and the value of empirical risk-sharing estimates for disciplining models of household consumption.

Fourth, we connect the literature on household risk sharing to household financial arrangements. Theoretical work has examined household models with "separate spheres" of decision-making and their implications for expenditure provision (Lundberg and Pollak, 1993; Browning et al., 2010). Past research shows that the allocation and control of household resources affect consumption (Lundberg et al., 1997; Ashraf, 2009), saving (Schaner, 2015; Choukhmane et al., 2025), and the allocation of productive inputs (Udry, 1996). Spouses may also sacrifice resources to retain control over their own consumption decisions (Afzal

et al., 2022). However, the role of financial separation in shaping consumption insurance remains unexplored. We provide new survey evidence showing that separate financial arrangements can be quite prevalent in modern economies and vary between personal and joint expenditures. We further demonstrate the importance of this mechanism by showing that budget separation is necessary to reproduce the limited cross-spousal responses in personal consumption observed in the data.

The rest of the paper proceeds as follows. Section 2 describes the data and the measurement of household, personal, and joint consumption. Section 3 presents descriptive evidence on household and individual pass-through. Section 4 develops the BPP-IV identification strategy. Section 5 estimates causal pass-through to household, personal, and joint consumption. Section 6 evaluates how well the limited commitment and separate budget models fit our empirical results. Section 7 quantifies the insurance gains from marriage and welfare gains from public insurance. Section 8 concludes.

2 Data and Measurement

We first describe our spouse-level administrative expenditure transaction data and the linked administrative registers. Additionally, we characterize our analysis sample and define how we measure household, personal, and joint expenditures. We complement the administrative data with an original survey of Norwegian couples documenting financial pooling and the financing of personal and joint expenditures.

2.1 Administrative Data

2.1.1 Expenditure Transaction Data

Our primary measure of consumption comes from administrative transaction data collected by NETS Norway from 2006 to 2018. The data cover two distinct types of electronic payments. The first consists of debit-card transactions processed through BankAxept, the national debit-card payment network. During our sample period, nearly all debit-card transactions at domestic physical stores were processed through this system, and BankAxept payments accounted for more than 80% of card payments at point-of-sale terminals. The second type consists of direct wire transfers out of household bank accounts cleared through the Norwegian Interbank Clearing System (NICS). These transfers include invoice and other account-based payments; on average, the NICS data cover about 77% of total bank wire pay-

ments (Ahn et al., 2024).[4,5] Our data cover a large share of Norwegian taxpayers throughout the sample period, from nearly 80% in 2006 to more than 95% of taxpayers aged 25–75 after 2012.[6]

Spending is classified into 26 categories, including groceries, clothing, health, transport, and recreation. Debit-card payments are classified using Merchant Category Codes, and NICS payments are classified using the merchant's industry code. Appendix Table C.1 reports the categories used. Debit-card transactions are linked to the cardholder, and NICS transfers are linked to the bank-account owner. Transactions are aggregated by the data provider to the individual-week-postal code-category level.

This individual-level structure is central to our analysis of within-household spending. Transaction datasets from individual banks require both spouses to bank at the same institution and often require linked accounts to identify couples. Our data, with near-population coverage, allow us to observe couple spending without this additional requirement that might cause selection in ex ante financial integration. Survey measures of within-household consumption rely on respondents' recall of their own or assigned purchases and are collected at relatively low frequency. Our data capture spending directly for each spouse at a weekly frequency, reducing concerns about recall and self-reporting error.

Despite this broad coverage, the transaction data do not capture all household spending. One important limitation is that we do not observe all cash expenditures.[7] We argue this concern should be limited in our setting. In 2017, only about one in ten point-of-sale payments was made in cash, and Norway had one of the lowest rates of cash use among advanced economies.[8] This low reliance on cash makes Norway a particularly suitable setting for studying household dynamics using electronic expenditure data. In addition, our data do not cover mobile payments or online transactions processed outside of BankAxept. Finally, credit-card spending is imputed from repayments observed in bank transfers, but cannot be classified by expenditure category.

[4] Ahn et al. (2024) show that NETS expenditure data from BankAxept and NICS cover about 80% of household electronic payments in official statistics from 2006 to 2018 and closely track nominal consumption growth in the national accounts, with quarterly correlations of 0.83 for aggregate and 0.81 for per-capita growth.

[5] In U.S. terms, these transfers correspond to ACH payments, including invoices and electronic bill payments. In Norway, most such payments pass through NICS. We exclude person-to-person transfers to focus on consumption expenditures rather than risk-sharing transfers.

[6] Our data do not capture bank accounts that were closed before 2012. As shown in Appendix Figure C.1, this limitation results in lower coverage among older taxpayers prior to 2012. The gap disappears after 2012.

[7] Our data contain a limited measure of cash expenditures from "cash-back" transactions during check-out.

[8] In 2018, cash in circulation was 2 percent of M1 in Norway, compared with 15 percent in the Euro area (Norges Bank, 2019).

Benchmarking: In Appendix Table C.1, we benchmark the NETS expenditure data against Statistics Norway's 2012 Survey of Household Expenditures.[9] Because we do not impute rent from home values in our expenditure measure, we exclude this component from the survey total to ensure comparability. At the aggregate level, mean household expenditure is 361,000 NOK in the survey and 341,000 NOK in our data. The similarity in expenditure levels suggests that our data do capture a large share of household spending. Across expenditure categories, the coverage is stronger for some categories than others: grocery spending aligns very closely with the survey average, while our measure understates transportation spending. Overall, our data reproduce the broad composition of household expenditure, even while the coverage of some categories remains incomplete.

We also assess whether expenditure coverage varies with income. In Appendix Figure C.2, we show that we capture a stable share of total expenditure across the income distribution, both at the bottom and the top. Finally, we test whether the spending in our data exhibits standard Engel-curve relationships with income. As expected, budget shares for food and transportation decline with income, whereas those for clothing and restaurants increase (Appendix Figure C.3).

2.1.2 Additional Administrative Registers

We merge the NETS expenditure data with additional linked administrative registers from Statistics Norway. First, the earnings register records tax-reported individual labor income, pre-tax income, and after-tax income. The wealth register provides measures of liquid and financial wealth, as well as debt. The employment register allows us to identify the employer of each spouse, which we use to construct the firm-driven income shocks for our causal identification strategy. Demographic registers provide information on age, education, location, and family structure, including official records of marriage and cohabitation. Together, these linked datasets allow us to observe consumption, income, wealth, household structure, and employer characteristics for both spouses in the same panel from 2006 to 2018.

2.1.3 Sample Selection and Summary Statistics

We focus on heterosexual married and cohabiting couples, though we also consider same-sex couples in a heterogeneity exercise. Throughout the paper, we use the terms cohabitant and spouse interchangeably.[10] We restrict our sample to spouses aged 20–75. We require that

[9] 2012 is the most recent year available from Statistics Norway, as the survey was discontinued and not resumed until it was redesigned in 2022.

[10] Cohabiting couples are identified in the Central Population Register using Statistics Norway's household classification. Two opposite-sex partners registered at the same address are classified as cohabitants if they

the household has at least one working spouse, at least 1,000 2015 NOK in aggregate labor income and consumption, that the couple is observed together for at least three consecutive years, and that each spouse has non-missing consumption data. These restrictions ensure that we observe both spouses' spending, that the household faces income risk, and that household structure is stable enough to interpret consumption responses.

Table 1 reports summary statistics for our sample. Mean after-tax income is 327,000 NOK for women and 476,000 NOK for men. Labor force participation is similar across genders, at 86% for women and 87% for men. Women contribute 42% of household after-tax income and account for 42% of total household expenditure. The overall expenditure share masks differences across categories: women account for 54% of grocery spending and 73% of clothing spending, but only 31% of home spending. Our final sample includes 821,964 couples observed together for 9.4 years on average, with 7.7 million total couple-year observations.

2.2 Survey Data on Household Financial Arrangements

To complement the administrative data, we conducted an original survey of married and cohabiting Norwegian consumers through the Norstat panel in August 2026.[11] The survey asked respondents how they manage income, financial accounts, and expenditures with their partners, providing direct evidence on the extent of household financial pooling.[12] Our sample contains 1,028 married and cohabiting respondents and closely matches the national population of coupled adults aged 20–75 in age, gender, and region. Moreover, after restricting the survey sample to couples in which at least one spouse works, as in our main sample, after-tax income, labor-force participation, and the female share of household income are similar across the survey and administrative samples (Appendix Table D.1).

2.3 Constructing Measures of Personal and Joint Expenditures

Our transaction data allow us to separately observe expenditures for each spouse. To understand the consumption and welfare effects of income shocks within couples, it is important to distinguish between personal and joint expenditures. Personal consumption includes goods and services consumed primarily or exclusively by one spouse, such as clothing, personal care, and leisure goods and services. Joint consumption, by contrast, benefits multiple household members and includes housing, utilities, groceries, and other shared household necessities.

have a child in common or meet SSB's statistical criteria: they are the only two unrelated adults in the household, both are aged 18 or older, and their age difference is less than 16 years.

[11] Norstat is a market research firm that maintains a curated panel of Norwegian consumers for survey research.

[12] We outline the relevant parts of the survey questionnaire in Appendix D.1.

Table 1: Summary Statistics

	Women	Men
Age	44.6	47.2
After-tax income	327,327	475,790
Share of household after-tax income	0.42	0.58
Labor force participation	0.86	0.87
Total expenditure	237,396	356,880
Share of household expenditure	0.42	0.58
Groceries	0.54	0.46
Home	0.31	0.69
Clothing	0.73	0.27
Direct personal	0.51	0.49
Imputed personal	0.54	0.46
N years per couple	9.4	
N couples	821,964	
N observations	7,716,754	

Note: This table reports sample means for heterosexual married and cohabiting couples in the main analysis sample. Monetary amounts are deflated to 2015 NOK. Income and expenditure shares are spousal shares of the household total. Labor-force participation is the share of observations with positive labor earnings. Home expenditures are utilities and housing maintenance; rent and mortgage are not imputed. Direct personal expenditure sums spending in categories plausibly assignable to one spouse, including clothing, personal care, and leisure goods and services. Appendix Table C.2 reports the classification. Imputed personal consumption is predicted from clothing expenditure using a gender-specific relationship with the personal expenditure bundle estimated before union formation.

For couples with children, childcare and other child-related expenses are often modeled as joint expenditures under altruistic preferences (Blundell et al., 2005). Personal and joint expenditures may respond differently to income shocks, and these responses may have different welfare implications for each spouse.

A key measurement challenge for personal consumption is that the person making a purchase is not necessarily the person consuming the purchased good. Expenditures recorded for one spouse may therefore spill over to the consumption of the other spouse or of other household members. For example, if women purchase most of their spouse's clothing, assigning these expenditures to women's personal consumption would be misleading.

We therefore construct two measures of personal consumption: direct personal expenditure and imputed personal consumption. Direct personal expenditure sums each spouse's spending in categories that are plausibly assignable to an individual, including clothing,

personal care, and leisure goods and services, similar to the card-based approach in Kim (2024). Appendix Table C.2 reports the categories used. To address potential spillovers across spouses, we also construct an imputed measure of personal consumption that is more robust to this concern. In our data, we observe each spouse's expenditure patterns before union formation, when purchases are not subject to spillovers between spouses. If a good remains "assignable" after the union, we can use its pre-union relationship with personal consumption to impute personal consumption afterward. Thus, imputed personal consumption is the level of personal spending that would lead a single agent to choose the observed quantity of the assignable good.

We use clothing as the assignable good. Clothing is frequently used to identify individual consumption and resource shares within the household (Dunbar et al., 2013; Lechene et al., 2022). Appendix Figure C.4 presents a binscatter of clothing expenditure against direct personal expenditure among single individuals 1 to 5 years before the start of a union. The relationship is approximately log-linear and monotonic, with gender-specific intercepts and slopes. This monotonicity allows us to map observed clothing expenditure into imputed personal consumption after union formation.

Motivated by this pattern, we estimate the following gender-specific relationship in the pre-union sample:

$$\log c_{it}^{\text{personal}} = \alpha_s + \beta_s \log c_{it}^{\text{clothing}} + X_{it}'\gamma_s + \varepsilon_{it}, \qquad s \in \{f, m\},$$

where c_{it}^{personal} denotes direct personal expenditure, c_{it}^{clothing} denotes clothing expenditure, and X_{it} contains age-group and education-group indicators and their interactions with clothing expenditure. We estimate this relationship by OLS and then apply the estimated mapping to each spouse's clothing expenditure after union formation:

$$\widehat{\log c_{it}^{\text{personal}}} = \widehat{\alpha}_s + \widehat{\beta}_s \log c_{it}^{\text{clothing}} + X_{it}'\widehat{\gamma}_s,$$

where $\widehat{\log c_{it}^{\text{personal}}}$ is our imputed measure of log personal consumption. Appendix Table C.3 shows the estimated coefficients for each gender.

In our imputation, our main assumptions are the assignability of clothing expenditures and the stability of the mapping between clothing and personal expenditures before and after union. We support the first assumption by measuring the patterns of clothing expenditure around the start of the union. In Appendix Figure C.5, we show that clothing expenditures do not appear to change around the start of the partnership, either as a share of total expenditure or in levels. Spouses therefore do not appear to reallocate clothing purchases across partners after the union, which supports treating clothing as assignable. Although we

cannot assess preference stability directly, we consider the assumption credible because the mapping is estimated among the same individuals before union formation. This assumption is therefore milder than those requiring preference stability between singles and married people in cross-sectional data, or between women and men, as is common in the literature.

Women account for 51% of direct personal expenditure and 54% of imputed personal consumption. Both shares increase with the female income share, as expected (Appendix Figure C.6). The imputed relationship is flatter: imputation moves personal consumption shares toward one-half and reduces the prevalence of extreme shares. We define joint expenditure as total household expenditure minus the sum of both spouses' personal expenditures. Under either personal consumption measure, joint goods account for 77% of household expenditure, close to the public consumption shares reported by Lise and Yamada (2019).[13]

2.4 Constructing Income and Consumption Growth

We aggregate after-tax income and expenditure to the annual level and deflate both to 2015 NOK. Following the consumption insurance literature, we use percent changes rather than level changes, as in marginal propensity to consume (MPC) calculations, to measure proportional consumption responses and limit the influence of large income and consumption movements (Blundell et al., 2008).

To measure income growth for each spouse, we scale income changes by lagged total household income rather than by each spouse's own lagged income. This treatment accounts for differences in spouses' contributions to household income: a 10% change in her income and a 10% change in his income can have very different effects on the household budget. Scaling by household income allows us to compare equally sized changes in household resources originating from each spouse.

Let $\Delta Y_{F,t}$ and $\Delta Y_{M,t}$ denote the changes in household income attributable to each spouse.[14] We calculate:

$$\Delta Y_{F,t} = \frac{Y_{F,t} - Y_{F,t-1}}{Y_{HH,t-1}} \qquad \text{and} \qquad \Delta Y_{M,t} = \frac{Y_{M,t} - Y_{M,t-1}}{Y_{HH,t-1}}.$$

Finally, we winsorize all income and consumption changes at the 1st and 99th percentiles.

[13] Our definition of joint expenditure is conservative because it classifies all uncategorized expenditures as joint consumption. Rent payments are the largest uncategorized household expenditures, and including them helps capture the expected share of joint consumption in household spending. In the empirical analysis, we also report a narrower joint measure based only on observed categories plausibly classified as joint. Joint categories are listed in Appendix Table C.2.

[14] Because $\Delta Y_{F,t} = \frac{Y_{F,t}-Y_{F,t-1}}{Y_{F,t-1}} \cdot \frac{Y_{F,t-1}}{Y_{HH,t-1}}$, the scaled income change is approximately equivalent to the log change in female income multiplied by her lagged household income share, $\Delta Y_{F,t} \approx \Delta \log Y_{F,t} \cdot w^F_{t-1}$, and likewise for $\Delta Y_{M,t}$. We use this alternative format in the consumption function for the BPP-IV estimation.

3 Descriptive Evidence on Consumption and Spousal Income Changes

We start by characterizing the co-movement between income changes and consumption changes in the household using OLS. These descriptive patterns establish the main empirical facts that motivate our causal analysis. We begin with aggregate household consumption and then turn to personal and joint expenditures.

We estimate

$$\Delta \log C_t = \beta_F \Delta Y_{F,t} + \beta_M \Delta Y_{M,t} + \gamma' X_t + \varepsilon_{h,t}, \tag{1}$$

where $\Delta \log C_t$ is the change in log consumption at the household, personal, or joint level. $\Delta Y_{F,t}$ and $\Delta Y_{M,t}$ are changes in household income from the female and male spouses, respectively. The vector X_t includes year-by-municipality fixed effects, year-by-education-group fixed effects, spousal age groups, and nonparametric transition controls for marital status, the number of children, children under age five, and contracted hours for each spouse.[15] The coefficients β_F and β_M measure the pass-through to consumption of an equally sized change in household income originating from each spouse. We cluster standard errors at the household level.

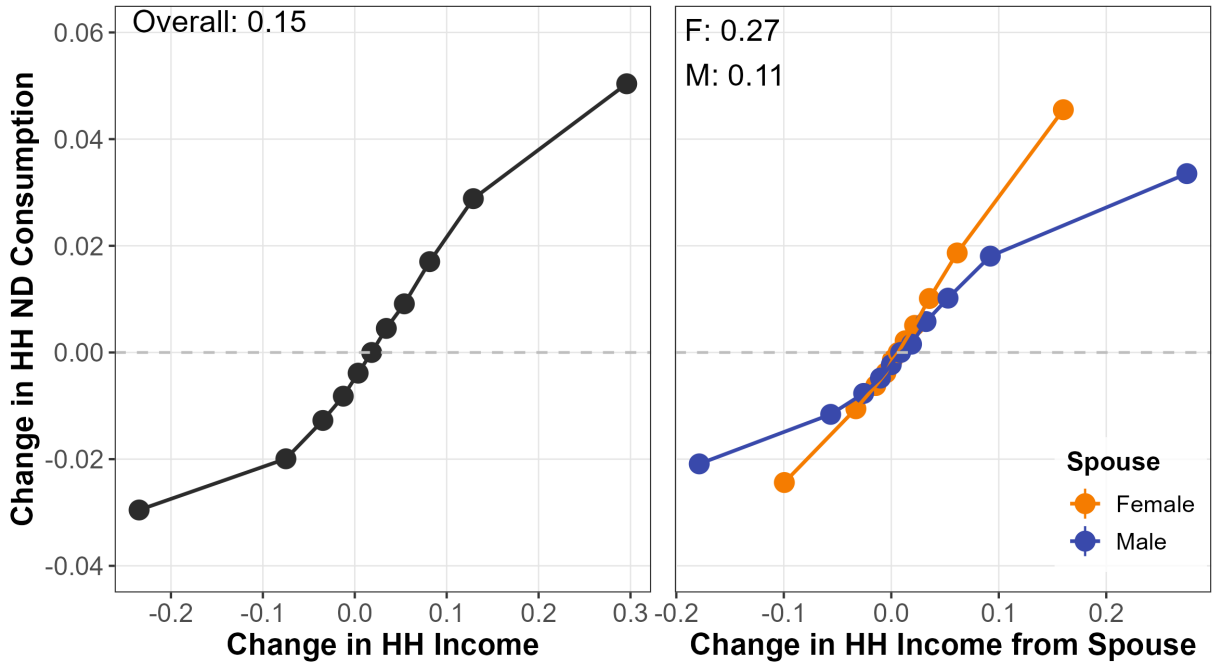
3.1 Household Consumption Responds More to Women's Income

The left panel of Figure 1 plots the aggregate response of household non-durable consumption to household income: a 10% change in household income is associated with a 1.5% change in household non-durable consumption.[16] But as we show in the right panel, the aggregate response combines two different slopes based on the source of the income change. The slope for female income changes is larger than for male income changes. A 10% increase in household income originating from changes in women's income is associated with a 2.7% change in non-durable consumption; a similarly sized increase originating from changes in men's income is associated with a 1.1% change in non-durable consumption. These estimates are our first indication that households may not treat equally sized income changes the same way when they originate from different spouses. Appendix Figure A.1 shows we obtain similar gaps in pass-through for total expenditures and most expenditure categories, including groceries, clothing, and health.

[15] For each variable we include indicators for every pair of lagged and current bins. For marital status these are cells such as married–married, cohabiting–cohabiting, and cohabiting–married. For contracted hours, the controls are indicators for every pair of lagged and current contracted-hours bins for each spouse.

[16] We define non-durable consumption as the sum of expenditures on groceries, transportation services, transportation operation, and utilities, among others. The full classification is in Appendix Table C.2.

Figure 1: OLS Pass-Through of Household Non-Durable Consumption



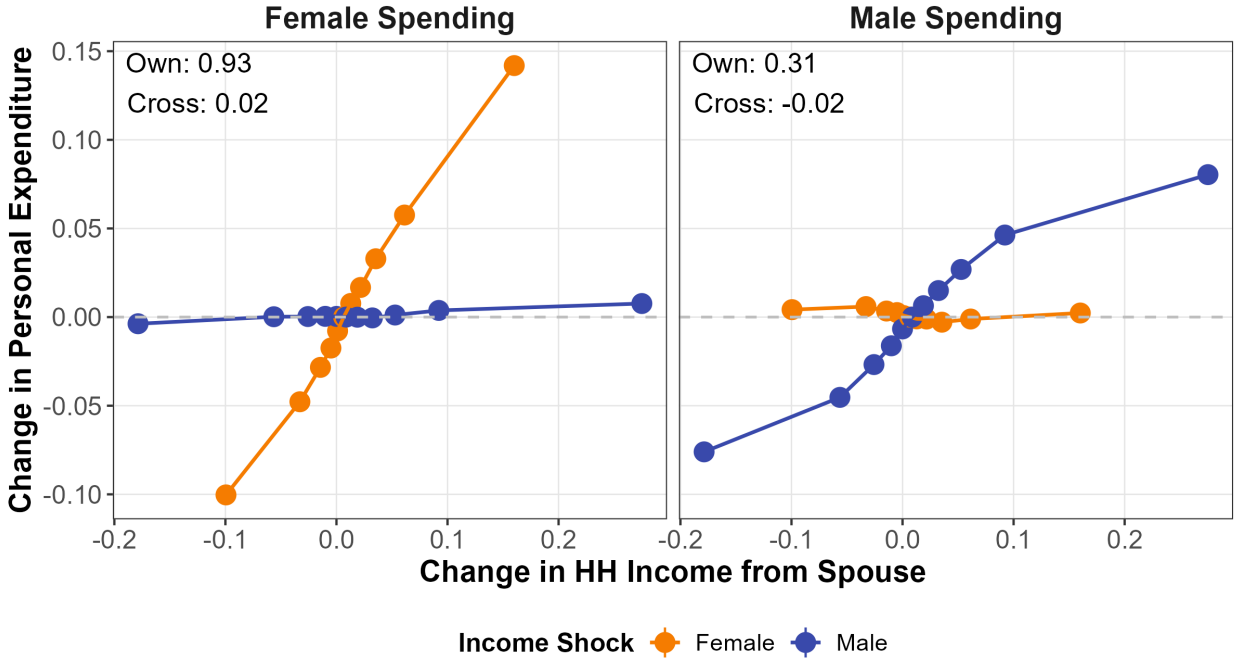
Note: Each panel shows a binscatter of changes in log household non-durable consumption against household income growth. The left panel uses total household income growth; the right panel includes changes in household income originating from each spouse. The binscatter is estimated with bin indicators in place of the linear income changes in equation (1). The slopes are the linear estimates from that equation. Standard errors are clustered at the household level.

3.2 Personal Consumption Tracks Own Income, Not the Partner's

Figure 2 plots changes in direct personal expenditure against female and male income changes, which enter jointly as in Equation (1). For both women and men, personal expenditure tracks own income but is essentially unrelated to the spouse's income changes. The OLS pass-through of female income to female personal expenditure is 0.93, compared with 0.02 for male income. The pass-through of male income to male personal expenditure is 0.31, compared with -0.02 for female income.[17] Personal expenditure is therefore associated primarily with the purchaser's own income rather than with that of the other spouse. This pattern holds even for large changes in household resources: when a decline in the husband's income reduces household income by 20%, his wife's personal expenditure remains essentially unchanged. Appendix Figure A.2 shows a similar pattern for clothing-imputed personal consumption.

[17] These slopes are not driven by any single personal category. We report estimates leaving out one personal category at a time in Appendix Figure A.3.

Figure 2: Changes in Direct Personal Expenditure and Spousal Income



Note: Each panel shows a binscatter of changes in log direct personal expenditure against household income changes from each spouse. The left panel uses female personal expenditure as the outcome and the right panel male personal expenditure; both income changes enter jointly. The binscatter is estimated with bin indicators in place of the linear income changes in equation (1). The slopes are the linear estimates from that equation. Standard errors are clustered at the household level.

3.3 Pass-Through Differs by Gender, Not by Who Earns More

One potential explanation for our results is that these asymmetries in pass-through reflect secondary-earner status rather than gender. On average, women earn 42% of household income and are more likely to be secondary earners. We present two pieces of evidence against this interpretation. First, we split the sample by average female income share. Appendix Table B.1 shows that even when women contribute more than half of household income, household consumption remains more sensitive to female income than to male income. Own personal consumption is also more sensitive to own income for women than for men. Cross-spousal responses remain small regardless of female income share.[18] Second, we compare heterosexual couples in which the woman earns less on average than the man with same-sex couples ranked by average income share (Appendix Figure A.4).[19] Although the same-sex estimates are noisier, the gap between the lower and higher earner is large in heterosexual

[18]The gap also remains when we cut our sample by liquid savings, civil status, and presence of children (Appendix Table B.1).

[19]The lower earner's average share of household income is similar in heterosexual and same-sex couples, at about 0.43.

couples and much smaller in same-sex couples; the gap does not exist in male same-sex couples. These differences suggest that the inequality in pass-through we document reflects a gender mechanism within heterosexual couples rather than a primary vs. secondary-earner mechanism.

4 Empirical Strategy

Our descriptive evidence documents asymmetry in household consumption pass-through and limited cross-spousal responses in personal consumption. However, these OLS consumption responses need not be causal. In our estimation, we need to contend with two challenges simultaneously: endogeneity and shock persistence.

First, observed income changes may be endogenous. Income and consumption may co-move as a result of omitted factors, such as health shocks or labor supply changes. Income may also adjust to consumption, rather than the reverse. Thus, raw OLS estimates may not isolate causal pass-through.

Second, income changes can differ in persistence. Some changes are long-lasting, such as promotions or labor-force exits; others are short-lived, such as bonuses or overtime. Standard consumption theory predicts larger responses to permanent shocks than to transitory shocks, because permanent shocks affect lifetime resources to a larger extent than transitory shocks (Hall, 1978). Changes in male and female income may differ in persistence. If women's income changes are more permanent, higher pass-through from their income could reflect the type of shock rather than differential treatment of male and female income. Appendix Figure A.5 shows that a larger share of women's income growth is indeed permanent.

To address these concerns, we adapt the approach of Blundell, Pistaferri, and Preston (BPP; Blundell et al., 2008) to a two-spouse setting and use firm-driven income shocks as plausibly exogenous sources of variation for identification. Building on the permanent-transitory income process in Meghir and Pistaferri (2004), BPP estimates how much of each type of shock passes through to consumption. The original approach identifies consumption pass-through under assumptions about the exogeneity of own income shocks. We maintain the permanent and transitory structure but instead exploit firm-driven income variation, proxied by changes in co-workers' income, to isolate plausibly exogenous shocks. This BPP-IV strategy allows us to estimate consumption pass-through while accounting for both the persistence and endogeneity of income changes.

4.1 Instrument

We instrument income growth for both spouses with the leave-one-out median change in co-workers' income, building on the identification strategy of Ganong et al. (2025).[20] The instrument is motivated by the Abowd, Kramarz, and Margolis (1999) decomposition of earnings into worker- and firm-specific components:

$$y_{it} = \alpha_i + \psi_{j(i,t),t} + \gamma_{it},$$

where α_i is a worker fixed effect, $\psi_{j(i,t),t}$ is the firm effect, and γ_{it} is an idiosyncratic worker residual.

Changes in the firm component generate income variation that is plausibly independent of the worker's labor-supply choices and idiosyncratic shocks, much like mass layoffs provide plausibly exogenous variation in unemployment. Previous research using Norwegian administrative data similarly exploits firm-level shocks as a source of plausibly exogenous variation in earnings risk (Fagereng et al., 2017). These firm-level changes may reflect changes in firm conditions, including firm productivity or value added, that affect workers' earnings broadly.

We construct the instrument using all workers linked to the same employer at $t-1$, regardless of whether they remain at the firm in period t. The instrument therefore captures income changes from both workers who exit the firm and those who stay. For workers in firms with at least ten co-workers,[21] we calculate the leave-one-out median change in income among co-workers:

$$\Delta Z^s_{it} = \operatorname{Med}_{\substack{k \neq i \\ j(k,t-1)=j(i,t-1)}} \Delta y_{kt}.$$

4.2 BPP-IV Identification Strategy

We next specify the stochastic processes assumed for the instrument and each spouse's income, as well as the consumption function used in estimation.

Firm pay Z^s_{it} for each spouse has a permanent component P^s and a transitory component v^s:

$$Z^s_{it} = P^s_{it} + v^s_{it}.$$

The permanent component is a random walk, $P^s_{it} = P^s_{i,t-1} + \zeta^s_{it}$. The transitory component

[20] Concurrent work by Galaasen et al. (2026) uses a switcher-based firm-pay instrument in the same data to study how firm pay inequality passes through to consumption, while our analysis focuses on intrahousehold risk sharing.

[21] Our BPP-IV results are robust to alternative firm-size cutoffs of 8, 15, or 20 co-workers.

follows an MA(1) process:

$$v_{it}^s = \varepsilon_{it}^s + \theta_s \varepsilon_{i,t-1}^s,$$

so the instrument, spousal firm-pay growth ΔZ_{it}^s, equals the permanent innovation ζ_{it}^s and the change in the transitory component Δv_{it}^s:

$$\Delta Z_{it}^s = \zeta_{it}^s + \varepsilon_{it}^s - (1-\theta_s)\varepsilon_{i,t-1}^s - \theta_s \varepsilon_{i,t-2}^s.$$

A transitory firm-level shock can therefore affect income growth for up to two periods.

The innovations ζ_{it}^s and ε_{it}^s are mean zero, mutually orthogonal, and serially uncorrelated, including across spouses.[22]

Spouse s's income growth Δy_{it}^s equals that firm-pay growth ΔZ_{it}^s plus an idiosyncratic worker residual η_{it}^s:[23]

$$\Delta y_{it}^s = \Delta Z_{it}^s + \eta_{it}^s.$$

We further assume, following Blundell et al. (2008),[24] that consumption growth responds to each spouse's current permanent ζ_{it}^s and transitory ε_{it}^s firm-level innovations,

$$\Delta c_{it} = \psi_F^P w_{it-1}^F \zeta_{it}^F + \psi_F^T w_{it-1}^F \varepsilon_{it}^F + \psi_M^P w_{it-1}^M \zeta_{it}^M + \psi_M^T w_{it-1}^M \varepsilon_{it}^M + u_{it}, \tag{2}$$

where w_{it-1}^s is spouse s's lagged income share and u_{it} is the consumption residual.

The estimands of interest are the four pass-through parameters $(\psi_F^P, \psi_F^T, \psi_M^P, \psi_M^T)$, which measure how consumption growth responds to permanent and transitory firm-driven income shocks of female and male spouses.

These parameters are identified under the following assumptions of income and consumption exogeneity:

[22] To support orthogonality across spouses, we residualize municipality-by-year effects from the instruments to remove common local labor-market shocks. Although our identification can accommodate correlated permanent and transitory shocks across spouses, we abstract from this extension because the estimated cross-instrument covariances are negligible.

[23] For exposition we assume that own income loads one-for-one on the firm shocks in ΔZ_{it}^s. If instead $\Delta y_{it}^s = \lambda_s \zeta_{it}^s + \lambda_s \Delta v_{it}^s + \eta_{it}^s$ and consumption responds to the realized shocks $\lambda_s \zeta_{it}^s$ and $\lambda_s \varepsilon_{it}^s$, identification of (ψ_s^P, ψ_s^T) is unchanged, as shown in Appendix E.

[24] BPP show that the household consumption growth equation, expressed in terms of permanent and transitory household income innovations, can be derived as an approximation to the solution of a consumption-saving problem with CRRA utility.

Assumption 1 (Income Exogeneity) For each spouse $s \in \{F, M\}$, current and lead instruments are uncorrelated with that spouse's idiosyncratic income residual:

$$E\left[\Delta Z^s_{i,t+\ell}\, \eta^s_{it}\right] = 0 \qquad \text{for } \ell = 0, 1, 2.$$

The worker's residual income change does not respond to the instrument. This restriction holds contemporaneously and at one- and two-year leads.[25]

Assumption 2 (Consumption Exogeneity) For each spouse $s \in \{F, M\}$, current and lead instruments are uncorrelated with the consumption residual:

$$E\left[\Delta Z^s_{i,t+\ell}\, u_{it}\right] = 0 \qquad \text{for } \ell = 0, 1, 2.$$

The instrument may affect consumption only by changing income, not through other channels. This restriction holds contemporaneously and at one- and two-year leads.[26]

The identification exploits the distinct patterns of income growth generated by permanent and transitory shocks. Under the assumed income process, permanent firm shocks raise income persistently, while transitory shocks raise income temporarily before reversing over subsequent periods. The covariances of current income growth with future co-worker income changes identify the variance of transitory shocks, while the remaining covariance at $t = 0$ identifies the variance of permanent shocks. The corresponding covariances between consumption growth and co-worker income changes identify the pass-through of each type of shock.

In summary, we adapt the BPP covariance approach to separate permanent and transitory shocks, but replace own income growth with co-worker firm shocks to isolate income variation plausibly exogenous to idiosyncratic worker shocks and labor-supply decisions. Under the assumptions of income and consumption exogeneity, our approach identifies the pass-through parameters for each spouse's permanent and transitory firm-level shocks. Appendix E provides the formal derivation.

[25] Assumption 1 applies only to the own-spouse income residual. If the partner adjusts hours in response to the firm shock, that channel is part of the estimated pass-through. In practice, these cross-spouse income responses are very small (Appendix Table B.2).

[26] Two threats to Assumption 2 exist. Co-worker pay may affect spending through consumption peer effects rather than through income (De Giorgi et al., 2020). Large firm shocks may also change local prices. Municipality-by-year fixed effects absorb common local conditions, but firm-specific price effects may still matter in small municipalities.

4.3 Sample Selection and Estimation

Our identification strategy requires additional sample restrictions. First, we restrict the sample to spouses employed at $t-1$ to identify their set of co-workers. We also require that each spouse remain at the same firm at $t=0$ and $t=+1$, so that the permanent and transitory firm variation we recover still governs the worker's income. If the worker moved to a new firm or became unemployed, those variances need not describe the income path the worker actually experiences.[27] Pass-through is therefore identified from within-firm income changes rather than from firm transitions or unemployment exits.

Second, we also drop couples in which both spouses work at the same firm, to avoid mechanical correlation between the two instruments. These restrictions reduce the original sample to 1.7 million observations. Even so, Appendix Table B.3 shows that the BPP-IV sample has higher income but is otherwise similar to the broader OLS sample on observable characteristics.

We residualize consumption, income and firm income growth using the same controls as in the OLS specification, including municipality-by-year fixed effects that absorb common local labor-market shocks. We estimate the parameters using two-step GMM and cluster standard errors at the household level.

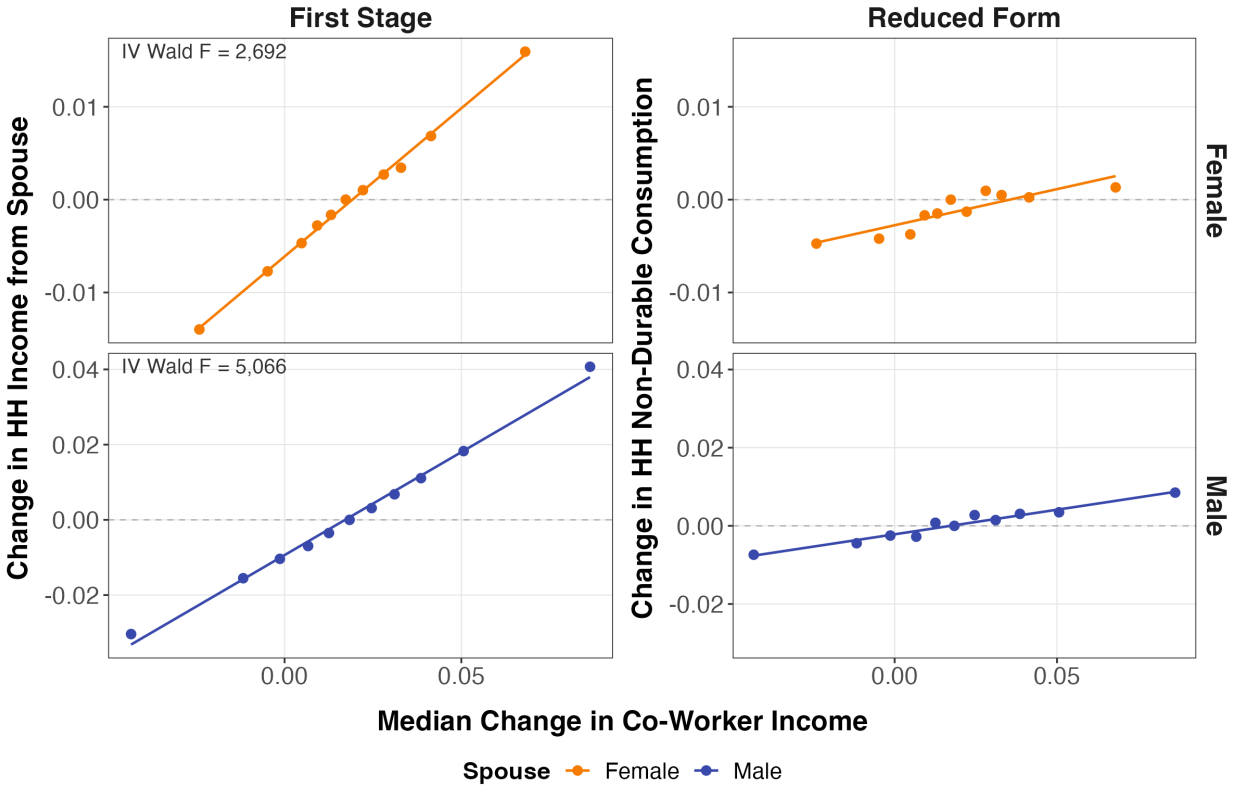
4.4 First Stage and Diagnostics

The instrument has a strong first stage. Figure 3 plots the first-stage relationship between co-worker income growth and the corresponding spouse's contribution to household income growth, together with the reduced-form relationship between co-worker income growth and household non-durable consumption growth, for both spouses. The first stage is tight and linear for both men and women. On the right, we observe that income changes driven by firm variation also translate into household consumption changes. Appendix Table B.2 reports the corresponding first-stage regressions. A 10% increase in female spouse's co-workers' income raises household income by 3.3%, while a 10% increase in male spouse's co-workers' income raises household income by 5.4%.[28] The first-stage Wald statistics are large: 2,692 for women and 5,066 for men. At the same time, cross-spouse first-stage responses are negligible: one spouse's firm shock has a negligible effect on the other spouse's income. Thus, cross-spouse income responses and any associated labor-supply adjustments are unlikely to drive our results.

[27] This restriction may introduce selection bias by conditioning on continued employment. To address this concern, in Section 5.4, we report pass-through estimates without imposing this restriction, imputing future income growth at the original firm, and find qualitatively similar results.

[28] This difference partly reflects the different shares of household income earned by women and men.

Figure 3: First Stage and Reduced Form for Co-Worker Income Shocks



Note: The left column shows the first-stage relationship between household income changes from each spouse and the leave-one-out median income growth of co-workers; IV Wald statistics are reported in each panel. The right column shows the corresponding reduced-form relationship between household non-durable consumption growth and the co-worker income instrument. Appendix Table B.2 shows first-stage regression estimates.

We next examine whether own labor-supply adjustments could explain the observed consumption responses. If consumption and leisure are non-separable in utility, a change in hours can shift consumption even under perfect insurance. Because we require employment at $t = 0$, the estimates are unlikely to be driven by extensive-margin exits, but they could still reflect changes in hours on the intensive margin. Appendix Figure A.6 plots the instrument against changes in contract hours. Contract hours vary little with the instrument. Even at its largest response, the instrument changes women's contract hours by only about two hours.[29] The instrument should therefore be interpreted as a change in earnings with contract hours held fixed. Own labor-supply changes are therefore unlikely to explain the observed consumption responses.

Finally, we assess our statistical assumptions about the income process. Appendix Fig-

[29]The lack of a response may reflect contractual rigidity. We cannot assess how actual hours worked vary with the instrument.

ure A.7 reports a local projection of own income growth on leads and lags of the co-worker instrument. Three features stand out. First, there is a large spike at $t = 0$, consistent with the instrument raising income contemporaneously. Second, income growth is negatively correlated with the instrument at $t = \pm 1$ and $t = \pm 2$. These reversals are consistent with a non-negligible transitory component and with an MA(1) transitory process in firm pay.[30]

Third, the coefficients are symmetric around $t = 0$: the responses at $t = -1$ and $t = 1$ nearly match, as do those at $t = -2$ and $t = 2$. Anticipation would mean that the worker's residual η responds to future firm shocks, which would break this symmetry.[31] Estimates at $t = -3$ are also close to zero. Workers therefore do not appear to adjust earnings in anticipation of the shock.

While informative, these diagnostics cannot assess exogeneity of the worker residual at $t = 0$. Workers could respond to a contemporaneous decline in firm pay by increasing their hours in their main job or taking on additional work to compensate for lost earnings, making the worker residual correlated with the firm shock. As we previously showed, contract hours move little with the instrument, suggesting negligible within-firm labor-supply adjustments. Moreover, dual job-holding is rare in Norway, making contemporaneous adjustments through secondary employment unlikely. Both facts support treating the instrument as plausibly income-exogenous at $t = 0$.

5 Consumption Responses to Spousal Income Shocks

We first estimate the response of aggregate household consumption to permanent and transitory income shocks. We then decompose that response into the personal consumption of each spouse and joint household expenditures. We also analyze heterogeneity across demographic groups and assess the robustness of our estimates to different specification choices.

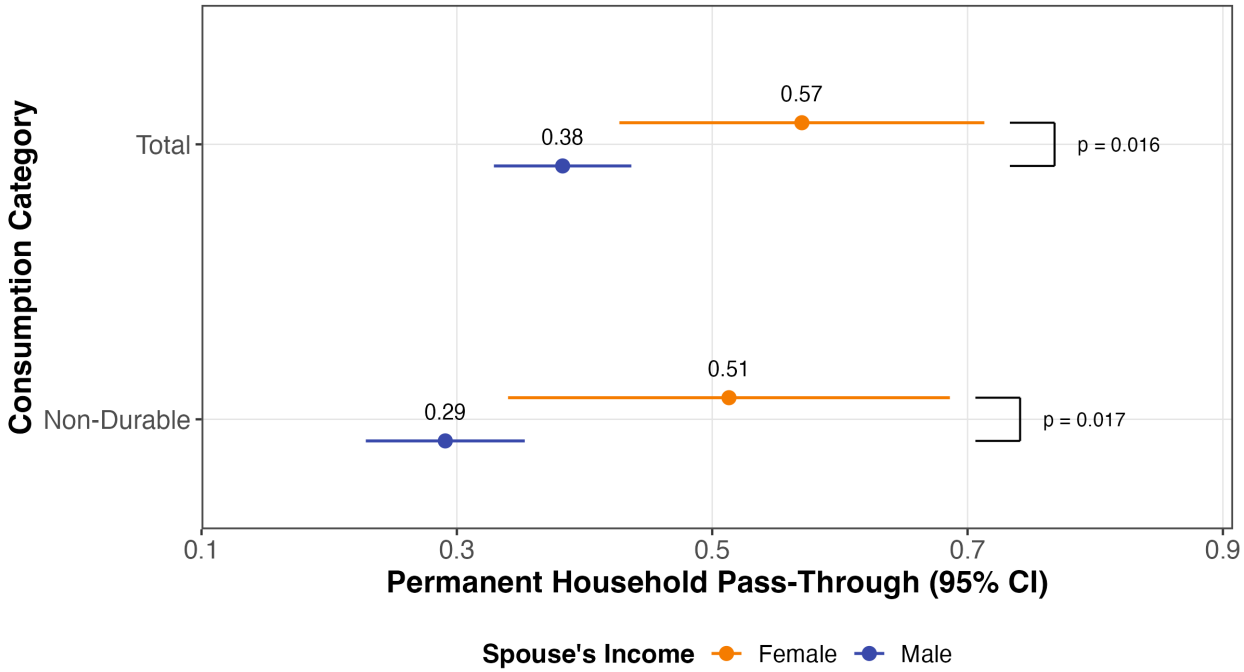
5.1 Aggregate Household Consumption

Figure 4 reports BPP-IV estimates of the pass-through of each spouse's permanent income shock into household consumption. Pass-through depends on the source of the shock. A 10% permanent change in household income originating from women's income raises total expenditure by 5.7%, compared with 3.8% for an equal-sized change originating from men's

[30] If the transitory component were MA(0) in levels, $E[\Delta y_t \Delta Z_{t+1}] < 0$ and $E[\Delta y_t \Delta Z_{t+2}] = 0$. Significant coefficients at $t = \pm 2$ are therefore inconsistent with MA(0).

[31] Without anticipation, the MA(1) process implies $E[\Delta y_t \Delta Z_{t-1}] = E[\Delta y_t \Delta Z_{t+1}]$ and $E[\Delta y_t \Delta Z_{t-2}] = E[\Delta y_t \Delta Z_{t+2}]$. An anticipatory response of η to future ΔZ would enter only the lag moments and break these equalities.

Figure 4: Pass-Through of Permanent Income Shocks to Household Consumption



Note: This figure shows BPP-IV estimates of the pass-through of permanent income shocks to household consumption, for the BPP-IV sample. Each spouse's shocks are scaled by pre-period budget shares as in equation (2). The horizontal axis identifies the spouse receiving the income shock. Horizontal bars show 95% confidence intervals; standard errors are clustered at the household level. The p-value corresponds to the difference between the male and female pass-through estimates. Point estimates and standard errors, including for transitory shocks, are reported in Appendix Table B.4.

income ($p = 0.016$ for the difference).[32] For non-durable consumption, the corresponding elasticities are 0.51 and 0.29. Appendix Table B.4 shows a similar gender gap across other categories, though not all differences are statistically significant. The gaps are especially large for personal categories such as clothing and health.

We find no comparable gender gap for transitory shocks. A 10% transitory increase in household income originating from women's income generates a 2.3% change in total expenditures, whereas an equivalent increase originating from men's income generates a 2.4% change. Most estimates are smaller than those of permanent shocks, consistent with a smaller consumption response to transitory shocks. However, we caution that these estimates are less precise than those for permanent shocks because most of the variation in firm pay is permanent (Appendix Table B.4).

[32]Female pass-through estimates are less precise than male estimates because the co-worker instrument is stronger for men: the first-stage Wald statistics are 2,692 for women and 5,066 for men (Appendix Table B.2).

Taken together, households insure less against female permanent income risk than against male permanent income risk. Methodologically, the differences between permanent and transitory responses show the importance of separating the responses to these two shock types rather than combining them in a single IV or OLS estimate.[33]

An important implication of these results is that the gender incidence of economic downturns can matter for aggregate consumption responses. Because the pass-through of permanent income shocks to spending is higher for women than for men, two downturns with the same decline in aggregate income can generate different consumption responses. Permanent income decreases that fall more heavily on women reduce spending by more.[34] The COVID-19 recession is a clear case of a gendered recession. Alon et al. (2022) document that postwar downturns were typically "man-cessions," disproportionately affecting male-dominated sectors such as manufacturing and construction, whereas the 2020 downturn was a service-led "she-cession" in most countries. As Patterson (2023) shows, when earners with a higher propensity to spend are more exposed to the downturn, the covariance between income losses and spending responses amplifies the aggregate consumption decline. We leave quantification of this amplification to future work.

5.2 Personal and Joint Consumption

We find significant differences in aggregate household spending across permanent changes in male and female income. We now decompose the total response into changes in personal consumption of each spouse and joint household expenditures. We use both direct and clothing-imputed measures of personal consumption that we outline in Section 2.3.[35]

Figure 5 reports the permanent pass-through into direct-personal and joint expenditures. First, we find that permanent shocks to the income of each spouse pass through strongly into that spouse's personal expenditure. The estimated pass-through is 1.44 for women and 0.88 for men, making women's response approximately 64% larger. Thus, for a given proportional change in permanent household income, women experience a larger change in personal expenditure. This difference also persists when shocks are expressed relative to each spouse's own income. Under this alternative scaling, a 10% permanent increase in own income raises personal consumption by 6.2% for women and 5% for men.[36] These findings

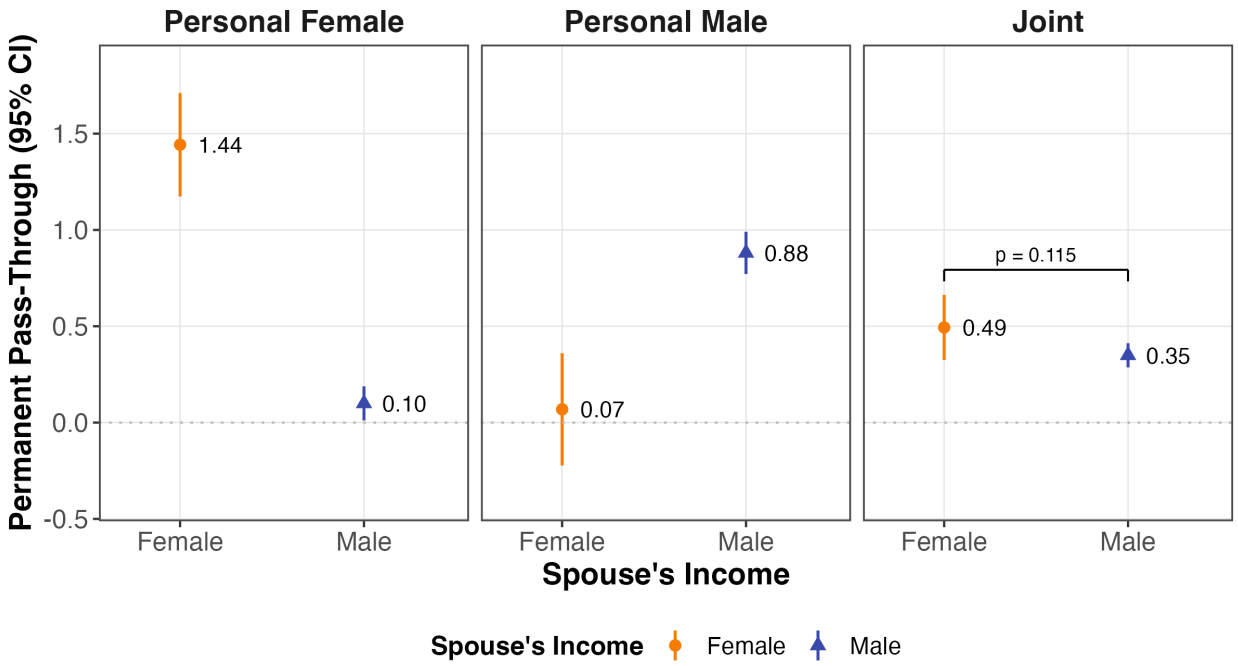
[33] Appendix Table B.6 makes this aggregation explicit. For total household consumption, standard IV estimates are 0.31 for women and 0.26 for men, which combine permanent and transitory pass-throughs into a single estimate for each spouse.

[34] This implication assumes that an extensive-margin income loss, such as job loss, passes through to spending at the same rate as an intensive-margin loss of the same size.

[35] Appendix Table B.5 reports the corresponding permanent and transitory pass-through estimates.

[36] The response to own income is approximately $\psi_s^P \bar{w}^s$, the reported pass-through times that spouse's average share of household income.

Figure 5: Pass-Through of Permanent Income Shocks to Personal and Joint Consumption



Note: This figure shows BPP-IV estimates of the pass-through of permanent income shocks to personal consumption of women, personal consumption of men, and joint expenditures, for the BPP-IV sample. Each spouse's shocks are scaled by pre-period budget shares as in equation (2). Personal consumption is direct expenditure in personal categories. Appendix Table C.2 reports the expenditure classification. The horizontal axis identifies the spouse receiving the income shock. Vertical bars show 95% confidence intervals; standard errors are clustered at the household level. The p-value corresponds to the difference between the male and female pass-through estimates. Point estimates and standard errors, including for transitory shocks, are reported in Appendix Table B.5.

indicate that women's personal consumption remains less insured against fluctuations in female income.

Second, we find that cross-spousal responses are small. A 10% change in permanent household income originating from women's income generates a 0.7% change in men's personal consumption, and the response is not statistically significant. Similarly, an equally sized change originating from men's income generates a 1.0% change in women's personal consumption. These responses are much smaller than the corresponding own-income pass-throughs. This large gap between own-income and cross-spousal pass-through indicates substantial failures of cross-spousal insurance in personal consumption.[37]

We reproduce our estimates with our clothing-imputed measure of personal consumption

[37]These conclusions are not driven by the BPP-IV method. Standard IV yields similar cross-spouse responses: 0.11 for women's consumption and 0.04 for men's consumption (Appendix Table B.6).

in Appendix Figure A.8. Imputation reduces own-income pass-through compared to the direct expenditure estimates, potentially reflecting expenditure spillovers across spouses. Nevertheless, the own-income pass-through of women remains approximately 37% larger than men's (0.96 vs. 0.70). Cross-spousal pass-through remains substantially smaller than own-income pass-through, although the cross-spousal response of women's personal consumption to male income rises from 0.10 to 0.23.[38] Across both measures, the exposure of personal consumption to household changes from own income remains larger for women than for men. Additionally, we conduct a bounding exercise to assess how expenditure spillovers across spouses affect the observed own–cross pass-through gap. We show in Appendix E.1 that, under fixed spillovers, i.e., a constant share of expenditures benefiting the other spouse, the observed pass-through gap is a lower bound on the true gap in personal consumption. Thus, fixed expenditure spillovers alone cannot explain the own–cross gap we document.

Third, we find that the pass-through into joint expenditures is more equal across spouses. Joint expenditures, including groceries and home spending, have a lower pass-through than personal consumption. The lower elasticity likely reflects the fact that joint expenditures capture household necessities. A 10% permanent change in household income originating from women's income generates a 4.9% change in joint expenditures, whereas an equivalent change originating from men's income generates a 3.5% change. Although the female estimate remains larger, the difference is not statistically significant, and it is much smaller than the differences in pass-through we observe for personal consumption.[39] These estimates indicate that joint expenditures are both more insured and less sensitive to the gender source of the household income shock.

5.3 Heterogeneity

Appendix Table B.7 examines heterogeneity in pass-through across household characteristics.

Liquidity: Household liquidity may affect the extent to which spouses can smooth consumption following income shocks. We show results by liquid asset tercile. The gap in aggregate household pass-through holds across liquidity groups. Own-income pass-through is larger when household liquid assets are lower. Between the bottom and top terciles, it falls from 1.83 to 1.22 for women and from 1.28 to 0.50 for men. Cross-spouse responses stay small relative to these own responses. The point estimate for women's shocks on men's

[38] The difference between the direct and implied cross-spousal estimates of male income shocks on female personal consumption is not statistically significant.

[39] We find similar estimates when we use the clothing-imputed residual for joint expenditures (0.56 vs. 0.39) and when we focus on joint expenditures categorized directly in the data (0.54 vs. 0.32).

personal consumption is larger when liquidity is low, which could indicate more pooling of men's personal consumption in those households, but these estimates are imprecise.

Female Income Share: Spouses' relative incomes may influence how income shocks are shared within the household. We cut the sample by terciles of the female income share. The gap in aggregate household pass-through also holds in each tercile. The own-income gap is larger when women earn a smaller share of household income. In the bottom tercile, own-income pass-through is 2.43 for women and 0.75 for men. In the top tercile, women's and men's own pass-throughs are close, 0.92 and 1.00. Cross-spouse responses remain small relative to these own responses in each tercile.

Civil Status: Whether partners are married or cohabiting may be associated with different levels of household risk-sharing. Both groups show the same pattern: women's pass-through is higher for household expenditure and for their own personal consumption, and cross-spouse responses remain small. The difference in own pass-through is larger among cohabiting couples than among married couples.

Children: Having children can create different levels of sharing and commitment in the household. Splitting the sample by whether spouses have children below age 18 yields similar patterns across these two groups.

Children can also complicate the measurement of personal expenditures. If women account for a larger share of these purchases, part of the gender difference in own-income pass-through could reflect spending on children rather than differences in personal consumption. Appendix Figure A.9 compares our main results to results in the sample without children. We find similar estimates of own- and cross-spouse pass-through. Although the estimates are smaller, the gender gap in own-income pass-through remains significant (1.30 vs. 0.58). The gender difference in individual pass-through is therefore unlikely to be driven primarily by child-related expenditure spillovers.

5.4 Robustness

Changes in Hours: Our income shocks may coincide with changes in contract hours. If consumption and labor supply are non-separable, changes in contract hours could affect consumption even in the absence of differences in consumption insurance.[40] In the main specification, we residualize with respect to indicators of transitions in contract hours, so that identifying variation comes from workers who experience similar changes in contract

[40] Appendix Figure A.6 shows that the instrument does not change contract hours significantly.

hours but different associated earnings changes. However, conditioning on contract hours transitions could induce collider bias if workers' hours decisions also depend on unobserved determinants of consumption. To assess this concern, we re-estimate without the contract hours transition controls and obtain similar estimates (Appendix Table B.8). We retain these controls in our preferred specification because they bring us closer to isolating lump-sum earnings changes.

Extensive Margin: In the main specification, we require workers to remain at the same firm through t+1. We impose this restriction so that the permanent and transitory firm shocks we identify continue to govern the worker's realized income path. The cost is that estimation focuses on intensive-margin responses among firm stayers.

We relax the restriction by requiring employment only at $t = -1$, thereby including workers who separate into non-employment after the shock. For these workers, we construct lead instruments from the leave-one-out median income growth of co-workers who remain at the origin firm. Re-estimating on this broader sample yields pass-through estimates similar to the main specification (Appendix Table B.8). Our main estimates are therefore not an artifact of conditioning on continued firm attachment. The similarity of the estimates also suggests that consumption responses along the extensive margin are not substantially different from those along the intensive margin.

Differences in Age: Given that women are on average younger than their male spouses, the pass-through gaps we document could reflect women's longer remaining working lives, as a permanent income shock would have a larger effect on lifetime resources and generate a larger consumption response. We restrict the sample to couples whose age difference is at most two years and obtain similar results, indicating that these life-cycle considerations do not drive our findings (Appendix Table B.8).

Over-identification Test: The GMM system is overidentified: we use two more moment conditions than parameters. This feature of our estimation allows us to compute Hansen's J-test of the full set of moments. We show the results of these tests in Appendix Tables B.4 and B.5. We fail to reject the test at the 5% level for most household, personal, and joint expenditure estimates. This result supports the joint validity of our assumptions about the income process, consumption function, and instrument exogeneity, but does not establish the validity of each assumption individually.

Construction of ΔZ: To construct the instrument, we select observations in which each spouse has at least 10 co-workers to calculate median co-worker income changes. Appendix Table B.9 shows that we obtain similar results for our main estimates when we construct the instrument with 8, 15, or 20 co-workers.

6 Models of Household Risk-Sharing

What mechanisms rationalize the pass-through patterns we observe? Permanent female income shocks raise household spending by more than similarly sized male shocks. Personal consumption tracks own income closely and responds much less to the partner's income, while joint spending is more insured and responds more equally to shocks from either spouse. These patterns point to limited intrahousehold risk sharing in personal consumption, with more pooling of risk in joint expenditures. In this section, we ask which household models can explain these patterns.

We start with the benchmark of full income pooling, under which household consumption allocations depend on total household resources, not on which spouse receives the income. Several canonical household models imply this property. Unitary models represent the household with one objective and one budget constraint. Therefore, for pure income shocks, only the size of the shock should matter for consumption allocations.[41] Collective models allow spouses to have distinct preferences and represent consumption allocations as the maximization of a Pareto-weighted sum of spouses' utilities. If Pareto weights remain fixed over marriage, as under "full commitment", the source of a pure income shock should again not affect the response of household consumption allocations (Mazzocco, 2007).[42]

Full income pooling has often been rejected empirically, both in the allocation of household resources and in responses to income shocks (Browning et al., 1994; Duflo and Udry, 2004; Lise and Yamada, 2019). Using transaction data, we show that household consumption also responds differently to year-to-year income shocks depending on which spouse experiences them in Norway. Moreover, we find that the degree of income pooling differs substantially between personal consumption and joint household expenditures.

In this paper, we sequentially introduce two risk-sharing frictions that can break full

[41] If consumption and labor supply are non-separable, a wage change can affect allocations even under the unitary model by changing the price of spousal leisure. Because we find that firm-pay changes have little effect on contract hours, our income shocks should be interpreted as changes in earnings with contract hours held fixed.

[42] Perhaps surprisingly, even non-cooperation need not break income pooling. In standard voluntary-contribution models, spouses independently choose contributions to a joint public good, yet equilibrium allocations can still depend only on total household resources (Bergstrom et al., 1986; Browning et al., 2010).

income pooling:

1. *Limited commitment:* Spouses cannot commit to future consumption allocations as under full commitment. A positive income shock can improve a spouse's outside option and, if the participation constraint of marriage binds, raise their Pareto weight. The resulting renegotiation thus allocates a larger share of household resources toward their personal consumption. Limited commitment with a common household budget is the workhorse model in the dynamic household consumption literature (Mazzocco, 2007; Voena, 2015; Theloudis et al., 2025).

2. *Separate budgets:* Spouses maintain separate budget constraints rather than pooling all resources. An increase in one spouse's income then increases that spouse's budget and need not be shared fully with their partner.[43]

First, we show that limited commitment weakens insurance, but with a joint budget, it still implies substantially more cross-spouse insurance of personal consumption than we observe. Because spouses renegotiate only when participation constraints bind, matching the large gaps between own- and cross-spouse pass-through would require frequent and substantial renegotiation. Second, our new survey of Norwegian couples documents that many spouses retain separate control over resources rather than pooling income fully. We therefore develop a model that allows spouses to maintain separate budgets while contributing a share of their resources to joint expenditures. We show that this model can reconcile very limited insurance of personal consumption with substantial insurance over joint consumption.

6.1 Limited Commitment with Single Budget

We first consider the standard collective model with a single budget and limited commitment. Income shocks may trigger renegotiation of spouses' bargaining weights, allowing personal consumption to respond differently to own and spouse's shocks. We assess whether renegotiation alone can reproduce the observed pass-through gap in personal and joint expenditures.

Household Problem: Spouses $s \in \{1, 2\}$ pool income and maintain a single joint asset a_t. Spouse s's marital flow utility combines effective consumption $C_{s,t}$ and match quality q_t, which captures the emotional and other non-pecuniary value of the union:

$$U^M_{s,t} = u_s(C_{s,t}) + q_t, \qquad u_s(C) = \frac{C^{1-\gamma_s} - 1}{1 - \gamma_s}.$$

[43] Private information can give rise endogenously to incomplete pooling: when income can be diverted toward unobserved personal consumption, incentive compatibility requires spouses to retain part of their marginal income (Bellue et al., 2026).

Following Lise and Yamada (2019), effective consumption is a CES aggregate of private consumption $c_{s,t}$ and the joint good g_t:

$$C_{s,t} = \left(\alpha_{cs} c_{s,t}^{\phi} + \alpha_{gs} g_t^{\phi}\right)^{1/\phi}$$

The parameter ϕ determines substitutability between the two goods, and α_{cs} and α_{gs} determine spouse s's relative tastes for private and joint consumption.

Match quality follows a random walk:

$$q_t = q_{t-1} + \varepsilon_{q,t}, \qquad \varepsilon_{q,t} \sim N(0, \sigma_q^2),$$

where $\varepsilon_{q,t}$ is a common match-quality shock.

Each spouse s receives income $y_{s,t} = \exp(p_{s,t})$, where permanent log income follows an independent random walk with i.i.d. innovations $\eta_{s,t}$:[44]

$$p_{s,t} = p_{s,t-1} + \eta_{s,t}, \qquad \eta_{s,t} \sim N(0, \sigma_{\eta_s}^2).$$

Given this married state $x_t = (a_t, p_{1,t}, p_{2,t}, q_t)$ and the household budget constraint, the household solves

$$\begin{aligned} \max_{c_{2,t}, g_t, a_{t+1}} \quad & \mu_t \left[U_{1,t}^M + \beta \mathbb{E}_t V_{1,t+1}^M(x_{t+1}, \mu_{t+1})\right] \\ & + (1-\mu_t)\left[U_{2,t}^M + \beta \mathbb{E}_t V_{2,t+1}^M(x_{t+1}, \mu_{t+1})\right] \\ \text{subject to} \quad & c_{1,t} + c_{2,t} + g_t + a_{t+1} = y_{1,t} + y_{2,t} + (1+r)a_t, \\ & a_{t+1} \geq \underline{a}. \end{aligned}$$

The solution to this problem defines a spouse-specific value of marriage $V_{s,t}^M$:

$$V_{s,t}^M(x_t, \mu_t) = U_{s,t}^M + \beta \mathbb{E}_t V_{s,t+1}^M(x_{t+1}, \mu_{t+1}).$$

The expected continuation value $\mathbb{E}_t V_{s,t+1}^M$ integrates over future income and match-quality shocks and allows for subsequent renegotiation of the Pareto weight and for separation.

If the couple separates, assets are split equally: spouse s enters singlehood with $a_{s,t}^S = \frac{1}{2} a_t$.[45] Given single state $x_{s,t}^S = (a_{s,t}^S, p_{s,t})$ and an individual budget constraint, the divorced spouse solves:

[44] We abstract from transitory income risk because the gender differences motivating the model arise primarily from permanent shocks.

[45] Under Norway's default divorce rules, marital wealth is split equally. We abstract from exceptions for clearly premarital or inherited wealth.

$$
\begin{aligned}
V^S_{s,t}(x^S_{s,t}) = \max_{g^S_{s,t}, a^S_{s,t+1}} \quad & u_s(C^S_{s,t}) + \beta \mathbb{E}_t V^S_{s,t+1}(x^S_{s,t+1}) \\
\text{subject to} \quad & c^S_{s,t} + g^S_{s,t} + a^S_{s,t+1} = y_{s,t} + (1+r)a^S_{s,t}, \\
& a^S_{s,t+1} \geq \underline{a}^S.
\end{aligned}
$$

Single flow utility contains no match-quality term. Singles consume the same categories of goods as married spouses, but $g^S_{s,t}$ is no longer shared with a partner.[46]

We now outline the renegotiation process under limited commitment. After period-t shocks are realized, both spouses must satisfy their participation constraints: each must prefer remaining married to separating.[47] Spouse s's outside option $\widetilde{V}^S_{s,t}(x_t)$ is the value of the single problem with post-divorce assets and current permanent income, minus a utility cost of divorce δ_s:

$$
\widetilde{V}^S_{s,t}(x_t) = V^S_{s,t}(\tfrac{1}{2}a_t, p_{s,t}) - \delta_s.
$$

Each period, the value of the inside option $V^M_{s,t}$ must be larger than $\widetilde{V}^S_{s,t}(x_t)$:

$$
V^M_{s,t}(x_t, \mu) \geq \widetilde{V}^S_{s,t}(x_t), \qquad s \in \{1, 2\}.
$$

Let $\underline{\mu}_t(x_t)$ be the lowest Pareto weight that satisfies spouse 1's participation constraint, and let $\overline{\mu}_t(x_t)$ be the highest weight that satisfies spouse 2's. If $\overline{\mu}_t(x_t) < \underline{\mu}_t(x_t)$, no feasible weight exists to keep both spouses in the marriage, and the couple separates.

If a feasible band exists, following Ligon et al. (2002), the updated weight μ_t is given by:

$$
\mu_t = \begin{cases}
\overline{\mu}_t(x_t) & \text{if } \mu_{t-1} > \overline{\mu}_t(x_t), \\
\mu_{t-1} & \text{if } \mu_{t-1} \in [\underline{\mu}_t(x_t), \overline{\mu}_t(x_t)], \\
\underline{\mu}_t(x_t) & \text{if } \mu_{t-1} < \underline{\mu}_t(x_t).
\end{cases}
$$

Thus, μ_t stays constant whenever the initial weight remains feasible, and otherwise is updated minimally to the nearest feasible boundary, i.e., renegotiation.

If the Pareto weight remains fixed, we devolve to full commitment: equal lump-sum changes in household income coming from either spouse generate the same consumption responses,

$$
\left.\frac{\partial c_s}{\partial y_1}\right|_{\mu_t = \mu_{t-1}} = \left.\frac{\partial c_s}{\partial y_2}\right|_{\mu_t = \mu_{t-1}}.
$$

[46] We also assume that singles cannot remarry in future periods.

[47] In contrast, under full commitment, separation is not allowed and renegotiation does not occur: Pareto weight μ_t stays fixed.

However, suppose a favorable permanent income shock raises spouse 1's outside option enough that their participation constraint binds. Restoring participation requires increasing μ_t, which raises $V_{1,t}^M$ by reallocating resources toward spouse 1's preferred bundle. An equal household income gain from spouse 2 expands the budget by the same amount but does not trigger this increase in μ_t. If personal consumption c_1 increases with μ_t[48], limited commitment predicts,

$$\frac{\partial c_1}{\partial y_1} \geq \frac{\partial c_1}{\partial y_2},$$

with equality when μ is unchanged and a strict gap when renegotiation occurs.

Limited commitment can therefore make personal consumption responses to equal-sized household income shocks depend on which spouse receives the shock, even though spouses pool their resources.

Stylized Model: We study a parsimonious three-period version of the model. The time structure isolates the limited-commitment mechanism and allows us to assess whether it can reproduce the observed gap between own- and cross-spouse pass-through.

$t = 1$: The couple forms and the initial Pareto weight is set.

$t = 2$: Income and match-quality shocks realize, spouses renegotiate or separate, and then consume and save.

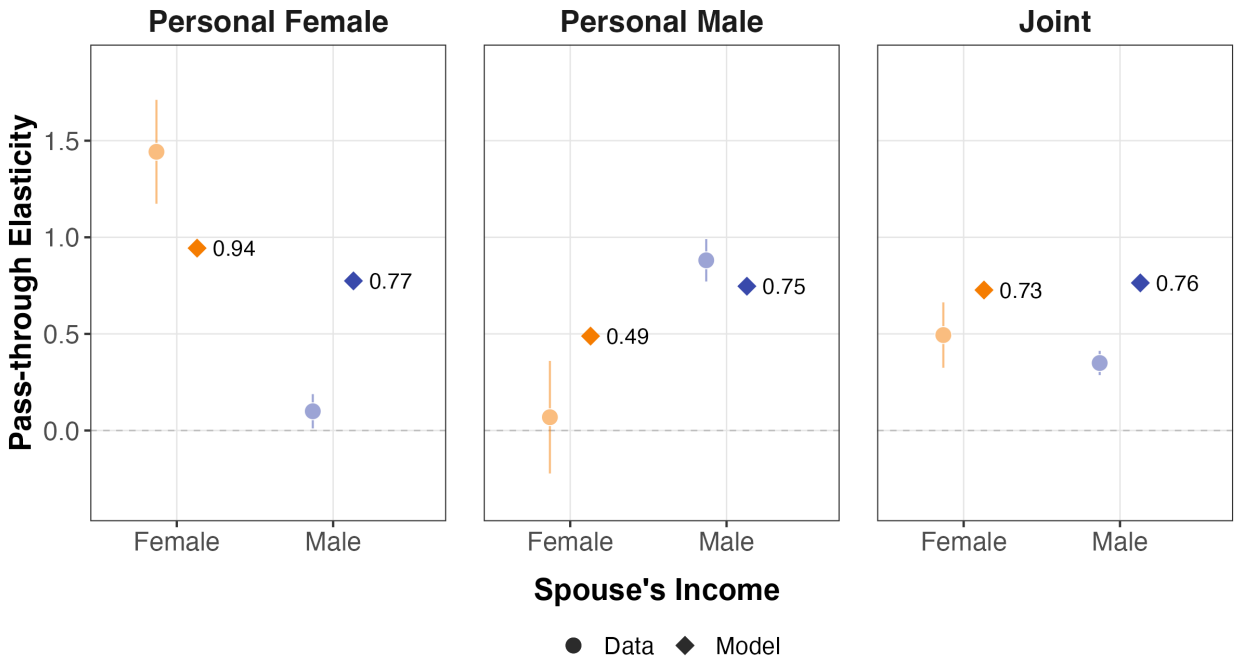
$t = 3$: Final shocks and participation decisions realize, and all remaining assets are consumed.

Estimation: We solve the model by backward induction and estimate it by simulated method of moments (SMM) to match pass-through into personal and joint consumption. Estimation proceeds in two stages. In the first stage, we follow Lise and Yamada (2019) and use the intratemporal first-order conditions for personal and joint consumption to estimate the joint-good taste weights. The resulting intratemporal condition does not depend on the Pareto weight, allowing us to estimate the taste weights without observing or estimating it beforehand. Fixing the substitutability parameter $\phi = 0.1$,[49] we obtain estimated joint-good taste weights of 0.67 for women and 0.62 for men, consistent with previous evidence that women place relatively greater weight on the joint good.

[48] If spouse 1 strongly prefers g and c_1 and g are substitutes, c_1 might fall even if μ_t increases.

[49] We choose $\phi = 0.1$ for both spouses, which gives an elasticity of substitution between private and joint consumption of 1.11, just above Cobb–Douglas. This parameter is close to the estimate of Lise and Yamada (2019) for women ($\phi = 0.15$) in a utility function that also includes leisure. Appendix Table F.1 reports the weights at alternative values of ϕ.

Figure 6: Pass-Through in the Data and Single-Budget Limited-Commitment Model



Note: Circles show BPP-IV estimates of permanent income pass-through in the data, with 95% confidence intervals. Diamonds show simulated pass-through from the estimated limited-commitment model with a single budget. Panels report pass-through on personal consumption of each spouse and joint expenditure; the horizontal axis is the spouse whose own income changes. Estimated parameters, targeted moments, and model elasticities are reported in Appendix Table F.3.

In the second stage, we hold these personal and joint weights fixed and use SMM to estimate spouse-specific risk-aversion parameters and divorce costs, together with the variance of the match-quality shock. We initialize the simulations by drawing from the empirical joint distribution of income and assets among new couples aged 25 to 45, then simulate permanent income innovations using the estimated gender-specific shock variances. We target the own- and cross-spouse pass-through elasticities and observed divorce rates. Appendix F provides further details on estimation, targeted moments, and parameter estimates for both the single-budget and separate-budget models.

Results: Figure 6 shows the main limitation of the single-budget model. In the data, own-income pass-through into personal consumption is high and cross-spouse pass-through is near zero. In its best-fitting parameterization, the model generates some gap between own- and cross-spouse pass-through, as expected from limited commitment. However, it substantially overstates cross-spouse responses in personal consumption. In the data, the

cross-spouse elasticities are 0.10 for women and 0.07 for men, whereas the best-fitting model predicts 0.77 and 0.49.

The model predicts substantially more insurance of personal consumption than we observe because all resources remain pooled until renegotiation. Appendix Table F.3 shows that the model's best fit occurs at zero divorce costs, the lowest value permitted in the estimation. Limited commitment could reproduce these findings only with negative divorce costs, so that remaining single is more attractive than marriage and participation constraints bind period by period.

We also test other values of ϕ, the complementarity parameter between joint and personal goods, and find similar pass-through elasticities. Targeting clothing-imputed personal consumption instead of direct personal expenditure produces a similar lack of fit (Appendix Table F.4).[50]

6.2 Survey Evidence on Budget Separation

Our findings show that personal expenditures are closely tied to own income but respond little to spouse's income shocks. As we have shown, these pass-through patterns are difficult to reconcile with models of household consumption in which spouses fully pool their resources. These facts raise an important question: how do couples organize their finances in practice, and can these arrangements help explain the pass-through patterns we document?

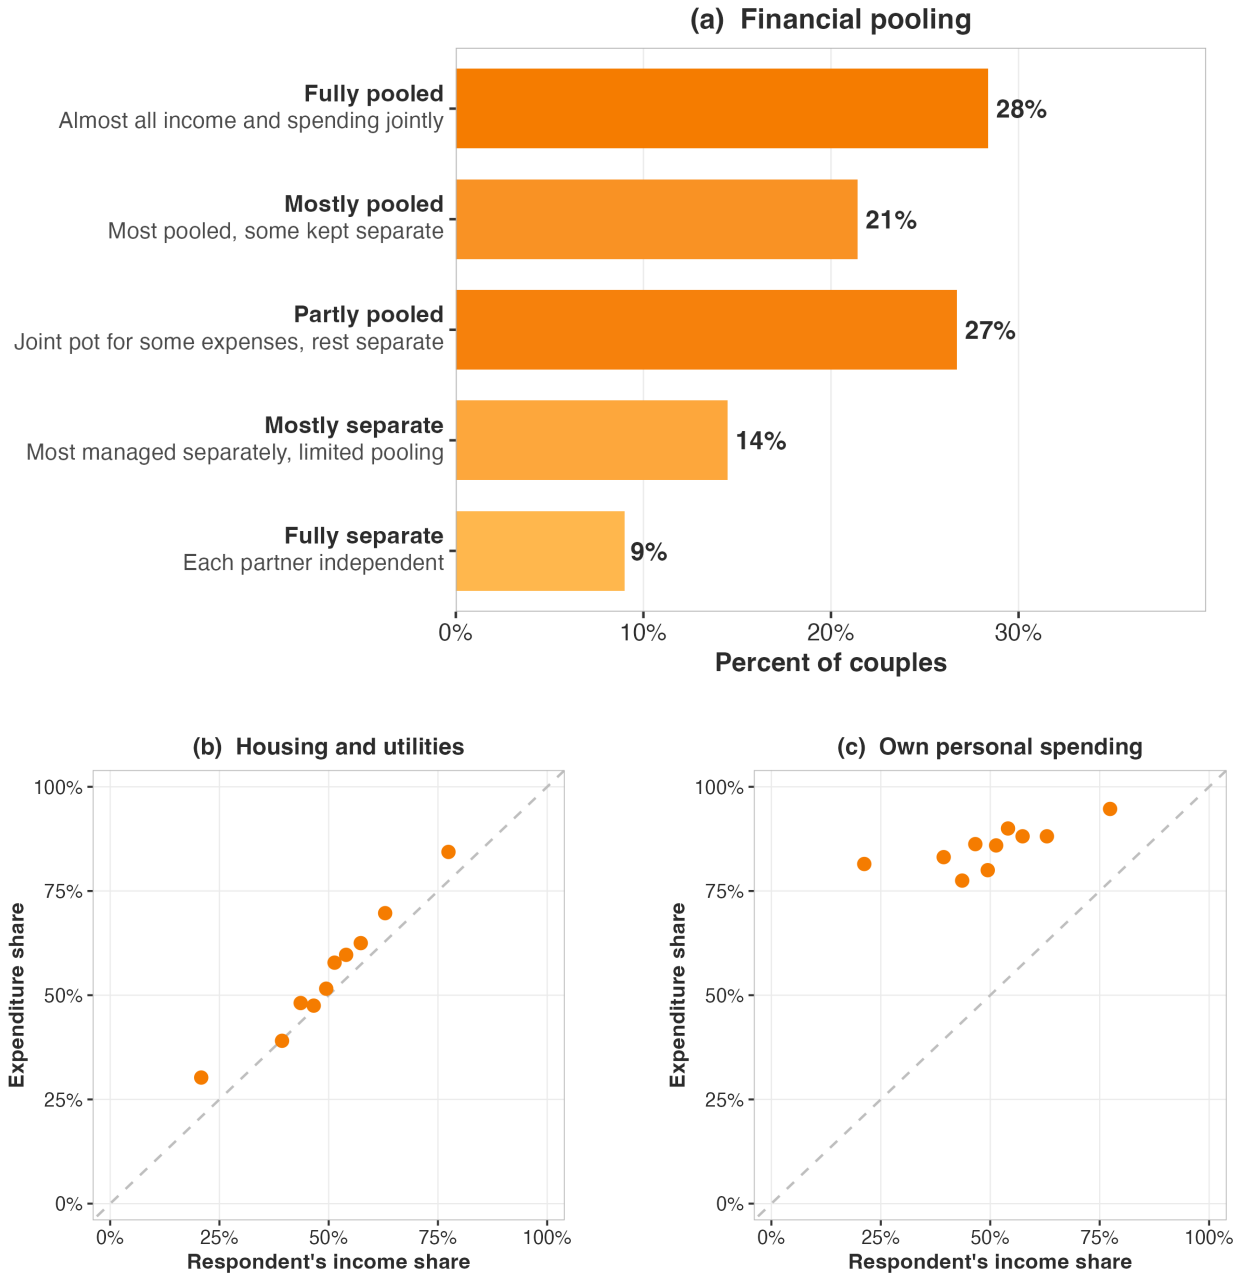
To shed light on household financial arrangements, we turn to our survey of Norwegian couples described in Section 2.2. We start by asking how couples pool their finances. Figure 7(a) shows that a single pooled budget is far from universal. Only 28% of respondents report pooling almost all resources. An additional 21% report pooling most resources while keeping some income separate. The remaining 51% report meaningful financial separation: 27% use a joint "pot" for shared expenses while keeping the rest of income separate, 14% manage most finances separately, and 9% maintain fully separate accounts. Thus, for almost 70% of couples, the natural description of the household is not one unique budget but a combination of shared and individually controlled resources.

The reasons households maintain some financial separation reflect both practical considerations and a desire for autonomy. On the practical side, 42% find separate accounts easier to manage and 40% retain accounts established before the relationship. At the same time, 42% of respondents cite independence in financial decisions and 40% cite different spending or saving preferences as reasons for maintaining separate accounts (Appendix Table D.2).

Financial separation is also not unique to Norway. A UK YouGov survey from 2023–2024

[50] In the baseline specification, we impose a zero borrowing limit. Appendix Table F.4 shows that allowing for borrowing ($\underline{a} = -0.5$) yields similar parameter estimates.

Figure 7: Survey Evidence on Budget Separation



Note: The data are from the authors' survey of Norwegian couples fielded in August 2026. The final sample included 1,028 respondents. Appendix Table D.1 shows that our sample is representative of national targets for married and cohabiting spouses. Panel (a) reports financial pooling arrangements. Panels (b) and (c) plot the respondent's share of category expenditure against their share of household income for housing and utilities and for own personal spending, respectively. Points are means within income-share deciles. Dashed lines are 45-degree lines. Full survey questions are reported in Appendix D.1.

among married and cohabiting respondents finds that only 33% of respondents pool all income and expenses, 27% contribute a fixed amount to a joint account, while 31% use only individual accounts (YouGov, 2024) (Appendix Figure D.1). This mix of arrangements in the UK and Norway echoes a broader sociological literature documenting substantial heterogeneity in household financial organization (Pahl, 1995; Ashby and Burgoyne, 2008; Addo and Sassler, 2010).

Is financial separation the same as budget separation? Even if spouses maintain separate accounts, they could still potentially pool resources in practice: two pots of money managed together instead of independently. Two pieces of evidence argue against this interpretation. First, we ask respondents to what extent they can *access* their partner's personal accounts. Thirty percent of respondents report having no access to their partner's personal account, and among those with access, most report never using that account to finance their own personal consumption (Appendix Figure D.2). Separate accounts represent more than accounting labels: they can create de facto limits on household resources across spouses.

Second, we ask respondents to what extent they *view* money as fungible across spouses. Half of Norwegian respondents agree that the money they earn feels like "mine," and a majority disagree that their partner's money is equally available to them, even though many respondents also express the view that household money should be "ours" (Appendix Figure D.3). Couples may therefore recognize money as not fungible across spouses, while still wishing to share more of it. Thus, financial separation is indeed linked to budget separation: a large majority of people maintain some separate financial accounts, access to which is limited, and money is not viewed as fungible across spouses.

Finally, we examine how couples finance different types of expenditure. Our pass-through estimates reveal substantially more risk-sharing for joint than for personal expenditures. One possible explanation is that spouses use different financing rules for these goods. We therefore ask respondents what share of each household expenditure they personally finance.[51]

For housing and utilities, typically the household's largest shared expenditure, payment shares closely track income shares and lie near the 45-degree line (Figure 7(b)). By contrast, respondents finance most of their own personal expenditures regardless of their income share (Figure 7(c)). The median payment share of own consumption is 100%. Respondents also finance little of their partner's personal expenditures (Appendix Figure D.4). The same couples therefore appear to apply different rules to different goods: shared expenditures are financed jointly and broadly in proportion to income, whereas personal expenditures are financed almost exclusively by the person consuming them.

[51]The survey asks which spouse ultimately finances each expenditure category, rather than which spouse makes the nominal payment (Appendix D.1).

Taken together, the survey evidence suggests that a single common budget is too coarse a description of financial arrangements among Norwegian couples. Spouses retain some control over separate resources and finance goods differently depending on whether they are consumed jointly or individually.

6.3 Limited Commitment with Separate Budgets

The single-budget limited-commitment model generates some dependence of personal consumption on the source of a household income shock but substantially larger cross-spouse responses than we observe in the data. Our survey of Norwegian households shows that spouses maintain substantial financial and budget separation, in contrast to the assumption underlying the single-budget model. We therefore ask whether separate budgets can reproduce the pass-through patterns.

We retain the limited-commitment structure of the single-budget model. Spouses continue to choose allocations jointly to maximize Pareto-weighted utility, and the Pareto weight remains persistent, changing only when participation constraints bind. We then add two departures.

First, we impose *budget separation*. Each spouse holds private assets and faces a separate budget constraint. An increase in spouse s's income expands the resources available in that spouse's budget and need not be available to the partner. Spouses can no longer make unrestricted transfers across their private budgets in the current period.

Second, we add a *financing rule* $\sigma_{s,t}$ for the joint good. Each spouse pays for the joint good g in proportion to current income. This treatment mirrors our survey evidence on the financing of housing and utility bills, in contrast to personal consumption. The contribution shares $\sigma_{s,t}$ are given by:

$$\sigma_{1,t} = \frac{y_{1,t}}{y_{1,t} + y_{2,t}}, \qquad \sigma_{2,t} = \frac{y_{2,t}}{y_{1,t} + y_{2,t}}.$$

We interpret this financing rule as a constraint on the provision of the public good. It can be understood as a fairness or no-free-riding restriction: spouses cannot rely entirely on their partner to finance joint consumption.

Household Problem: We now characterize the household problem under separate budgets. The primitives, the renegotiation process, and the three-period horizon are unchanged from the single-budget model. Because assets are now held separately, the marital state becomes $x_t = (a_{1,t}, a_{2,t}, p_{1,t}, p_{2,t}, q_t)$.

Given this state, the household solves,

$$\begin{aligned}
\max_{g_t, a_{1,t+1}, a_{2,t+1}} \quad & \mu_t \left[U^M_{1,t} + \beta \mathbb{E}_t V^M_{1,t+1}(x_{t+1}, \mu_{t+1}) \right] \\
& + (1-\mu_t) \left[U^M_{2,t} + \beta \mathbb{E}_t V^M_{2,t+1}(x_{t+1}, \mu_{t+1}) \right] \\
\text{subject to} \quad & c_{s,t} + \sigma_{s,t}\, g_t + a_{s,t+1} = y_{s,t} + (1+r) a_{s,t}, \qquad s \in \{1,2\}, \\
& \sigma_{1,t} = \frac{y_{1,t}}{y_{1,t} + y_{2,t}}, \qquad \sigma_{2,t} = 1 - \sigma_{1,t}, \\
& a_{s,t+1} \geq \underline{a}.
\end{aligned}$$

Note that we assume spouses choose personal consumption and the joint good jointly. The household problem is therefore solved cooperatively, despite the two distinct budgets.[52,53]

With separate budgets, a positive income shock to spouse 1 relaxes only spouse 1's resource constraint. As a result, spouse 1's personal consumption increases, and generally the level of joint consumption g as well. Due to the joint-good financing rule, spouse 1 finances a larger share of g as their income share increases. On the other hand, spouse 2's contribution, $\sigma_2 g$, may increase, but this increase is dampened by the rise in spouse 1's income share. Thus, even a larger g need not increase spouse 2's contribution and may even decrease it. Budget separation and the joint-good financing rule therefore limits how much of spouse 1's shock passes through to spouse 2's personal consumption. At the same time, spouse 2 still benefits from the higher g following spouse 1's shock, even without direct transfers between spouses.

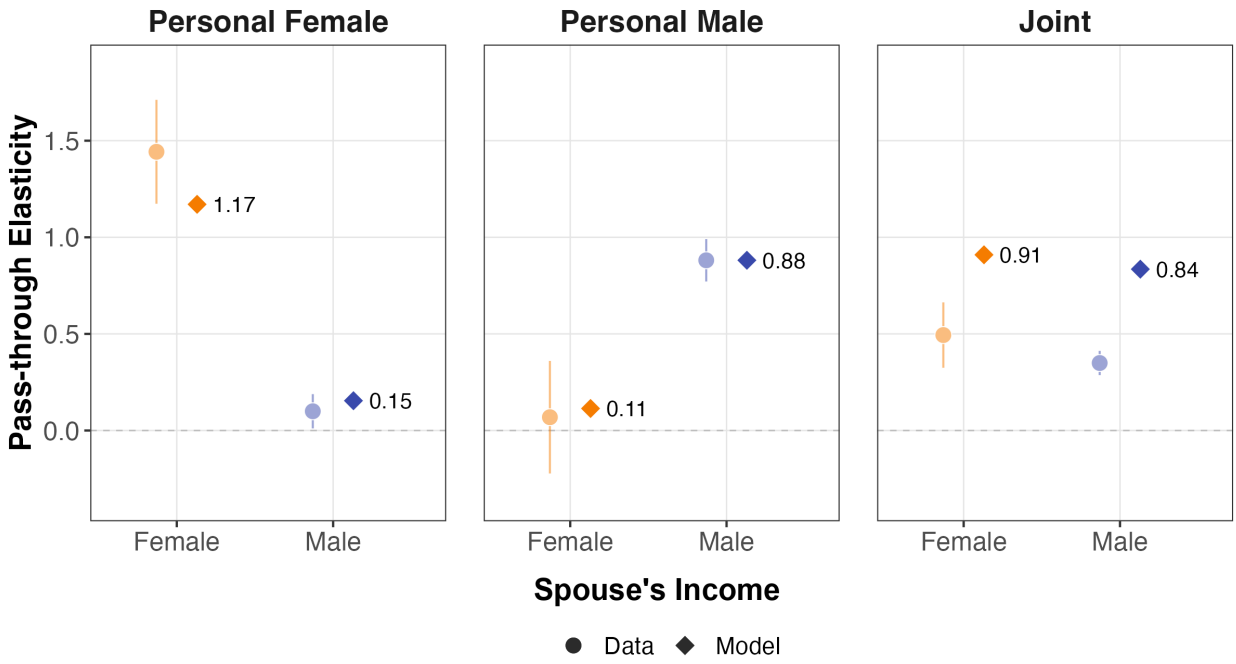
Renegotiation can further change consumption allocations. For example, if the shock increases spouse 1's Pareto weight and spouse 1 places greater weight on the joint good, g can rise by more than the increase in resources alone would imply. Joint-good shares are not renegotiated: spouse 1 still pays a larger share of the joint good, and spouse 2's contribution rises only if the increase in g outweighs the decline in their income share. If renegotiation cannot satisfy spouse 1's participation constraint, the couple separates.

Separate budgets and limited commitment are two distinct mechanisms that affect how income risk is shared within the household. Separate budgets tie own income to personal consumption even absent of renegotiation. Limited commitment adds the same friction as before: when a spouse's participation constraint binds, bargaining weight shifts toward that spouse, and consumption allocations change towards his or her preference.

[52] A non-cooperative model would also allow strategic behavior regarding the accumulation of personal assets, which can improve the spouse's outside option and affect future renegotiation. We abstract from this behavior.

[53] Additionally, on separation, these assets are still pooled and split equally: spouse s enters the single state with $a^S_{s,t} = \frac{1}{2}(a_{1,t} + a_{2,t})$.

Figure 8: Pass-Through in the Data and Separate-Budgets Limited-Commitment Model



Note: Circles show BPP-IV estimates of permanent income pass-through in the data, with 95% confidence intervals. Diamonds show simulated pass-through from the estimated limited-commitment model with separate budgets. Panels report pass-through on personal consumption of each spouse and joint expenditure; the horizontal axis is the spouse whose own income changes. Estimated parameters, targeted moments, and model elasticities are reported in Appendix Table F.5.

Results: We estimate the separate-budgets model holding preference parameters fixed. The first-stage public-good weights remain at 0.67 for women and 0.62 for men.[54] Given these weights, we re-estimate spouses' risk-aversion parameters, divorce costs, and the variance of the match-quality shock by SMM.

Figure 8 shows that allowing separate budgets sharply improves the model's ability to reproduce the personal-consumption moments. In the best-fitting model, own-income pass-through is 1.17 for women and 0.88 for men, compared with 1.44 and 0.88 in the data. Cross-spouse pass-through falls to 0.15 for women's personal consumption and 0.11 for men's, close to matching the 0.10 and 0.07 in the data. These cross-spouse responses are much smaller than in the single-budget model, which generates corresponding values of 0.77 and 0.49. Separate budgets therefore reproduce the central empirical feature that the single-budget assumption misses: personal consumption responds strongly to own permanent income and

[54]The first-order condition used to estimate these weights does not hold under separate budgets. We retain the weights for comparability and find equivalent consumption shares across the single-budget and separate-budget specifications.

very little to the partner's income.

Our estimates of pass-through to joint consumption are higher than in the data but still nearly equal.[55] Joint goods are largely household necessities, so their income elasticities likely differ from those of personal goods, a distinction the model does not include.

Why is pass-through to women's personal consumption higher than to men's? The model parameters shed light on the mechanisms behind this difference (Appendix Table F.5). Differences in risk aversion across spouses are small: the estimate risk aversion parameters are 2.45 and 2.85 for women and menr respectively. In contrast, separation costs are substantially different: $\delta_F = 2.17$ for women and $\delta_M \approx 0$ for men. These separation costs capture both the pecuniary and non-pecuniary costs of separation, which may differ by gender. Moreover, women begin marriage with lower assets (Appendix Table F.2), which limits their ability to smooth shocks under budget separation.

To quantify the impact of these parameter differences, Appendix Figure F.1 reports counterfactual simulations in which we set separately each female parameter equal to its male counterpart, holding all other parameters fixed. The baseline model difference in own-income pass-through is 0.29. Equalizing risk aversion or the utility weight for the public good leaves the gap largely unchanged, at 0.31 and 0.35 respectively. In contrast, equalizing initial income and assets reduces the gap to 0.13, while equalizing divorce costs reduces it to 0.08. These results show that gender differences in initial resources and outside options, rather than preferences, generate much of the model's predicted gap in the exposure of personal consumption to income shocks.

Overall, separate budgets accomplish what renegotiation within a single budget could not. Because personal consumption is financed primarily from own resources, own-income pass-through remains high while cross-spouse pass-through is close to zero. At the same time, income-based contributions to the joint good preserve substantial risk-sharing over joint expenditures. In this way, our model with separate budgets can reproduce the sharp distinction between personal and joint insurance observed in our empirical estimates.

7 Welfare and Policy Implications

Our empirical analysis shows that spouses share very little risk in personal consumption with each other, while displaying more income pooling in joint expenditures. What are the welfare implications of this behavior for private and public insurance? We use our estimated separate-budgets model to answer this question.

[55] This model reproduces the small gap between the point estimates of the female and male joint-good elasticities, though this gap is imprecisely estimated in the data.

First, we ask what the insurance gains from marriage implied by the model are, i.e., the gain in welfare from smoother consumption in marriage versus being single. We show empirically that women's personal consumption remains more exposed to own income variation than men's. At the same time, women also enjoy the pooling of resources from joint expenditure, for which they have a higher weight. Furthermore, women also face more permanent income risk (Appendix Figure A.5). Thus, even if their pass-through in personal consumption is higher, the welfare benefits can potentially be larger for women than for men.

We then compare the insurance gains in our estimated model with the insurance gains under full commitment and a single budget, i.e., full income pooling. The full-commitment insurance gain measures the total potential value of marriage to smooth consumption. Thus, our exercise recovers the extent to which the frictions we document decrease this surplus.

Second, we assess how these risk-sharing frictions shift the benefit of public transfers. The value of individual public insurance, such as welfare or unemployment benefits, depends on how far it replaces household insurance. When that insurance falls short of full risk sharing, the public gains can be larger. A budget-balanced income floor highlights the comparison, measuring welfare gains under our separate-budgets model and under full commitment.

7.1 Insurance Gains from Marriage

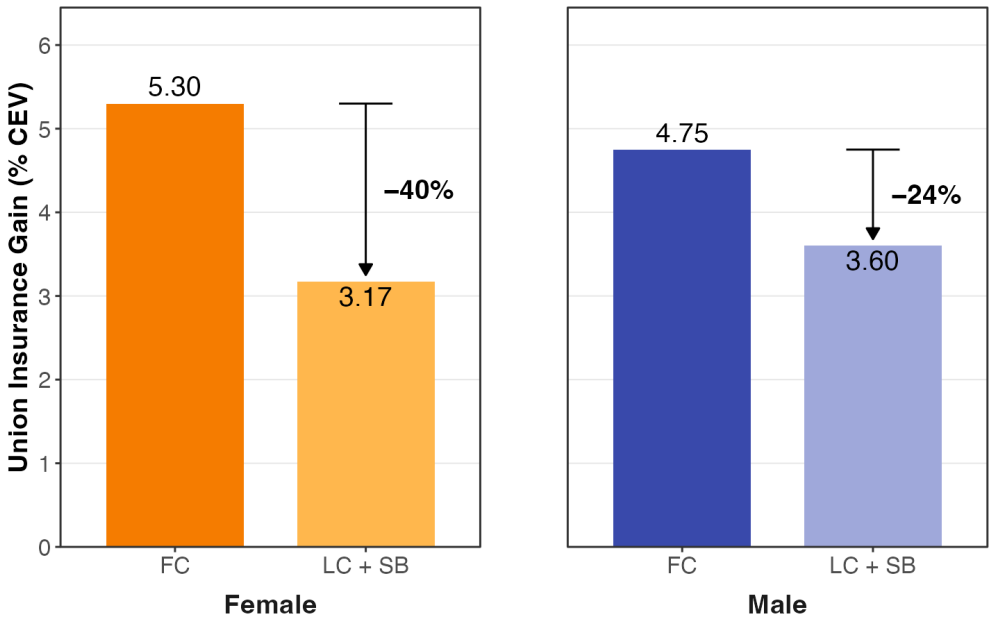
For each spouse s, we calculate the welfare loss L_s from experiencing permanent income risk σ_s in consumption-equivalent terms, keeping the other spouse's shocks active. For a spouse with ex-ante permanent income variance $\bar{\sigma}_s$, the CEV calculation converts the change from $V_s(\sigma_s = \bar{\sigma}_s)$ to $V_s(\sigma_s = 0)$ into a share of lifetime consumption.[56]

We calculate this object under three financial regimes. Under autarky (A), the individual remains single and consumes from own income and assets. Under full commitment with a single budget (FC), spouses face a single household budget and a fixed Pareto weight. We take this regime as the frictionless benchmark in which idiosyncratic income risk can be shared without participation constraints or restrictions on transfers. Finally, we calculate L_s under our estimated limited commitment with separate budgets (LC+SB) model.

The insurance gain from marriage Δ_s is the reduction in the welfare loss from income

[56] Because of Jensen's inequality, average log income rises with the variance of log income. Removing the variance while leaving average log income unchanged would lower expected income. We instead raise average log income by enough to hold expected income fixed, so the consumption-equivalent gain is the value of removing the risk at the same mean income.

Figure 9: Insurance Gains from Union



Note: The figure reports the median consumption-equivalent insurance gains from marriage, the reduction in the welfare loss from income risk relative to autarky. LC+SB is the estimated limited-commitment model with separate budgets. FC keeps those preferences and parameters and imposes full commitment with a single budget. The label next to the arrow reports the percent reduction in the insurance gain relative to FC.

risk relative to autarky,[57]

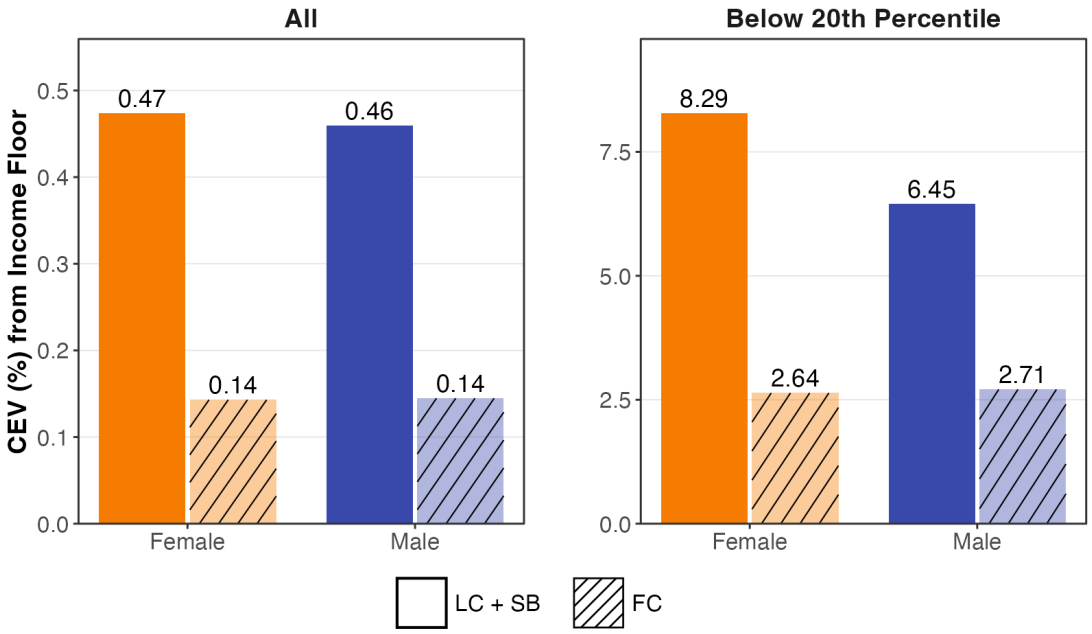
$$\Delta_s^{\mathrm{FC}} = L_s^A - L_s^{\mathrm{FC}}, \qquad \Delta_s^{\mathrm{LC+SB}} = L_s^A - L_s^{\mathrm{LC+SB}}$$

Figure 9 reports the simulated median value of Δ_s^{FC} and $\Delta_s^{\mathrm{LC+SB}}$ for each spouse. Both spouses enjoy insurance gains from marriage even in the separate-budgets model. The insurance gain for men is 14% larger than for women (3.60% vs. 3.17%). Men therefore enjoy higher welfare gains from sharing risk in marriage than women, even though women face higher variance from permanent shocks.

Under full commitment, the insurance gains are larger for women than for men: 5.30% compared with 4.75%. This difference likely reflects the estimated income process for women, together with gender differences in assets and income. If spouses shared risk fully, women would enjoy larger welfare gains than men. The risk-sharing frictions in the separate-budgets model reduce these potential gains, and the reductions are unequal: women lose 40% of the

[57]Note that our welfare exercise isolates the insurance value of marriage rather than its total surplus. Marriage may provide other benefits for spouses, including match-quality utility and gains from economies of scale that are not captured in L_s.

Figure 10: Welfare Gains from an Income Floor



Note: The figure reports the median consumption-equivalent gain from an income floor that tops up own income below the 20th percentile q_{20} of the pre-transfer distribution, $T_s = \tau \max(0,\, q_{20} - y_s)$, with replacement rate $\tau = 0.25$. The floor is financed by a tax on income above the 80th percentile, with the tax rate set so that the present value of taxes equals the present value of transfers. Solid bars are the estimated limited-commitment model with separate budgets. Hatched bars keep those preferences and parameters and impose full commitment with a single budget. The right panel restricts to median gains for spouses below the 20th percentile of own income.

full-commitment insurance gain, whereas men lose 24%. These frictions not only decrease the total insurance provided by marriage, but also change how that insurance is distributed across spouses.

7.2 Welfare Gains from Public Insurance

The value of social income insurance, such as welfare, unemployment benefits, or a negative income tax, depends on the extent to which it replaces private insurance provided by the household. When spouses pool assets and smooth shocks together, a public transfer can replace resources the other spouse would have provided, so it can crowd out household insurance (Attanasio and Ríos-Rull, 2000). If spouses maintain separate budgets and the insurance gains from marriage are smaller, public insurance crowds out less household insurance and can yield larger welfare gains.

We assess the value of an additional income floor in our separate-budgets model and compare it with its value under full risk sharing. We abstract from labor-supply responses

to the transfer and assume spouses remain attached to the labor market.[58]

The transfer is an individual top-up of own income. If income falls below the 20th percentile q_{20} of the pre-transfer income distribution, spouse s receives

$$T_s = \tau \max(0,\, q_{20} - y_s),$$

where the replacement rate τ is 0.25.[59] The floor is financed by a tax on income above the 80th percentile of the pre-transfer distribution, with the tax rate set so that the present value of taxes equals the present value of transfers.

Figure 10 shows the consumption-equivalent gains from this transfer for married spouses. Under full risk sharing, the median woman and the median man gain the same amount, about 0.14% of lifetime consumption. Under the estimated separate-budgets model, the gains are more than three times as large: about 0.47% for women and 0.46% for men. Less household insurance therefore raises the median benefit of the additional income floor. The differences are larger for spouses below the 20th percentile, the intended beneficiaries (8.29% versus 2.64% for women and 6.45% versus 2.71% for men). At the bottom, the gains are also asymmetric: women benefit more, which we do not see under full commitment, where the transfer is shared. The asymmetric gains reflect initial differences in assets: women at the bottom of the distribution have lower assets to insure privately than comparable men and thus benefit more from the additional transfer.

This exercise shows that the average and distributional welfare benefits of public insurance can change substantially with how much risk spouses share. Under separate budgets, public insurance crowds out less household insurance and raises its benefit, particularly for women with fewer resources.

8 Conclusion

We show that households do not share income risk fully or symmetrically. Within the household, personal consumption remains closely tied to own income, especially for women, and responds little to the partner's shock. Joint consumption is shared more evenly. These patterns are difficult to reconcile with models in which household resources form a single unrestricted pool. Instead, both our administrative expenditure data and our survey evidence point to financial arrangements in which spouses cooperate over shared expenditures while retaining meaningful control over personal resources. A limited-commitment model with

[58] In the separate-budgets model, we allow for renegotiation of the Pareto weight and for separation when assessing the welfare gain from this policy. Under full commitment, Pareto weights remain fixed.

[59] Other thresholds and replacement rates for the transfer give the same conclusions (Appendix Table F.6).

separate budgets and income-based contributions to joint consumption can reproduce these patterns.

Our findings have important policy implications. As we show, the extent to which spouses share risk alters the welfare benefits of public transfers. A related question concerns the joint taxation of spouses. Almost all calculations of the marginal cost of taxing an additional earner assume that spouses pool the reduction in resources under a common household budget. With budget separation, the cost incidence across spouses is a first-order question that could affect the optimality of different tax regimes.

Our empirical results highlight the importance of understanding risk sharing empirically rather than through assumptions about household behavior. Norway and the United Kingdom display strong norms of independent financial management in marriage, whereas spouses in countries with a more unitary conception of marriage could engage in greater risk sharing and income pooling. Financial arrangements can change: rising labor force participation and a growing desire for female autonomy can shift traditional financial management in marriage over time. The extent to which spouses share risk, and what that implies for the design of social insurance, are questions policy-makers and researchers have to assess empirically.

References

Abowd, J. M., Kramarz, F., and Margolis, D. N. (1999). High wage workers and high wage firms. *Econometrica*, 67(2):251–333.

Addo, F. R. and Sassler, S. (2010). Financial arrangements and relationship quality in low-income couples. *Family Relations*, 59(4):408–423.

Afzal, U., d'Adda, G., Fafchamps, M., and Said, F. (2022). Intrahousehold consumption allocation and demand for agency: A triple experimental investigation. *American Economic Journal: Applied Economics*, 14(3):400–444.

Ahn, S., Galaasen, S. M., and Mæhlum, M. (2024). The cash-flow channel of monetary policy evidence from billions of transactions. Working Paper 20, Norges Bank Research, Oslo, Norway.

Alon, T., Coskun, S., Doepke, M., Koll, D., and Tertilt, M. (2022). From Mancession to Shecession: Women's employment in regular and pandemic recessions. *NBER Macroeconomics Annual*, 36(1):83–151.

Arduini, F. (2026). Estimating intra-household inequality from time-use data. IFS Working Paper 26/24, Institute for Fiscal Studies.

Arellano, M., Blundell, R., and Bonhomme, S. (2017). Earnings and consumption dynamics: A nonlinear panel data framework. *Econometrica*, 85(3):693–734.

Ashby, K. J. and Burgoyne, C. B. (2008). Separate financial entities?: Beyond categories of money management. *Journal of Socio-Economics*, 37(2):458–480.

Ashraf, N. (2009). Spousal control and intra-household decision making: An experimental study in the Philippines. *American Economic Review*, 99(4):1245–1277.

Attanasio, O. and Davis, S. (1996). Relative wage movements and the distribution of consumption. *Journal of Political Economy*, 104(6):1227–1262.

Attanasio, O. and Ríos-Rull, J.-V. (2000). Consumption smoothing in island economies: Can public insurance reduce welfare? *European Economic Review*, 44(7):1225–1258.

Attanasio, O. P. and Lechene, V. (2002). Tests of income pooling in household decisions. *Review of Economic Dynamics*, 5(4):720–748.

Bargain, O., Lacroix, G., and Tiberti, L. (2022). Intrahousehold resource allocation and individual poverty: Assessing collective model predictions using direct evidence on sharing. *The Economic Journal*, 132(643):865–905.

Barigozzi, F., Di Timoteo, C., and Monfardini, C. (2023). The gender gaps in time-use within Italian households during 2002–2014. *Italian Economic Journal*, 9(3):1263–1296.

Beblo, M. and Robledo, J. R. (2008). The wage gap and the leisure gap for double-earner couples. *Journal of Population Economics*, 21(2):281–304.

Bellue, S., Doepke, M., and Tertilt, M. (2026). Private information in the family. Discussion Paper 738, CRC TR 224.

Bergstrom, T. C., Blume, L. E., and Varian, H. R. (1986). On the private provision of public goods. *Journal of Public Economics*, 29(1):25–49.

Bertrand, M., Kamenica, E., and Pan, J. (2015). Gender identity and relative income within households. *The Quarterly Journal of Economics*, 130(2):571–614.

Blau, D. M. and Goodstein, R. M. (2016). Commitment in the household: Evidence from the effect of inheritances on the labor supply of older married couples. *Labour Economics*, 42:123–137.

Blundell, R., Chiappori, P.-A., and Meghir, C. (2005). Collective labor supply with children. *Journal of Political Economy*, 113(6):1277–1306.

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.

Browning, M., Bourguignon, F., Chiappori, P.-A., and Lechene, V. (1994). Income and outcomes: A structural model of intrahousehold allocation. *Journal of Political Economy*, 102(6):1067–1096.

Browning, M., Chiappori, P.-A., and Lechene, V. (2010). Distributional effects in household models: Separate spheres and income pooling. *The Economic Journal*, 120(545):786–799.

Calvi, R. (2020). Why are older women missing in India? The age profile of bargaining power and poverty. *Journal of Political Economy*, 128(7):2453–2501.

Calvi, R., Penglase, J., Tommasi, D., and Wolf, A. (2023). The more the poorer? Resource sharing and scale economies in large families. *Journal of Development Economics*, 160:102986.

Chiappori, P.-A. and Mazzocco, M. (2017). Static and intertemporal household decisions. *Journal of Economic Literature*, 55(3):985–1045.

Chiappori, P.-A. and Meghir, C. (2015). Intrahousehold inequality. In Atkinson, A. B. and Bourguignon, F., editors, *Handbook of Income Distribution*, volume 2, chapter 16, pages 1369–1418. Elsevier, Amsterdam.

Choukhmane, T., Goodman, L., and O'Dea, C. (2025). Efficiency in household decision-making: Evidence from the retirement savings of US couples. *American Economic Review*, 115(5):1485–1519.

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

De Giorgi, G., Frederiksen, A., and Pistaferri, L. (2020). Consumption network effects. *Review of Economic Studies*, 87(1):130–163.

Dubois, P. and Ligon, E. (2009). Nutrition and risk sharing within the household. Discussion Paper 7574, Centre for Economic Policy Research.

Duflo, E. and Udry, C. (2004). Intrahousehold resource allocation in Côte d'Ivoire: Social norms, separate accounts and consumption choices. Working Paper 10498, National Bureau of Economic Research.

Dunbar, G. R., Lewbel, A., and Pendakur, K. (2013). Children's resources in collective households: Identification, estimation, and an application to child poverty in Malawi. *American Economic Review*, 103(1):438–471.

Eika, L., Mogstad, M., and Vestad, O. L. (2020). What can we learn about household consumption expenditure from data on income and assets? *Journal of Public Economics*, 189:104163.

Fagereng, A., Guiso, L., and Pistaferri, L. (2017). Firm-related risk and precautionary saving response. *American Economic Review*, 107(5):393–397.

Fukuda, K. (2024). Job loss, consumption insurance, and household time allocation. *Labour Economics*, 91:102645.

Galaasen, S., Kostøl, A. R., and Simmons, M. (2026). Firm pay and consumption inequality. *Working Paper*.

Ganong, P., Jones, D., Noel, P., Greig, F., Farrell, D., and Wheat, C. (2025). Liquid wealth and consumption smoothing of typical labor income shocks. *Working Paper*.

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

Hall, R. E. (1978). Stochastic implications of the life cycle-permanent income hypothesis: Theory and evidence. *Journal of Political Economy*, 86(6):971–987.

Hancock, K., Lafortune, J., and Low, C. (2025). Winning the bread and baking it too: Gendered frictions in the allocation of home production. NBER Working Paper 33393, National Bureau of Economic Research.

Kaplan, G. and Violante, G. (2010). How much consumption insurance beyond self-insurance? *American Economic Journal: Macroeconomics*, 2(4):53–87.

Kim, O. S. (2024). Credit and the family: The economic consequences of closing the credit gap of U.S. couples. *Working Paper*.

Lechene, V., Pendakur, K., and Wolf, A. (2022). Ordinary least squares estimation of the intrahousehold distribution of expenditure. *Journal of Political Economy*, 130(3):681–731.

Ligon, E., Thomas, J. P., and Worrall, T. (2002). Informal insurance arrangements with limited commitment: Theory and evidence from village economies. *Review of Economic Studies*, 69(1):209–244.

Lise, J. and Seitz, S. (2011). Consumption inequality and intra-household allocations. *Review of Economic Studies*, 78(1):328–355.

Lise, J. and Yamada, K. (2019). Household sharing and commitment: Evidence from panel data on individual expenditures and time use. *Review of Economic Studies*, 86(5):2184–2219.

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

Lundberg, S. and Pollak, R. A. (1993). Separate spheres bargaining and the marriage market. *Journal of Political Economy*, 101(6):988–1010.

Lundberg, S. J., Pollak, R. A., and Wales, T. J. (1997). Do husbands and wives pool their resources? Evidence from the United Kingdom child benefit. *Journal of Human Resources*, 32(3):463–480.

Mazzocco, M. (2007). Household intertemporal behaviour: A collective characterization and a test of commitment. *Review of Economic Studies*, 74(3):857–895.

Meghir, C. and Pistaferri, L. (2004). Income variance dynamics and heterogeneity. *Econometrica*, 72(1):1–32.

Norges Bank (2019). Retail payment services 2018. Norges Bank Papers 1, Norges Bank, Oslo, Norway.

Pahl, J. (1995). His money, her money: Recent research on financial organisation in marriage. *Journal of Economic Psychology*, 16(3):361–376.

Patterson, C. (2023). The matching multiplier and the amplification of recessions. *American Economic Review*, 113(4):982–1012.

Robinson, J. (2012). Limited insurance within the household: Evidence from a field experiment in Kenya. *American Economic Journal: Applied Economics*, 4(4):140–164.

Schaner, S. (2015). Do opposites detract? Intrahousehold preference heterogeneity and inefficient strategic savings. *American Economic Journal: Applied Economics*, 7(2):135–174.

Stephens, Jr., M. (2002). Worker displacement and the added worker effect. *Journal of Labor Economics*, 20(3):504–537.

Theloudis, A., Velilla, J., Chiappori, P.-A., Giménez-Nadal, J. I., and Molina, J. A. (2025). Commitment and the dynamics of household labour supply. *The Economic Journal*, 135(665):354–386.

Udry, C. (1996). Gender, agricultural production, and the theory of the household. *Journal of Political Economy*, 104(5):1010–1046.

Voena, A. (2015). Yours, mine, and ours: Do divorce laws affect the intertemporal behavior of married couples? *American Economic Review*, 105(8):2295–2332.

Vogler, C., Lyonette, C., and Wiggins, R. D. (2008). Money, power and spending decisions in intimate relationships. *The Sociological Review*, 56(1):117–143.

YouGov (2024). What is the fairest way for couples to split household bills? YouGov Survey Results. Fieldwork: July 21, 2023–February 5, 2024. Sample of GB adults living with a romantic partner.

Supplementary Appendix

A Additional Figures . . . 53

B Additional Tables . . . 61

C Data Appendix . . . 69

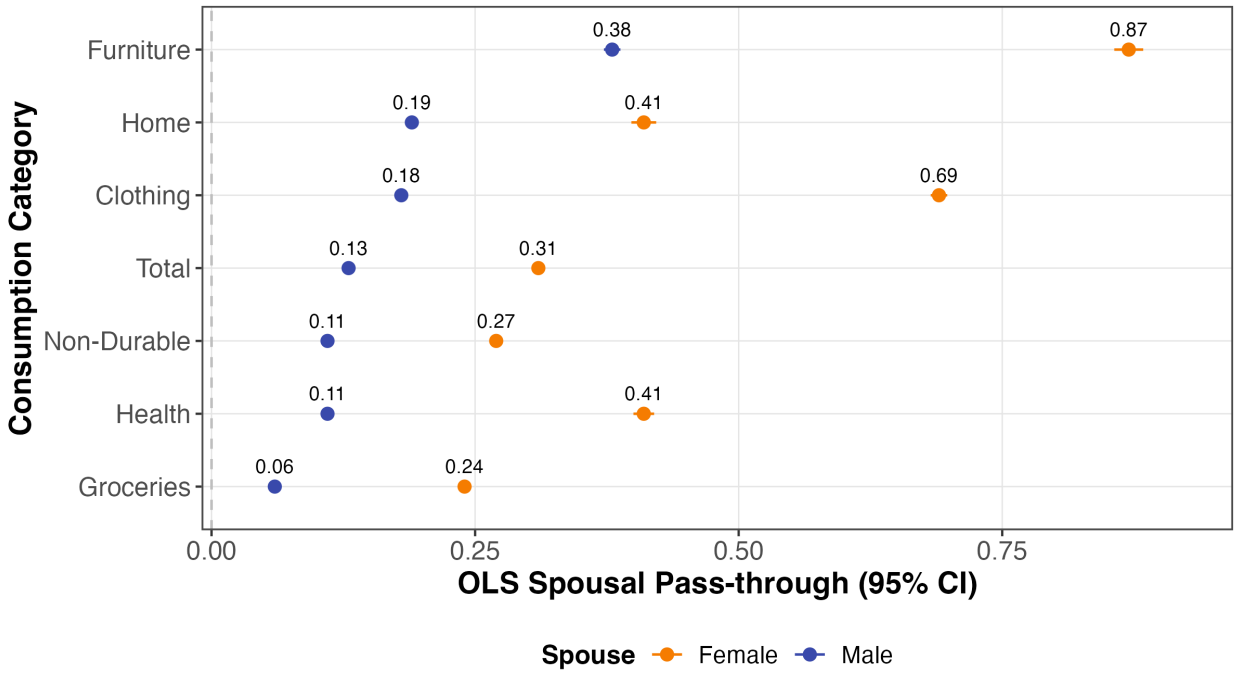
D Survey Appendix . . . 76

E Identification of the Pass-Through Parameters . . . 82

F Solution and Estimation of Household Models . . . 86

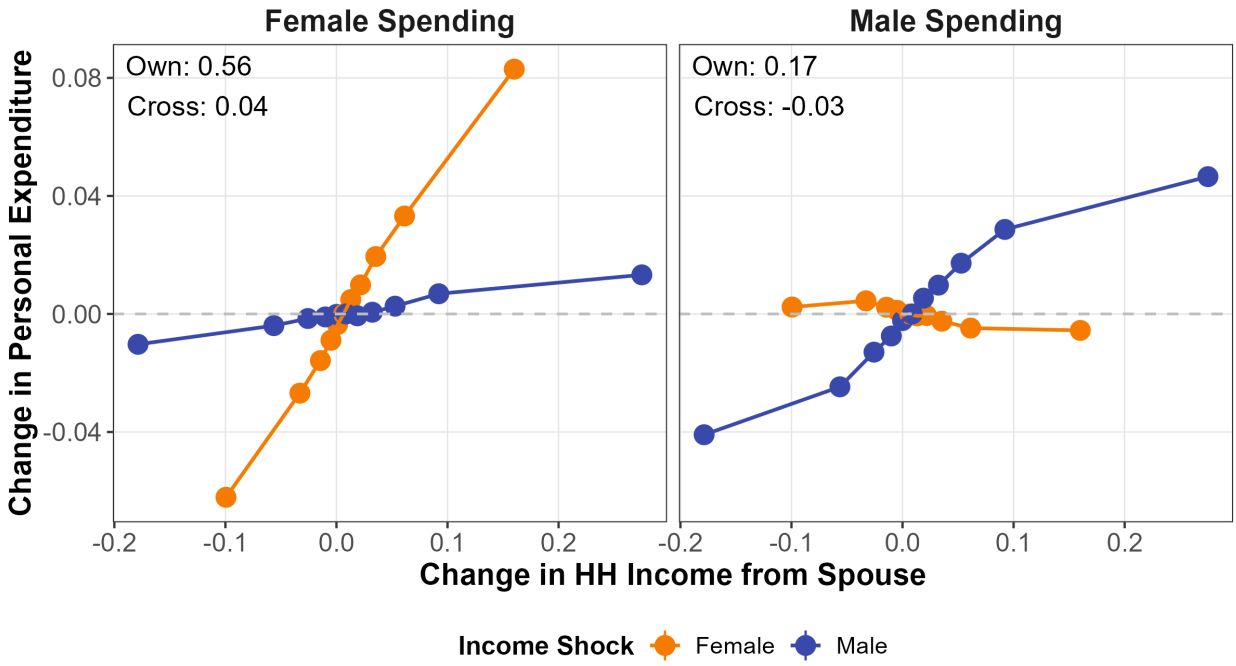
A Additional Figures

Figure A.1: OLS Pass-Through by Consumption Category



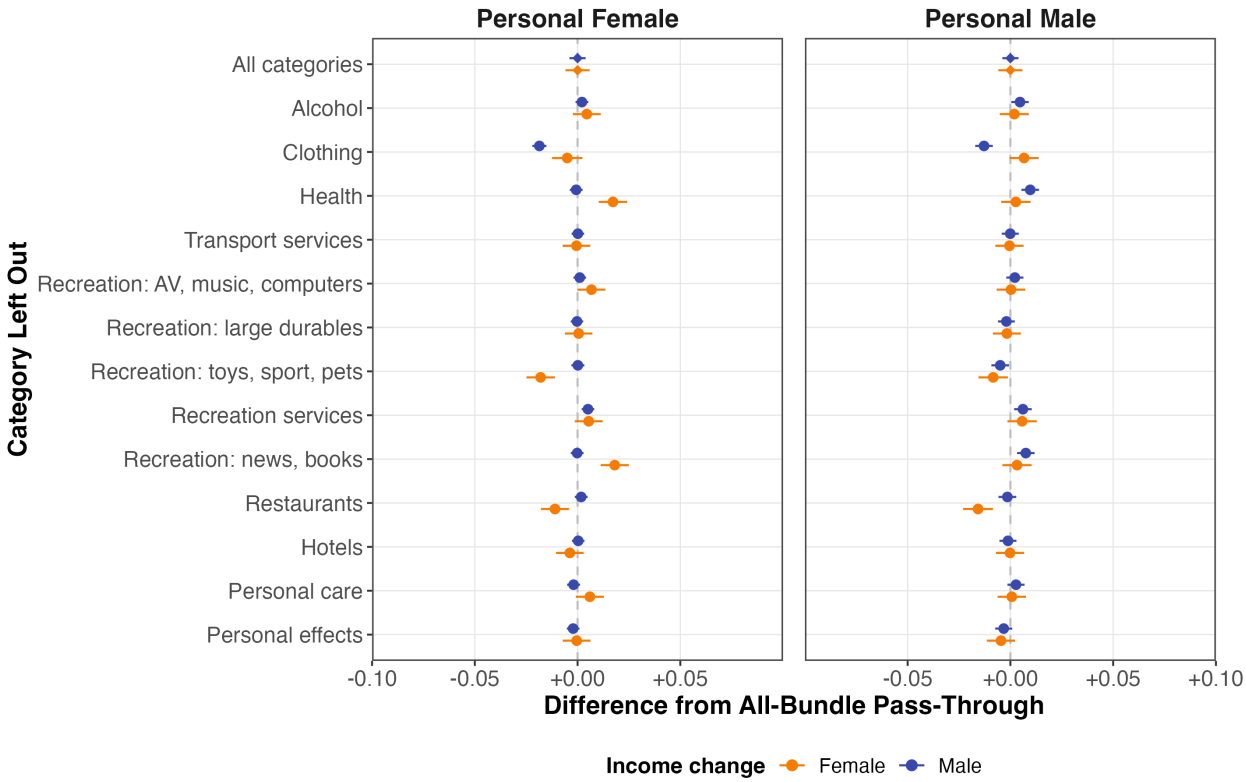
Note: This figure shows OLS estimates of the pass-through of changes in household income originating from each spouse to household expenditure by category. These changes enter jointly, and controls are those in equation (1). Horizontal bars show 95% confidence intervals; standard errors are clustered at the household level. Home expenditures are utilities and housing maintenance; rent is not imputed.

Figure A.2: Changes in Clothing-Imputed Personal Consumption and Spousal Income



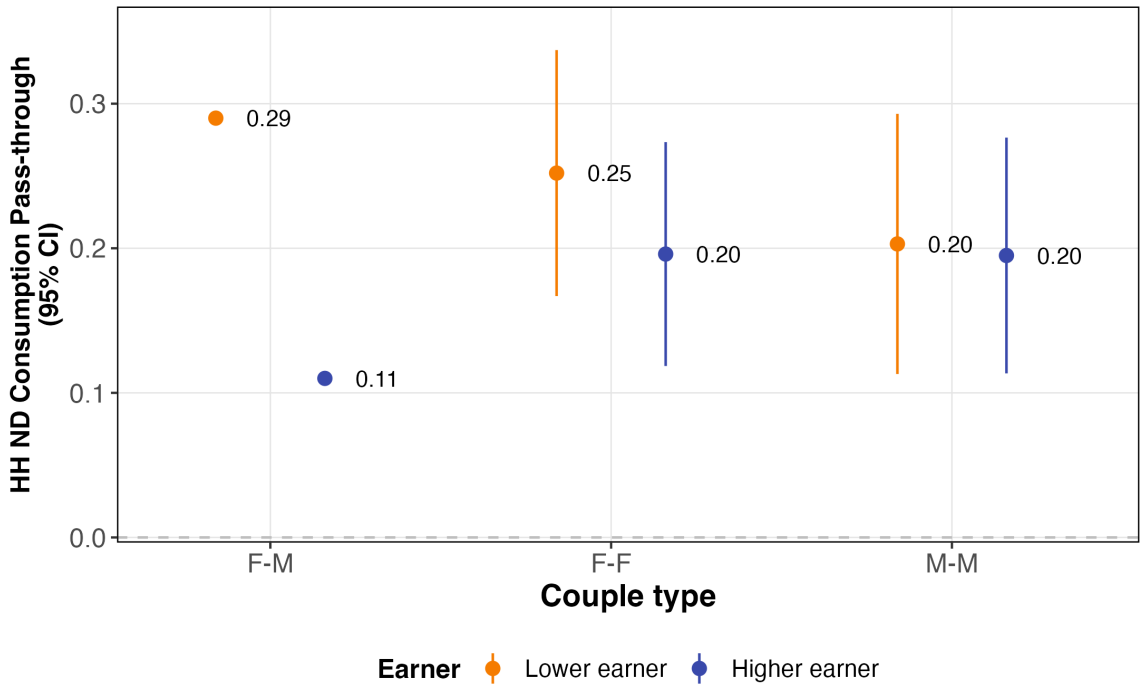
Note: Each panel shows a binscatter of changes in log clothing-imputed personal consumption against household income changes from each spouse. The left panel uses female personal consumption as the outcome and the right panel male personal consumption; household income changes from each spouse enter jointly in both panels, controlling for the covariates in equation (1). Standard errors are clustered at the household level.

Figure A.3: Leave-One-Out Robustness of Personal Expenditure Pass-Through



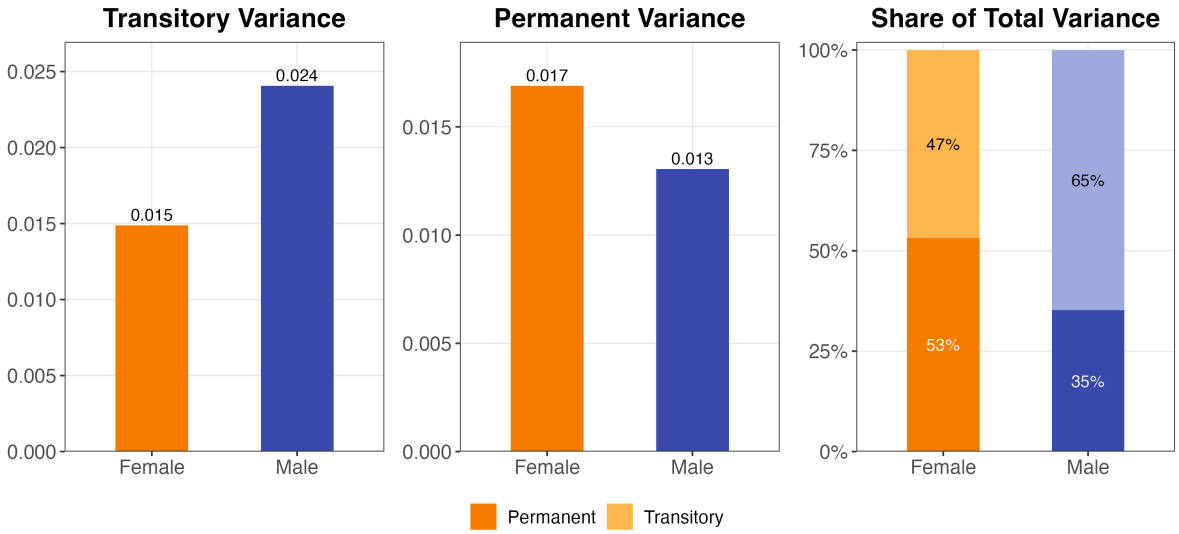
Note: Each point is the OLS pass-through from equation (1) after omitting one personal-expenditure category, minus the pass-through estimated on the full personal bundle. The left panel is female personal expenditure and the right panel is male personal expenditure. Female and male markers are the coefficients on changes in household income originating from each spouse, which enter jointly. Circles leave out the category on the vertical axis. Diamonds are the all-bundle estimates, normalized to zero. Vertical bars are 95% confidence intervals for these pass-through coefficients, recentered on the all-bundle estimate. Category labels follow Appendix Table C.2. Standard errors are clustered at the household level.

Figure A.4: OLS Pass-Through for Lower and Higher Earners, by Couple Type



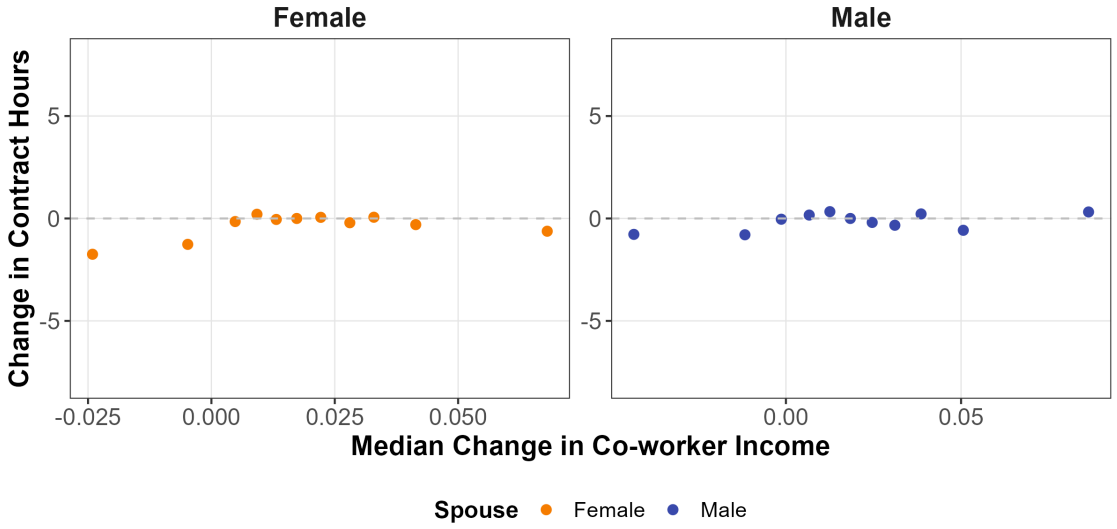
Note: This figure shows OLS estimates of the pass-through of income changes to household nondurable consumption for the lower- and higher-earning spouse. The female–male (F–M) sample includes only couples in which the woman earns on average less than the man. Female–female (F–F) and male–male (M–M) couples are ranked by average income share into a lower and a higher earner. The lower earner's average share of household income is similar in heterosexual and same-sex couples, at about 0.43. Lower- and higher-earner income changes enter jointly, controlling for the covariates in equation (1). Vertical bars show 95% confidence intervals for the difference between each same-sex estimate and the F–M estimate for the same earner rank. Standard errors are clustered at the household level.

Figure A.5: Decomposition of Income Growth into Permanent and Transitory Components, by Gender



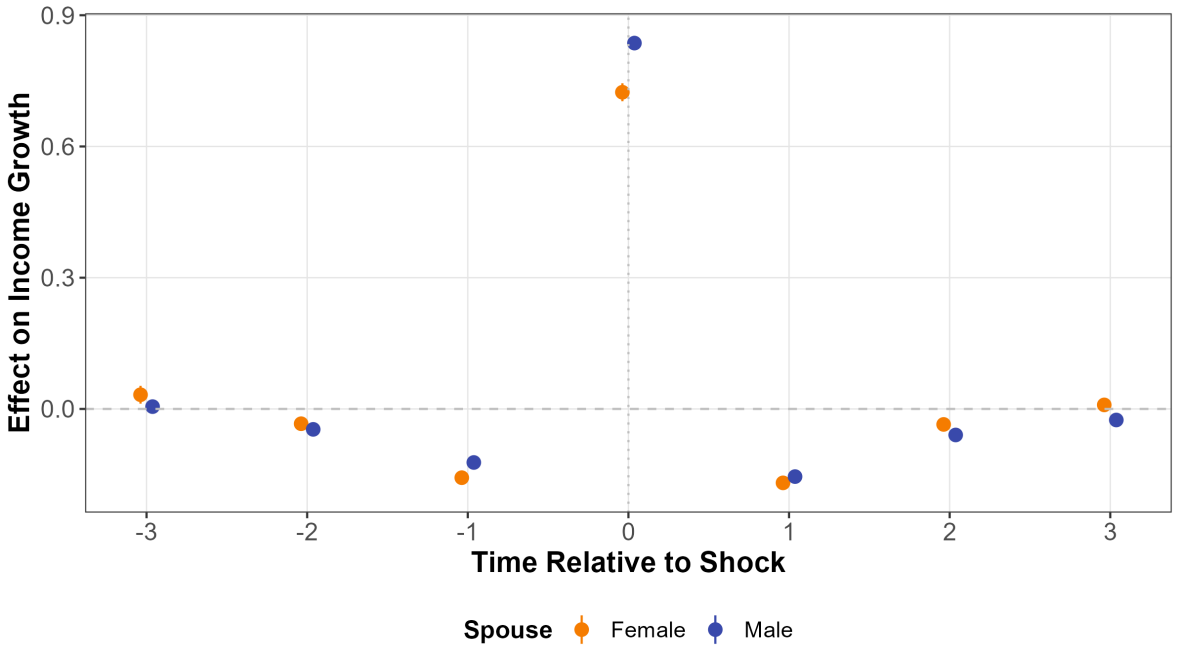
Note: The figure reports a Blundell, Pistaferri, and Preston (2008) decomposition of income-growth variance for women and men on the overall sample. The permanent component is a random walk and the transitory component is an MA(1) process. The left panel is the total transitory component and the middle panel is the permanent variance. The right panel reports each component as a share of total income-growth variance.

Figure A.6: Contract Hours around Co-Worker Income Shocks



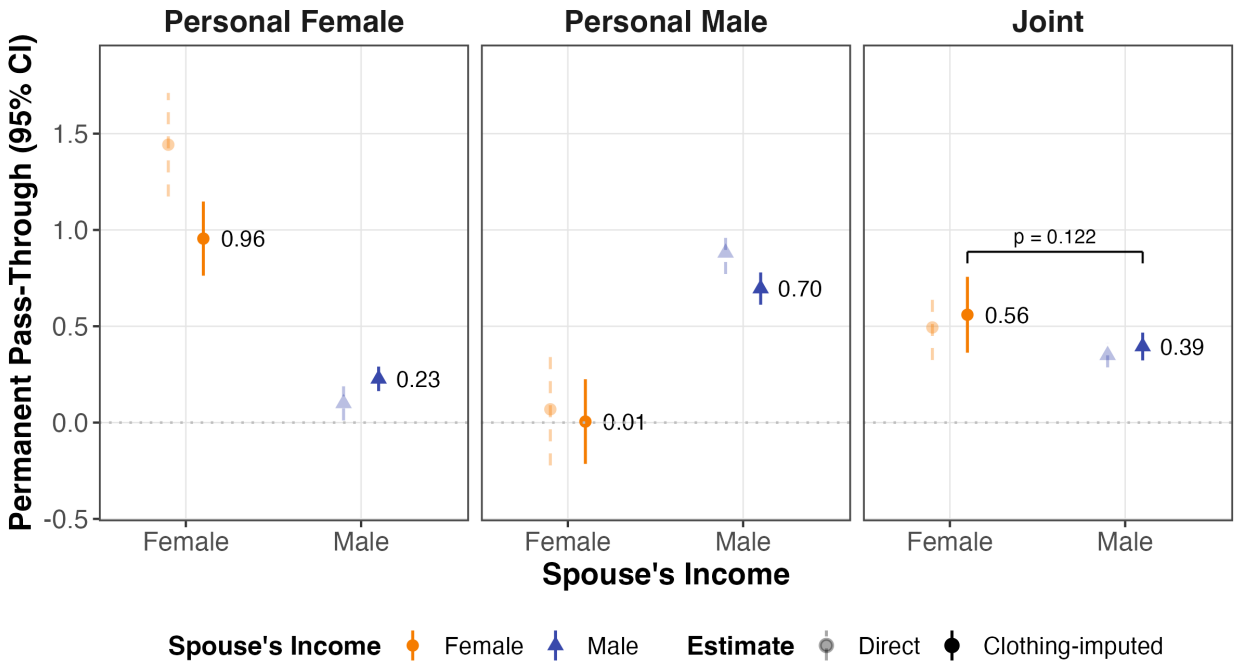
Note: Each panel is a binned scatter of the change in contract hours against the leave-one-out median co-worker income change, controlling for the covariates in equation (1) except the hours transition controls. The left panel is female spouses and the right panel is male spouses. The outcome is contract hours, not worked hours.

Figure A.7: Own Income Growth around Co-Worker Income Shocks



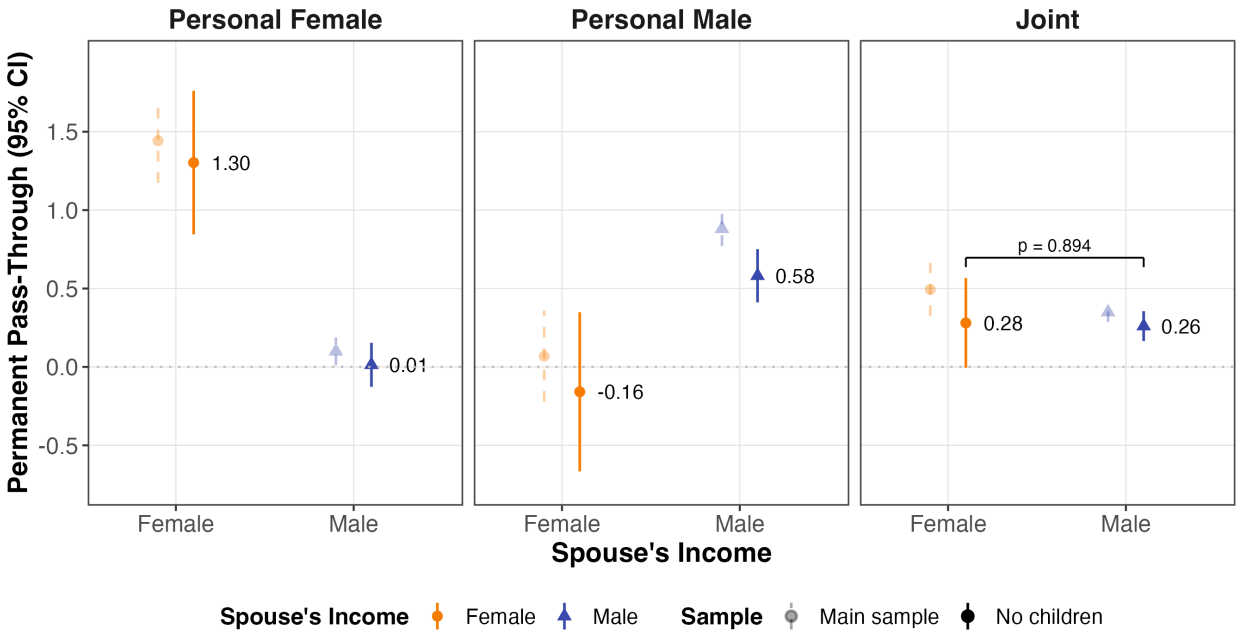
Note: The figure reports local projections of log own income growth on leads and lags of the leave-one-out co-worker income changes, separately for female and male spouses, controlling for the covariates in equation (1). Vertical bars show 95% confidence intervals. The coefficient at $t = 0$ is the contemporaneous response. Negative effects at $t = -1$ and $t = -2$ come from the transitory component. Significant effects at $t = -2$ and $t = 2$ provide evidence of an MA(1) process. Anticipatory effects would break the symmetry of the lead and lag coefficients at $t = \pm 1$ and $t = \pm 2$.

Figure A.8: Permanent Pass-Through with Clothing-Imputed Personal Consumption



Note: This figure shows BPP-IV estimates of the pass-through of permanent income shocks to personal consumption of women, personal consumption of men, and joint expenditures, for the BPP-IV sample. Each spouse's shocks are scaled by pre-period budget shares as in equation (2). Light markers use direct personal expenditure and dark markers use clothing-imputed personal consumption. Joint consumption is residual household expenditure under the corresponding personal-consumption measure. Appendix Table C.2 reports the expenditure classification. The horizontal axis identifies the spouse receiving the income shock. Vertical bars show 95% confidence intervals; standard errors are clustered at the household level. The p-value corresponds to the difference between the male and female pass-through estimates. Point estimates and standard errors, including for transitory shocks, are reported in Appendix Table B.5.

Figure A.9: Permanent Pass-Through among Couples without Children



Note: This figure shows BPP-IV estimates of the pass-through of permanent income shocks to personal consumption of women, personal consumption of men, and joint expenditures, for the BPP-IV sample. Each spouse's shocks are scaled by pre-period budget shares as in equation (2). Light markers use the main sample and dark markers use couples without children. Personal consumption is direct expenditure in personal categories. Appendix Table C.2 reports the expenditure classification. The horizontal axis identifies the spouse receiving the income shock. Vertical bars show 95% confidence intervals; standard errors are clustered at the household level. The p-value corresponds to the difference between the male and female pass-through estimates. Point estimates and standard errors are reported in Appendix Table B.7.

B Additional Tables

Table B.1: OLS Pass-Through Heterogeneity

	Total		Non-durable		Female Personal		Male Personal	
	Female	Male	Female	Male	Female	Male	Female	Male
Overall	0.309	0.133	0.266	0.111	0.932	0.016	−0.015	0.311
	(0.002)	(0.001)	(0.002)	(0.001)	(0.003)	(0.002)	(0.003)	(0.002)
Asset Quintile								
Q1	0.502	0.322	0.430	0.275	1.474	0.076	0.004	0.772
	(0.004)	(0.003)	(0.004)	(0.004)	(0.009)	(0.006)	(0.009)	(0.007)
Q2	0.390	0.229	0.323	0.189	1.147	0.060	−0.018	0.543
	(0.004)	(0.003)	(0.005)	(0.003)	(0.008)	(0.005)	(0.008)	(0.006)
Q3	0.331	0.179	0.281	0.146	0.952	0.032	0.004	0.409
	(0.004)	(0.003)	(0.005)	(0.003)	(0.007)	(0.004)	(0.008)	(0.005)
Q4	0.267	0.124	0.232	0.100	0.798	0.020	−0.006	0.291
	(0.004)	(0.002)	(0.005)	(0.003)	(0.007)	(0.004)	(0.008)	(0.005)
Q5	0.148	0.063	0.137	0.057	0.485	0.004	−0.003	0.134
	(0.004)	(0.001)	(0.004)	(0.002)	(0.006)	(0.002)	(0.007)	(0.003)
Female Income Share								
$[0.00, 0.35]$	0.331	0.113	0.304	0.089	1.391	0.012	−0.021	0.198
	(0.004)	(0.002)	(0.004)	(0.002)	(0.008)	(0.002)	(0.007)	(0.003)
$(0.35, 0.43]$	0.323	0.135	0.278	0.112	1.004	0.018	−0.017	0.314
	(0.004)	(0.002)	(0.004)	(0.002)	(0.007)	(0.003)	(0.007)	(0.004)
$(0.43, 0.49]$	0.316	0.145	0.263	0.127	0.825	0.012	−0.003	0.380
	(0.004)	(0.002)	(0.004)	(0.003)	(0.007)	(0.004)	(0.007)	(0.005)
$(0.49, 1.00]$	0.276	0.168	0.226	0.152	0.595	0.019	−0.027	0.537
	(0.003)	(0.003)	(0.004)	(0.003)	(0.006)	(0.004)	(0.007)	(0.006)
Civil Status								
Married	0.296	0.122	0.259	0.102	0.920	0.016	−0.009	0.283
	(0.002)	(0.001)	(0.002)	(0.001)	(0.004)	(0.002)	(0.004)	(0.002)
Cohabiting	0.337	0.166	0.278	0.138	0.956	0.016	−0.031	0.396
	(0.003)	(0.002)	(0.003)	(0.002)	(0.006)	(0.003)	(0.006)	(0.004)
Children								
With children	0.331	0.145	0.285	0.121	1.015	0.023	−0.025	0.334
	(0.002)	(0.001)	(0.003)	(0.002)	(0.004)	(0.002)	(0.004)	(0.003)
No children	0.274	0.119	0.236	0.100	0.799	0.010	0.006	0.284
	(0.003)	(0.001)	(0.003)	(0.002)	(0.006)	(0.002)	(0.006)	(0.003)

Note: Each cell is the OLS pass-through of the change in household income from each spouse to the indicated expenditure. Changes in household income originating from each spouse enter jointly, controlling for the covariates in equation (1). Total is all household expenditures and non-durable is household non-durable consumption. Female Personal and Male Personal are each spouse's direct personal expenditure. Rows other than Overall are estimated on the indicated subsample. Asset rows are household liquid savings quintiles. Female income share bins are quartiles of the female spouse's average household income share. Standard errors, clustered at the household level, are in parentheses.

Table B.2: First Stage Table

	Δ Female Income	Δ Male Income
Δ Co-worker Income F	0.329	−0.004
	(0.005)	(0.004)
Δ Co-worker Income M	−0.003	0.541
	(0.001)	(0.005)
Observations	1,609,514	1,609,514
R^2	0.093	0.066
Wald (1st stage)	2,692	5,066

Note: The table reports OLS regressions of the change in household income originating from each spouse on the leave-one-out median co-worker income changes of each spouse, for the BPP-IV sample. Both co-worker changes enter jointly, controlling for the covariates in equation (1). Each Wald statistic jointly tests the two instruments in that spouse's regression. Standard errors, clustered at the household level, are in parentheses.

Table B.3: Summary Statistics: OLS and BPP-IV Samples

	OLS sample		BPP-IV sample	
	Women	Men	Women	Men
Age	44.6	47.2	44.8	47.2
After-tax income	327,327	475,790	352,833	495,149
Share of household after-tax income	0.42	0.58	0.43	0.57
Labor force participation	0.86	0.87	1.00	1.00
Total expenditure	237,396	356,880	245,125	347,300
Share of household expenditure	0.42	0.58	0.43	0.57
Groceries	0.54	0.46	0.55	0.45
Home	0.31	0.69	0.31	0.69
Clothing	0.73	0.27	0.74	0.26
N years per couple	9.4		4.1	
N couples	821,964		413,953	
N observations	7,716,754		1,687,221	

Note: The table reports sample means for heterosexual married and cohabiting couples. Monetary amounts are deflated to 2015 NOK. Income and expenditure shares are each spouse's share of the household total. Labor-force participation is the share of observations with positive labor earnings. Home is utilities and housing maintenance; rent is not imputed. The BPP-IV sample requires both spouses to be employed and to have a valid co-worker instrument. Both working shares therefore equal one, while couple spells are shorter and income is slightly higher than in the OLS sample.

Table B.4: Pass-Through of Permanent and Transitory Income Shocks to Household Consumption

	HH Total	HH N-D	Groceries	Home	Furniture	Health	Clothing
Permanent							
Male Shock	0.383	0.291	0.077	0.518	0.999	0.245	0.756
	(0.027)	(0.032)	(0.033)	(0.096)	(0.104)	(0.073)	(0.054)
Female Shock	0.570	0.513	0.148	0.902	1.381	0.670	1.435
	(0.073)	(0.088)	(0.090)	(0.269)	(0.292)	(0.211)	(0.151)
p-value Diff	0.016	0.017	0.457	0.177	0.214	0.056	< 0.001
Transitory							
Male Shock	0.235	0.282	0.215	0.496	0.405	0.369	0.003
	(0.062)	(0.073)	(0.070)	(0.225)	(0.244)	(0.176)	(0.120)
Female Shock	0.235	0.034	0.227	0.291	0.982	0.079	0.205
	(0.104)	(0.126)	(0.124)	(0.400)	(0.427)	(0.313)	(0.211)
N	1,609,173	1,609,138	1,607,827	1,605,905	1,600,588	1,603,388	1,600,345
p-value OverID	0.063	0.200	0.482	0.908	0.767	0.456	0.558

Note: Two-step GMM estimates of the BPP-IV pass-through of permanent and transitory income shocks to aggregate household expenditure, for the BPP-IV sample. Each spouse's shocks are scaled by pre-period budget shares as in equation (2). HH Total is total household expenditure, HH N-D is household non-durable expenditure, and the remaining columns are category-specific household expenditures. Home is utilities and housing maintenance; rent is not imputed. Standard errors, in parentheses, are clustered at the household level. N is the number of couple-year observations. The overidentification p-value is from Hansen's J test. The p-value of the difference tests equality of permanent pass-through for male and female income shocks.

Table B.5: Pass-Through of Permanent and Transitory Income Shocks

	Direct			Implied		
	Personal M	Personal F	Joint	Personal M	Personal F	Joint
Permanent						
Male Shock	0.881	0.100	0.349	0.696	0.227	0.395
	(0.056)	(0.045)	(0.032)	(0.043)	(0.032)	(0.037)
Female Shock	0.069	1.443	0.494	0.005	0.955	0.560
	(0.149)	(0.137)	(0.086)	(0.112)	(0.098)	(0.100)
p-value Diff	< 0.001	< 0.001	0.115	< 0.001	< 0.001	0.122
Transitory						
Male Shock	0.220	0.255	0.246	-0.075	0.045	0.283
	(0.127)	(0.106)	(0.074)	(0.093)	(0.072)	(0.082)
Female Shock	0.080	0.510	0.197	-0.079	0.209	0.299
	(0.219)	(0.184)	(0.126)	(0.167)	(0.132)	(0.145)
N	1,602,747	1,598,956	1,609,160	1,402,070	1,402,070	1,402,070
p-value OverID	0.071	0.548	0.160	0.238	0.764	0.139

Note: Two-step GMM estimates of the BPP-IV pass-through of permanent and transitory income shocks to personal consumption of women, personal consumption of men, and joint expenditures, for the BPP-IV sample. Each spouse's shocks are scaled by pre-period budget shares as in equation (2). Direct personal consumption is observed personal expenditure; implied personal consumption is predicted from clothing expenditure using the pre-union clothing relationship. Joint consumption is residual household expenditure under the corresponding personal-consumption measure. Standard errors, in parentheses, are clustered at the household level. N is the number of couple-year observations. The overidentification p-value is from Hansen's J test. The p-value of the difference tests equality of permanent pass-through for male and female income shocks.

Table B.6: Comparison of OLS, IV, and BPP-IV Pass-Through

	OLS		IV		BPP-IV			
					Permanent		Transitory	
	M	F	M	F	M	F	M	F
HH Total	0.133	0.322	0.260	0.313	0.383	0.570	0.235	0.235
	(0.002)	(0.004)	(0.012)	(0.026)	(0.027)	(0.073)	(0.062)	(0.104)
HH Non-Durable	0.107	0.285	0.217	0.231	0.291	0.513	0.282	0.034
	(0.003)	(0.005)	(0.014)	(0.033)	(0.032)	(0.088)	(0.073)	(0.126)
Personal M	0.297	−0.017	0.529	0.039	0.881	0.069	0.220	0.080
	(0.005)	(0.009)	(0.025)	(0.057)	(0.056)	(0.149)	(0.127)	(0.219)
Personal F	0.026	0.872	0.110	0.782	0.100	1.443	0.255	0.510
	(0.004)	(0.008)	(0.020)	(0.049)	(0.045)	(0.137)	(0.106)	(0.184)
Joint	0.125	0.291	0.244	0.269	0.349	0.494	0.246	0.197
	(0.003)	(0.005)	(0.014)	(0.032)	(0.032)	(0.086)	(0.074)	(0.126)

Note: Each entry is the pass-through of the change in household income from the spouse's source into the row outcome, estimated on the BPP-IV sample. OLS and IV regress consumption growth on the change in household income from each spouse's source. IV instruments that change with the leave-one-out median co-worker income change. BPP-IV separates the permanent and transitory parts of that change in the baseline two-step GMM specification. Each spouse's shocks are scaled by pre-period budget shares as in equation (2). Personal consumption is direct personal expenditure. M and F denote male and female income shocks. Standard errors, in parentheses, are clustered at the household level.

Table B.7: Causal Permanent Pass-Through Heterogeneity

	HH total		Personal M		Personal F		Joint	
	Male	Female	Male	Female	Male	Female	Male	Female
Overall	0.383	0.570	0.881	0.069	0.100	1.443	0.349	0.494
	(0.027)	(0.073)	(0.056)	(0.149)	(0.045)	(0.137)	(0.032)	(0.086)
Liquid Asset Tercile								
T1	0.393	0.661	1.276	0.416	0.042	1.834	0.321	0.524
	(0.055)	(0.138)	(0.120)	(0.298)	(0.101)	(0.283)	(0.063)	(0.159)
T2	0.459	0.581	1.141	0.262	0.228	1.459	0.390	0.470
	(0.048)	(0.123)	(0.098)	(0.252)	(0.079)	(0.220)	(0.056)	(0.147)
T3	0.307	0.485	0.503	−0.247	0.060	1.218	0.311	0.462
	(0.042)	(0.118)	(0.080)	(0.240)	(0.066)	(0.217)	(0.050)	(0.143)
Female Income Share								
T1	0.363	0.602	0.751	−0.159	0.136	2.425	0.335	0.530
	(0.033)	(0.181)	(0.066)	(0.365)	(0.052)	(0.395)	(0.039)	(0.212)
T2	0.409	0.575	1.080	−0.228	0.016	1.471	0.371	0.516
	(0.052)	(0.141)	(0.110)	(0.283)	(0.091)	(0.259)	(0.060)	(0.168)
T3	0.367	0.513	1.002	0.314	0.038	0.917	0.324	0.434
	(0.080)	(0.087)	(0.160)	(0.182)	(0.130)	(0.152)	(0.092)	(0.103)
Civil Status								
Married	0.359	0.564	0.852	0.054	0.091	1.291	0.326	0.540
	(0.032)	(0.083)	(0.065)	(0.172)	(0.053)	(0.157)	(0.037)	(0.099)
Cohabiting	0.447	0.591	0.958	0.093	0.121	1.843	0.411	0.364
	(0.055)	(0.153)	(0.111)	(0.300)	(0.087)	(0.292)	(0.065)	(0.179)
Children								
With children	0.454	0.693	1.074	0.204	0.153	1.529	0.405	0.617
	(0.037)	(0.092)	(0.075)	(0.181)	(0.059)	(0.170)	(0.043)	(0.109)
No children	0.271	0.355	0.581	−0.158	0.014	1.303	0.260	0.281
	(0.041)	(0.122)	(0.087)	(0.259)	(0.072)	(0.233)	(0.049)	(0.146)

Note: Each cell is the BPP-IV pass-through of a permanent income shock to the indicated expenditure, for the BPP-IV sample. Each spouse's shocks are scaled by pre-period budget shares as in equation (2). HH total is household expenditure. Personal M and Personal F are each spouse's direct personal expenditure. Joint is residual household expenditure. Rows other than Overall are estimated on the indicated subsample. Liquid Asset Tercile and Female Income Share rows are terciles, with T1 the lowest. Female income share is the female spouse's average household income share. With children is couples with a child below age 18. Standard errors, clustered at the household level, are in parentheses.

Table B.8: Robustness: Hours Controls, Same-Age Sample, and Broad Sample

	Main		No Hours Control		Same Age		Broad	
	M	F	M	F	M	F	M	F
HH Total	0.383	0.570	0.374	0.517	0.452	0.546	0.413	0.634
	(0.027)	(0.073)	(0.027)	(0.064)	(0.039)	(0.091)	(0.024)	(0.067)
HH Non-Durable	0.291	0.513	0.280	0.458	0.345	0.410	0.316	0.536
	(0.032)	(0.088)	(0.031)	(0.078)	(0.044)	(0.110)	(0.027)	(0.077)
Personal Male	0.881	0.069	0.861	0.056	0.927	−0.041	0.896	0.128
	(0.056)	(0.149)	(0.054)	(0.132)	(0.077)	(0.186)	(0.050)	(0.128)
Personal Female	0.100	1.443	0.110	1.320	0.092	1.206	0.080	1.383
	(0.045)	(0.137)	(0.044)	(0.121)	(0.061)	(0.167)	(0.039)	(0.128)
Joint	0.349	0.494	0.340	0.447	0.430	0.499	0.391	0.598
	(0.032)	(0.086)	(0.031)	(0.077)	(0.045)	(0.108)	(0.028)	(0.077)

Note: Two-step GMM estimates of the pass-through of the permanent change in household income from the spouse's source. Each spouse's shocks are scaled by pre-period budget shares as in equation (2). Main is the baseline specification. No Hours Control omits the contracted-hours transition controls. Same Age restricts to couples with an age difference of at most two years. Broad requires employment only at $t = -1$ and includes separators into non-employment; lead instruments for those workers use co-workers who remain at the origin firm. M and F denote male and female income shocks. Personal consumption is direct personal expenditure. Standard errors, in parentheses, are clustered at the household level.

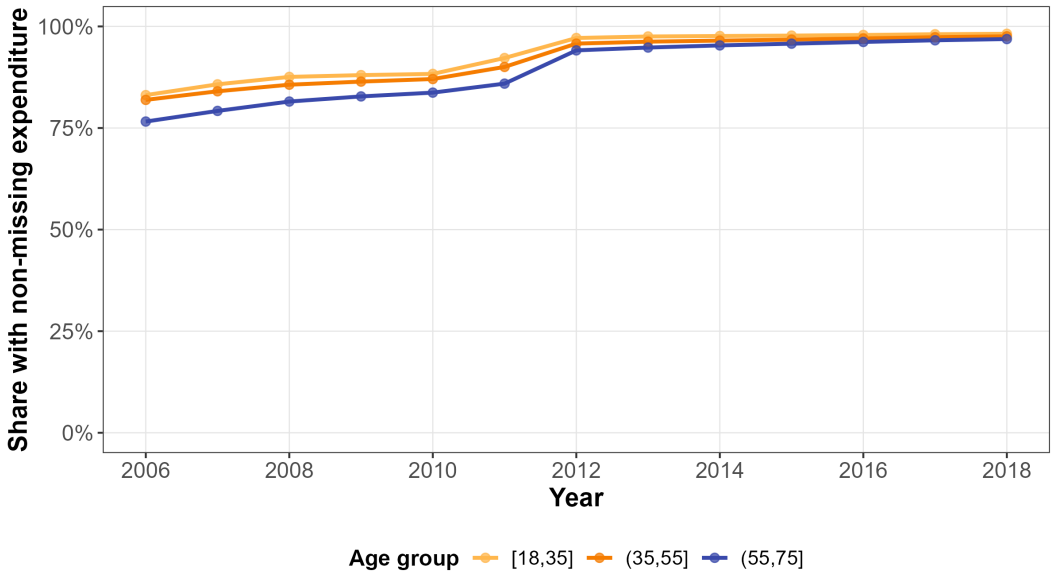
Table B.9: Robustness: Permanent Pass-Through by Minimum Number of Co-Workers

	$n = 8$		$n = 10$		$n = 15$		$n = 20$	
	M	F	M	F	M	F	M	F
HH Total	0.395	0.578	0.383	0.570	0.419	0.608	0.424	0.597
	(0.027)	(0.073)	(0.027)	(0.073)	(0.029)	(0.079)	(0.030)	(0.083)
HH Non-Durable	0.294	0.529	0.291	0.513	0.305	0.566	0.306	0.589
	(0.031)	(0.088)	(0.032)	(0.088)	(0.033)	(0.097)	(0.035)	(0.103)
Personal Male	0.863	0.045	0.881	0.069	0.879	0.099	0.847	0.106
	(0.055)	(0.148)	(0.056)	(0.149)	(0.058)	(0.161)	(0.060)	(0.168)
Personal Female	0.107	1.461	0.100	1.443	0.107	1.531	0.112	1.482
	(0.044)	(0.136)	(0.045)	(0.137)	(0.047)	(0.150)	(0.050)	(0.158)
Joint	0.364	0.492	0.349	0.494	0.395	0.532	0.405	0.518
	(0.032)	(0.086)	(0.032)	(0.086)	(0.033)	(0.094)	(0.035)	(0.098)

Note: Two-step GMM estimates of the pass-through of the permanent change in household income from the spouse's source. Each spouse's shocks are scaled by pre-period budget shares as in equation (2). Each pair of columns varies the minimum number of co-workers used to construct the leave-one-out median income instrument. All columns use the hours-controlled specification. The main specification requires at least 10 co-workers. M and F denote male and female income shocks. Personal consumption is direct personal expenditure. Standard errors, in parentheses, are clustered at the household level.

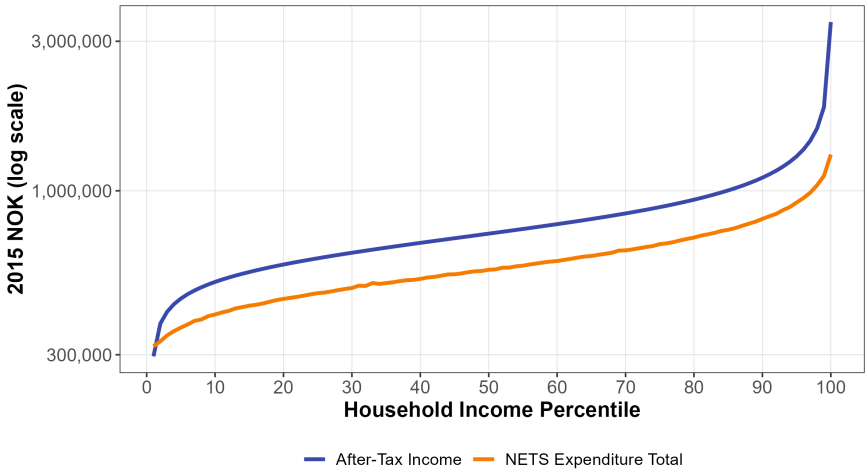
C Data Appendix

Figure C.1: Share of Norwegian Taxpayers with Expenditure Data, by Age Group and Year



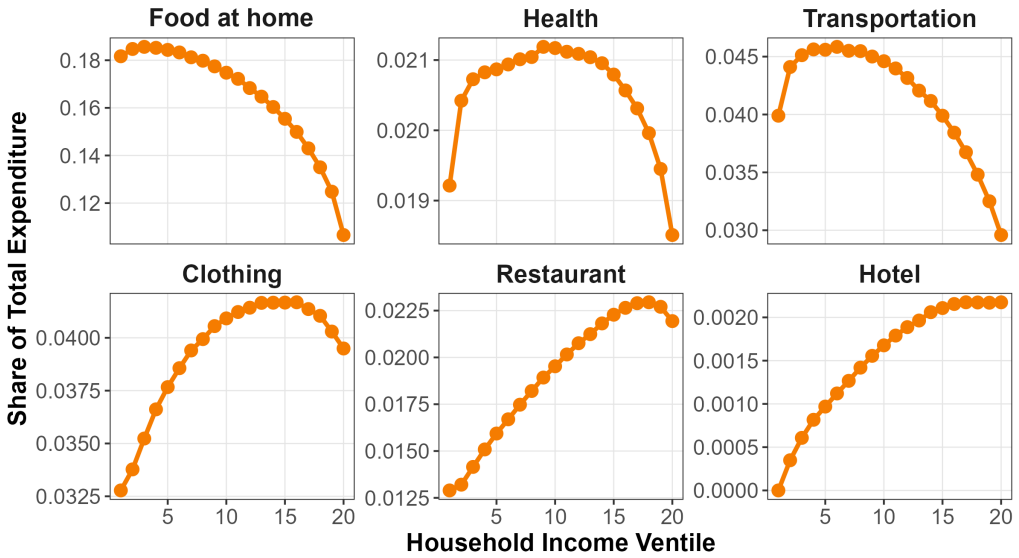
Note: The figure reports, for each year, the share of taxpayers in each age group with non-missing expenditure in the NETS–NICS data. Age groups are $[18, 35]$, $(35, 55]$, and $(55, 75]$. The average coverage rate is 96%. The increase in coverage in 2012 reflects the retention of bank-account linkages for accounts closed before that year; see Ahn et al. (2024).

Figure C.2: Household Income and Consumption by Income Percentile



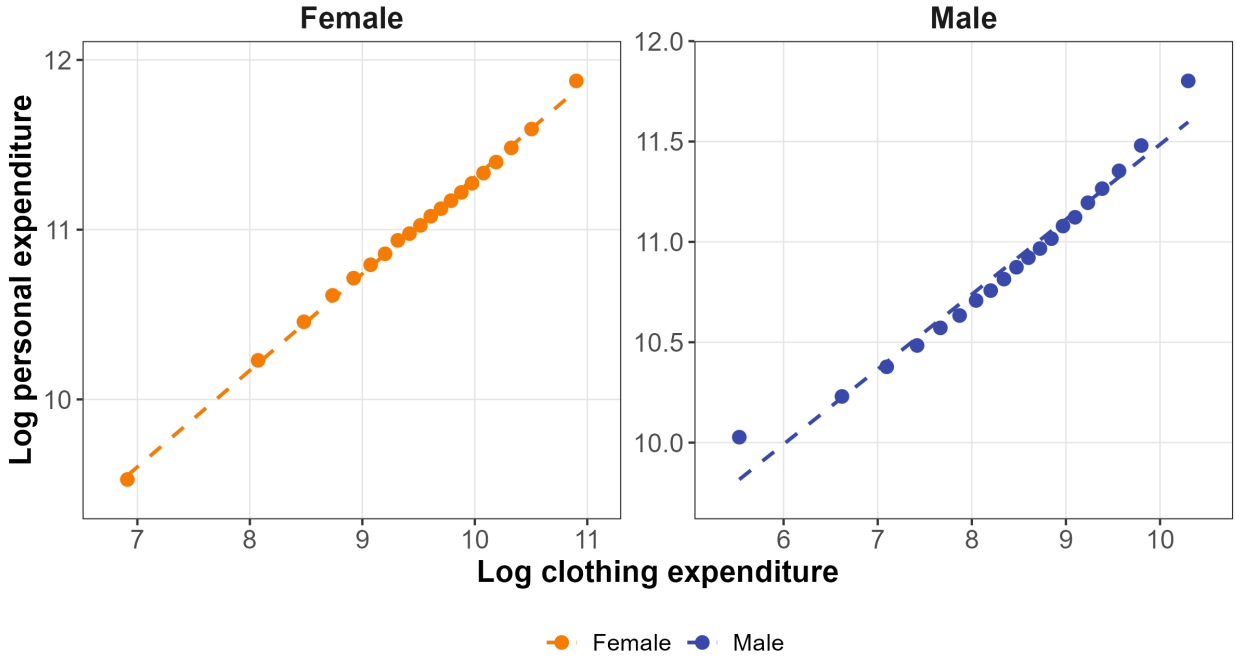
Note: The figure plots mean after-tax household income and mean NETS expenditure, in 2015 NOK, by household income percentile. Both series are on a log scale. Expenditure tracks income across the distribution, and the consumption-to-income ratio is stable except at the very top.

Figure C.3: Engel Curves for Selected Consumption Categories



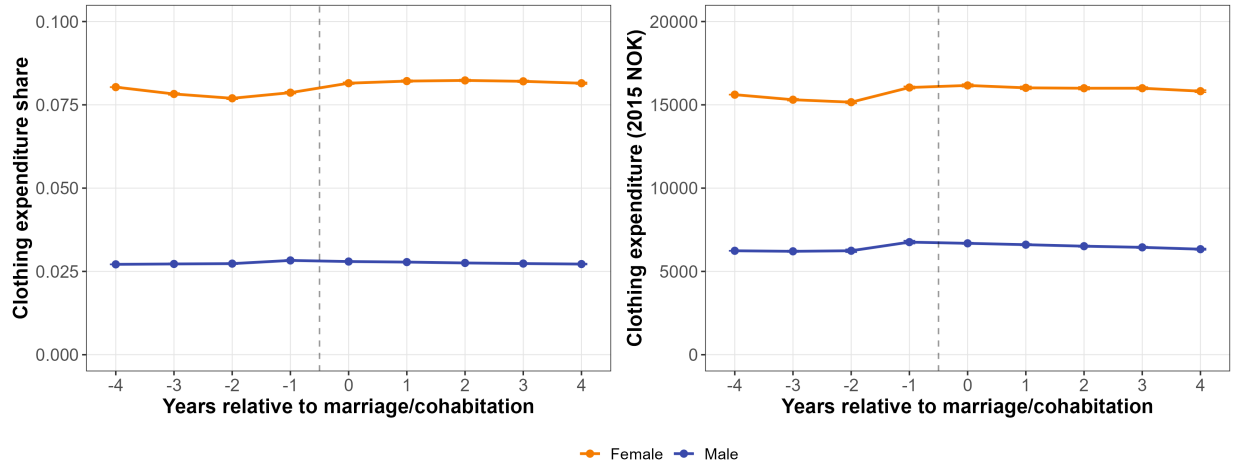
Note: Each panel plots the share of total expenditure in the indicated category against household income ventiles. Points are binned means. Food at home and transportation shares decline with income over most of the distribution. Clothing, restaurant, and hotel shares rise. Health rises through the middle of the distribution and then falls. Food at home is the grocery category used throughout the paper.

Figure C.4: Clothing and Personal Expenditure among Singles



Note: Each panel is a binscatter of log direct personal expenditure against log clothing expenditure, separately for women and men. Personal expenditure uses the personal categories in Appendix Table C.2. The sample is individuals later observed in a couple, in years $t = -5$ through $t = -1$, where $t = 0$ is union formation. Dashed lines are linear fits.

Figure C.5: Clothing Expenditures around Union Formation



Note: The left panel plots the clothing share of total expenditure, and the right panel plots clothing expenditure in 2015 NOK. Both are shown from four years before through four years after marriage or cohabitation, separately for women and men. The vertical dashed line marks the last year before union formation.

Figure C.6: Income Shares and Personal Consumption Shares



Note: The figure is a binscatter of the female share of personal expenditure against the female share of household income. Observed personal expenditure assigns spending to the purchaser, using the personal categories in Appendix Table C.2. Imputed personal expenditure is predicted from clothing expenditure using the pre-union clothing relationship. The average female share is 0.51 for observed personal expenditure and 0.54 for imputed personal expenditure. The dashed line is the 45-degree line.

Table C.1: Benchmark of Consumption Categories against the 2012 SSB Survey of Household Expenditures

Category	NETS–NICS	SSB
Total	341,088	360,793
Groceries	51,345	51,429
Alcohol	5,108	7,656
Clothing	14,346	23,618
Home	39,248	50,236
Furnishings	15,092	24,495
Health	9,115	11,421
Transport	21,574	81,574
Vehicles	6,002	43,226
Vehicle operation	12,888	26,038
Transport services	2,684	12,310
Communication	10,831	8,253
Recreation	31,517	34,960
Audio-visual equipment	2,738	7,515
Recreation durables	794	3,888
Gardens and pets	13,239	8,570
Recreation services	9,658	10,103
Books and stationery	5,088	4,884
Education	1,101	985
Restaurants	8,724	14,304
Accommodation	1,516	1,253
Personal care	3,676	9,705
Personal effects	2,364	2,859
Social protection	823	4,345
Insurance	14,793	8,869

Note: The table compares mean annual household expenditure in the NETS–NICS data with Statistics Norway's 2012 Survey of Household Expenditures. Category labels follow the 1999 COICOP classification used by Ahn et al. (2024). Groceries is food and non-alcoholic beverages; restaurants and alcohol are reported separately. Indented rows are components of Transport and Recreation. Home is utilities and housing maintenance; actual and imputed rent are excluded from the survey measure because owner-occupied rent is not imputed in NETS–NICS. Amounts are in 2012 NOK.

Table C.2: Consumption Categories across Expenditure Definitions

Category	Personal	Joint Narrow	Joint	Non-durable
Alcohol	✓			✓
Clothing	✓			
Health	✓			✓
Transport services	✓			✓
Recreation: AV, music, computers	✓			
Recreation: large durables	✓			
Recreation: toys, sport, pets	✓			
Recreation services	✓			✓
Recreation: news, books	✓			
Restaurants	✓			✓
Hotels	✓			✓
Personal care	✓			✓
Personal effects	✓			
Groceries		✓	✓	✓
Home		✓	✓	✓
Furniture		✓	✓	
Transport durables		✓	✓	
Transport operation		✓	✓	✓
Communications		✓	✓	✓
Education		✓	✓	✓
Social services		✓	✓	✓
Insurance		✓	✓	✓
Public sector			✓	
Imputed credit card			✓	
Uncategorized			✓	

Note: Check marks indicate that the category is included in the expenditure measure. Personal expenditure is the sum of assignable categories recorded for each spouse. The narrow joint bucket includes observed shared categories. Joint expenditure is residual household spending net of both spouses' personal expenditure, and therefore also includes public-sector payments, imputed credit-card spending, and uncategorized expenditure. Home is utilities and housing maintenance; rent is not imputed.

Table C.3: Pre-Union Clothing Imputation of Personal Consumption

	Women	Men
Log clothing	0.549	0.319
	(0.002)	(0.003)
Log clothing × Age: (25,35]	−0.019	0.011
	(0.002)	(0.003)
Log clothing × Age: (35,45]	−0.039	0.036
	(0.003)	(0.003)
Log clothing × Age: (45,55]	−0.059	0.029
	(0.003)	(0.003)
Log clothing × Age: (55,65]	−0.092	0.012
	(0.006)	(0.005)
Log clothing × Age: (65,75]	−0.117	−0.018
	(0.029)	(0.016)
Log clothing × Education: Basic	0.142	0.135
	(0.017)	(0.018)
Log clothing × Education: High school	0.048	0.030
	(0.002)	(0.002)
Log clothing × Education: No education	0.113	0.180
	(0.022)	(0.026)
Log clothing × Education: Post-bachelor	−0.021	−0.022
	(0.003)	(0.003)
Observations	446,904	487,507
R^2	0.487	0.284
Adj. R^2	0.487	0.284
RMSE	0.53	0.68

Note: Each column reports gender-specific OLS estimates of log direct personal expenditure on log clothing expenditure. The sample is individuals later observed in a couple, in years $t = -5$ through $t = -1$, where $t = 0$ is union formation. Personal expenditure uses the personal categories in Appendix Table C.2. Reported coefficients are the log-clothing slope and its interactions with age and education groups. The regressions include age-group and education-group base effects, as in the imputation equation in the main text. Standard errors are in parentheses.

D Survey Appendix

D.1 Selected Survey Questionnaire

Financial Pooling Arrangement. How would you describe the way you and your partner organize income and spending?
Responses:

- Fully joint — almost all income and all spending are pooled and managed jointly.
- Mostly joint — most income and spending are pooled, but some is kept separate.
- Partly joint — we put money in a joint pot for some shared expenses, while the rest of our finances are kept separate.
- Mostly separate — most income and spending are managed separately, with limited pooling.
- Fully separate — each partner manages their own income and spending independently.
- Prefer not to answer.

Reasons for Keeping Finances Separate. Which of the following are reasons that you and/or your partner keep parts of your finances separate? Select all that apply.
Responses:

- Because it is easier to manage money with some accounts separate.
- Our accounts were already set up this way before the relationship.
- To keep some independence in financial or spending decisions.
- Because we have different preferences for spending or saving.
- To keep some privacy in financial or spending decisions.
- To reduce disagreement or tension about spending.
- To have a personal financial buffer.
- To keep personal debt or obligations separate.
- Other.

Expenditure Financing Shares. For each of the following types of expenses, indicate roughly what share is usually covered by your income. Think about whose income ultimately covers these expenses, even if the payments are made from a joint account. If you have no expenses in a category, select not applicable.

- Rent, mortgage, and electricity/municipal charges.
- My personal expenses (clothes, hobbies, and so on).
- My partner's personal expenses (clothes, hobbies, and so on).

Response Scale: 0% (Only the partner) / 25% (Mostly the partner) / 50% (Equal share) / 75% (Mostly me) / 100% (Only me) / Not applicable.

Access to Partner Accounts. How often do you use money from your partner's personal accounts for your own expenses without consulting your partner first?
Responses: I do not have access to my partner's personal accounts / Never / Rarely / Sometimes / Often / Very often / Prefer not to answer.

Access to Joint Accounts. How often do you use money from joint or shared accounts for your own expenses without consulting your partner first?
Responses: I do not have access to our joint accounts / Never / Rarely / Sometimes / Often / Very often / Prefer not to answer.

Perceived Ownership of Income. People also think differently about money in a relationship — whether income feels as if it belongs to the person who earned it, or to the household as a whole. Indicate how much you agree or disagree with each statement.

- Money I earn feels like "my" money to decide over.
- Money my partner earns feels as available to me as money I earn myself.
- In a couple, money should be seen as "ours," not "mine" or "yours."

Response Scale: Strongly disagree / Somewhat disagree / Neither agree nor disagree / Somewhat agree / Strongly agree.

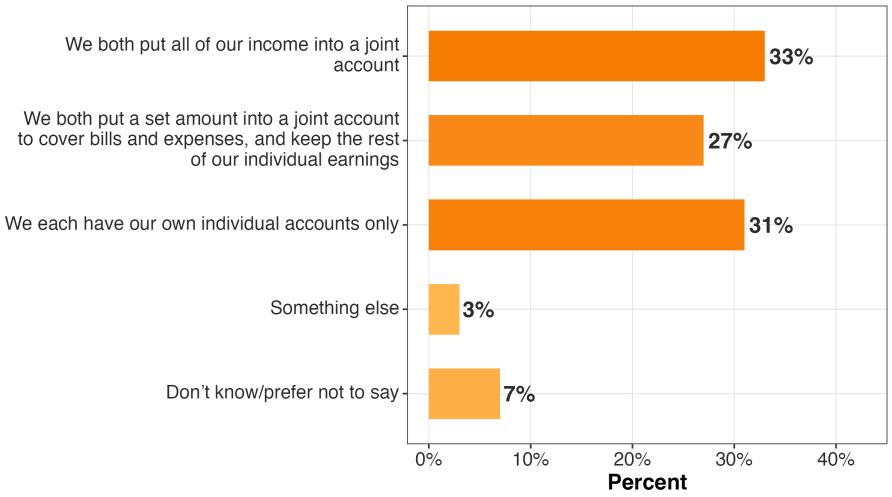
UK Financial Pooling Arrangement. Thinking about you and your partner's finances, which of the following comes closest to describing how they work?
Responses:

- We both put all of our income into a joint account.
- We both put a set amount into a joint account to cover bills and expenses, and keep the rest of our individual earnings.
- We each have our own individual accounts only.
- Something else.
- Don't know/prefer not to say.

Note: This question is from a YouGov survey of UK respondents living with a romantic partner, fielded between July 21, 2023, and February 5, 2024 (YouGov, 2024).

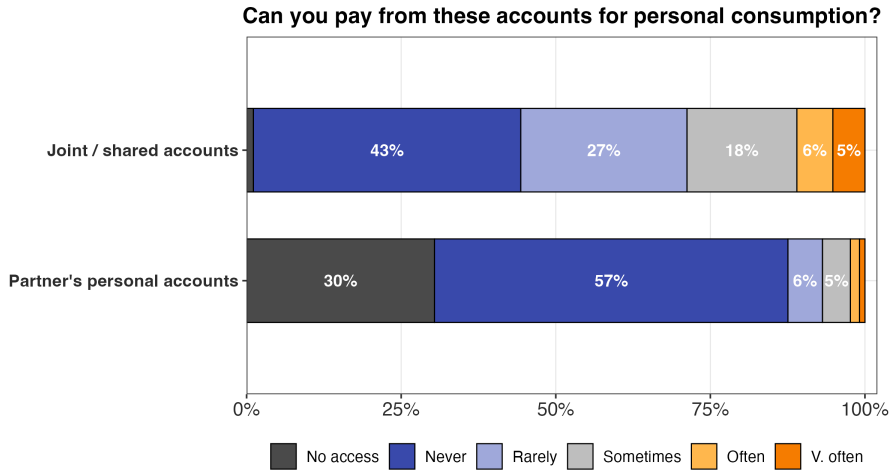
D.2 Additional Survey Exhibits

Figure D.1: Financial Pooling Arrangements among Married and Cohabiting Couples in the United Kingdom



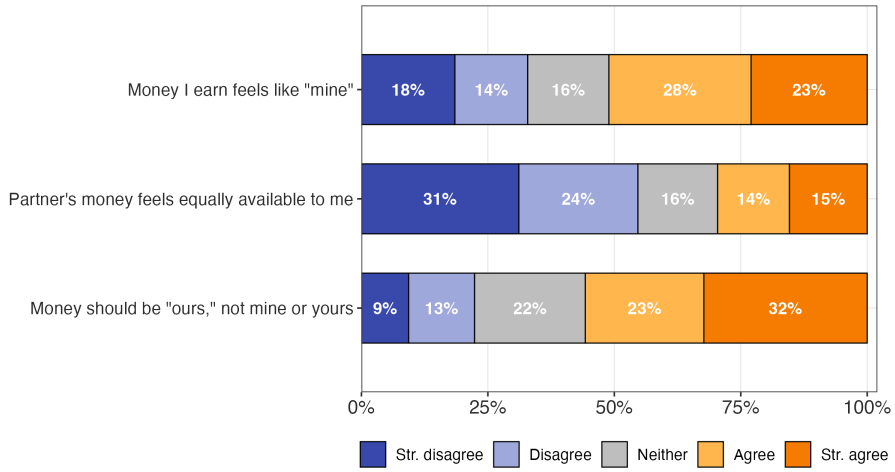
Note: The figure reports the share of respondents in each financial pooling arrangement. Data are from a YouGov survey of UK respondents living with a romantic partner, fielded between July 21, 2023, and February 5, 2024 (YouGov, 2024). The sample includes married respondents and cohabitants. The survey question is reported in Appendix D.1.

Figure D.2: Use of Joint and Partner Accounts for Personal Consumption



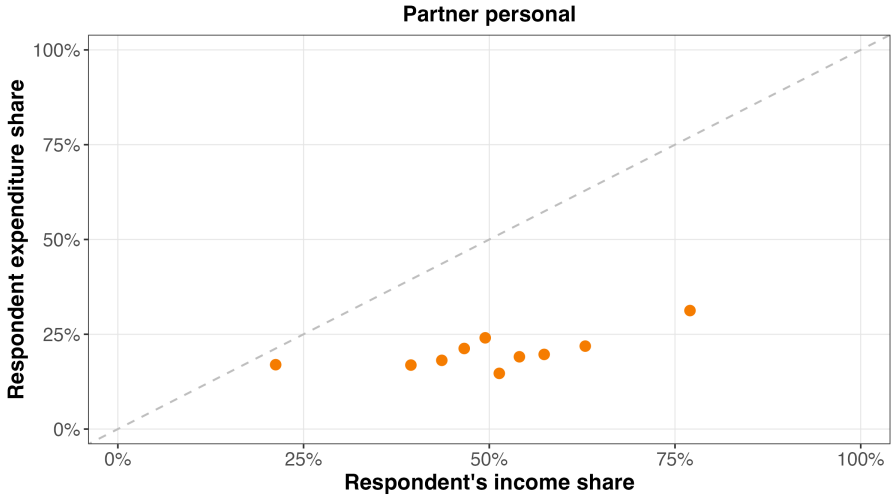
Note: Data are from the authors' survey of Norwegian couples, fielded through the Norstat panel in August 2026. The final sample includes 1,028 respondents. The figure reports how often respondents use money from joint or shared accounts, and from their partner's personal accounts, for their own expenses without consulting their partner first. The survey questions are reported in Appendix D.1 and Appendix D.1.

Figure D.3: Attitudes toward Privately Owned versus Shared Income



Note: Data are from the authors' survey of Norwegian couples, fielded through the Norstat panel in August 2026. The final sample includes 1,028 respondents. The figure reports agreement with three statements: money the respondent earns feels like their own, the partner's money feels equally available, and money in a couple should be seen as shared. The survey question is reported in Appendix D.1.

Figure D.4: Financing of Partner's Personal Expenditures



Note: Data are from the authors' survey of Norwegian couples, fielded through the Norstat panel in August 2026. The final sample includes 1,028 respondents. The figure plots the share of the partner's personal expenses usually covered by the respondent's income against the respondent's share of household income. Points are means within deciles of the respondent's income share. The dashed line is a 45-degree line. The survey question is reported in Appendix D.1.

Table D.1: Survey Sample versus National Targets

Panel A: Demographic Composition

	Sample %	National %
Age		
20–29	11.7	9.4
30–44	31.1	32.7
45–66	43.9	44.8
67–75	13.3	13.1
Gender		
Female	50.6	50.0
Male	49.2	50.0
Other	0.2	—
Region		
Oslo & Akershus	23.2	25.0
Eastern Norway excl. Oslo/Akershus	26.3	25.0
Agder & Rogaland	15.4	15.0
Western Norway	17.6	17.0
Trøndelag	9.1	9.0
Northern Norway	8.5	9.0

Panel B: Comparison with Administrative Sample

	Survey	Table 1
Women's After-Tax Income	323,000	327,000
Men's After-Tax Income	425,000	476,000
Women's Share of Household Income	0.44	0.42
Women's Labor-Force Participation	0.82	0.86
Men's Labor-Force Participation	0.85	0.87

Note: Data are from the authors' survey of Norwegian couples, fielded through the Norstat panel in August 2026. The final sample includes 1,028 respondents. Panel A compares the sample shares for age, gender, and region with the fielding targets for coupled adults aged 20–75. Panel B compares survey means with the administrative sample in Table 1. The survey sample in Panel B is restricted to couples in which at least one spouse works. After-tax income entries are sample means in 2015 NOK, rounded to 1,000 NOK. The remaining entries are shares.

Table D.2: Motives for Financial Separation

Reasons for Separate Accounts	Share
Simpler to manage money with some accounts apart	42%
Accounts were already set up that way before the relationship	40%
Keep some independence in money decisions	42%
Different preferences for spending or saving	40%
Keep some privacy in money decisions	33%
Keep a personal financial buffer	30%
Keep personal debt or obligations separate	25%
Reduce disagreement or tension about spending	22%

Note: Data are from the authors' survey of Norwegian couples, fielded through the Norstat panel in August 2026. The final sample includes 1,028 respondents. Each entry is the share of respondents selecting that reason for keeping parts of their finances separate. Respondents could select more than one reason, so the shares need not sum to 100. The survey question is reported in Appendix D.1.

E Identification of the Pass-Through Parameters

We show that the persistence parameter θ_s, the permanent and transitory covariance exposures (P^s, V^s), and the four pass-through parameters $(\psi_F^P, \psi_F^T, \psi_M^P, \psi_M^T)$ in the instrument-process and consumption equations above are identified in closed form from the covariances of income and consumption with the current, lead-one, and lead-two co-worker instruments.

Setup: Recall the MA(1) statistical model in the main text. The co-worker instrument is

$$\Delta Z_{it}^s = \zeta_{it}^s + \Delta v_{it}^s$$

where the transitory component is MA(1) in levels, $v_{it}^s = \varepsilon_{it}^s + \theta_s \varepsilon_{i,t-1}^s$, so that

$$\Delta Z_{it}^s = \zeta_{it}^s + \varepsilon_{it}^s - (1-\theta_s)\varepsilon_{i,t-1}^s - \theta_s \varepsilon_{i,t-2}^s$$

Own income growth equals the same firm-driven permanent and transitory changes plus a worker-specific residual,[60]

$$\Delta y_{it}^s = \zeta_{it}^s + \Delta v_{it}^s + \eta_{it}^s$$

and household consumption responds to the current permanent and transitory firm-level shocks of both spouses,

$$\Delta c_{it} = \psi_F^P w_{it-1}^F \zeta_{it}^F + \psi_F^T w_{it-1}^F \varepsilon_{it}^F + \psi_M^P w_{it-1}^M \zeta_{it}^M + \psi_M^T w_{it-1}^M \varepsilon_{it}^M + u_{it}$$

where w_{it-1}^s is spouse s's lagged income share and u_{it} is the consumption residual. We assume the innovations ζ_{it}^s and ε_{it}^s are i.i.d., mean zero, and mutually orthogonal across spouses and that Assumptions 1–2 hold.

Let

$$P^s \equiv \mathrm{Var}(\zeta_{it}^s) \qquad V^s \equiv \mathrm{Var}(\varepsilon_{it}^s)$$

denote the permanent and transitory firm-shock variances.

We write $\bar{w}^s \equiv E[w_{it-1}^s]$, and let $\mathcal{I}_{i,t-1}$ denote household i's information at $t-1$, so that $w_{it-1}^s \in \mathcal{I}_{i,t-1}$.

[60] The derivation treats own income as loading one-for-one on the firm shocks. If instead $\Delta y_{it}^s = \lambda_s \zeta_{it}^s + \lambda_s \Delta v_{it}^s + \eta_{it}^s$ and consumption responds to the realized worker-level shocks $\lambda_s \zeta_{it}^s$ and $\lambda_s \varepsilon_{it}^s$, the income moments absorb λ_s into (P^s, V^s) and identification of (ψ_s^P, ψ_s^T) is unchanged. The same cancellation holds if the permanent and transitory loadings differ.

Income moments: Under the income process above, the current and lead covariances of income growth with the firm instrument are:

$$E[\Delta y^s_{it}\Delta Z^s_{it}] = P^s + 2(1-\theta_s+\theta_s^2)V^s \tag{3}$$

$$E[\Delta y^s_{it}\Delta Z^s_{i,t+1}] = -(1-\theta_s)^2 V^s \tag{4}$$

$$E[\Delta y^s_{it}\Delta Z^s_{i,t+2}] = -\theta_s V^s \tag{5}$$

If $E[\Delta y^s_{it}\Delta Z^s_{i,t+1}] \neq 0$, the persistence parameter θ_s is uniquely determined by

$$\frac{E[\Delta y^s_{it}\Delta Z^s_{i,t+2}]}{E[\Delta y^s_{it}\Delta Z^s_{i,t+1}]} = \frac{\theta_s}{(1-\theta_s)^2}$$

Then V^s follows from either lead income moment. Using the first lead,

$$V^s = -\frac{E[\Delta y^s_{it}\Delta Z^s_{i,t+1}]}{(1-\theta_s)^2}$$

The permanent variance P^s is

$$P^s = E[\Delta y^s_{it}\Delta Z^s_{it}] - 2(1-\theta_s+\theta_s^2)V^s$$

Consumption moments: The consumption equation and Assumptions 1–2 leave only the own-spouse firm shocks, weighted by the lagged income share:

$$E[\Delta c_{it}\Delta Z^s_{it}] = \psi^P_s E\big[w^s_{it-1}(\zeta^s_{it})^2\big] + \psi^T_s E\big[w^s_{it-1}(\varepsilon^s_{it})^2\big] \tag{6}$$

$$E[\Delta c_{it}\Delta Z^s_{i,t+1}] = -\psi^T_s(1-\theta_s)\, E\big[w^s_{it-1}(\varepsilon^s_{it})^2\big] \tag{7}$$

$$E[\Delta c_{it}\Delta Z^s_{i,t+2}] = -\psi^T_s\theta_s\, E\big[w^s_{it-1}(\varepsilon^s_{it})^2\big] \tag{8}$$

Because the firm-level innovations are i.i.d., they are independent of predetermined information:

$$(\zeta^s_{it}, \varepsilon^s_{it}) \perp \mathcal{I}_{i,t-1}$$

Hence

$$E\big[(\zeta^s_{it})^2 \mid \mathcal{I}_{i,t-1}\big] = P^s \qquad E\big[(\varepsilon^s_{it})^2 \mid \mathcal{I}_{i,t-1}\big] = V^s$$

The lagged share then factors out of the consumption moments by the law of iterated expectations,

$$E\big[w^s_{it-1}(\zeta^s_{it})^2\big] = E\Big[E\big[w^s_{it-1}(\zeta^s_{it})^2 \mid \mathcal{I}_{i,t-1}\big]\Big] = E\Big[w^s_{it-1}\,E\big[(\zeta^s_{it})^2 \mid \mathcal{I}_{i,t-1}\big]\Big] = E[w^s_{it-1}]P^s = \bar{w}^s P^s$$

and likewise $E[w_{it-1}^s(\varepsilon_{it}^s)^2] = \bar{w}^s V^s$. Substituting these factorizations yields:

$$E[\Delta c_{it} \Delta Z_{it}^s] = \psi_s^P \bar{w}^s P^s + \psi_s^T \bar{w}^s V^s \tag{9}$$

$$E[\Delta c_{it} \Delta Z_{i,t+1}^s] = -\psi_s^T \bar{w}^s (1-\theta_s) V^s \tag{10}$$

$$E[\Delta c_{it} \Delta Z_{i,t+2}^s] = -\psi_s^T \bar{w}^s \theta_s V^s \tag{11}$$

Closed-form identification: The transitory pass-through parameter follows from either lead consumption moment. For $\theta_s \in (0,1)$,

$$\psi_s^T = -\frac{E[\Delta c_{it} \Delta Z_{i,t+2}^s]}{\bar{w}^s \theta_s V^s} = -\frac{E[\Delta c_{it} \Delta Z_{i,t+1}^s]}{\bar{w}^s (1-\theta_s) V^s}$$

Substituting ψ_s^T into the contemporaneous consumption moment and solving for the permanent parameter gives

$$\psi_s^P = \frac{E[\Delta c_{it} \Delta Z_{it}^s] - \psi_s^T \bar{w}^s V^s}{\bar{w}^s P^s}$$

Because (θ_s, P^s, V^s) are pinned down by the income moments above, these expressions identify ψ_s^T and ψ_s^P for each spouse $s \in \{F, M\}$ from population moments alone, provided $\bar{w}^s > 0$ and $P^s, V^s > 0$, i.e. the first stage is non-degenerate.

These moment conditions establish the identification of $(\theta_s, P^s, V^s, \psi_s^P, \psi_s^T)$ for each spouse. In practice we estimate the system jointly by two-step GMM, stacking the income and consumption moment conditions above for both spouses and clustering at the household level.

E.1 A Bound on the Own–Cross Gap under Fixed Spillovers

Direct personal expenditure may include purchases for the other spouse. We show that under fixed spillovers and own-shock dominance—true personal consumption responds at least as strongly to own income as to the spouse's—the observed own–cross pass-through gap in personal consumption is a lower bound on the true consumption gap.

Let c_j denote spouse j's true personal consumption, and suppose observed female personal expenditure includes a fixed share of her husband's consumption,

$$\widetilde{c}_F = c_F + \lambda c_M, \qquad \lambda \geq 0,$$

where λ is fixed. Let z_s denote spouse s's income shock, scaled by household income, and

define the pass-through into true log personal consumption as

$$\psi_{j\leftarrow s} = \frac{\partial \log c_j}{\partial z_s}, \qquad j, s \in \{F, M\}.$$

Let $\widetilde{\psi}_{F\leftarrow s}$ denote the corresponding pass-through into observed log female personal expenditure. Holding λ fixed, differentiation gives

$$\begin{aligned}
\widetilde{\psi}_{F\leftarrow s} &= \frac{\partial \log \widetilde{c}_F}{\partial z_s} \\
&= \frac{1}{c_F + \lambda c_M}\left(\frac{\partial c_F}{\partial z_s} + \lambda \frac{\partial c_M}{\partial z_s}\right) \\
&= \frac{c_F}{c_F + \lambda c_M}\psi_{F\leftarrow s} + \frac{\lambda c_M}{c_F + \lambda c_M}\psi_{M\leftarrow s}.
\end{aligned}$$

Defining $\alpha = \frac{\lambda c_M}{c_F + \lambda c_M}$ as the share of observed female personal expenditure attributable to male consumption, we obtain

$$\widetilde{\psi}_{F\leftarrow s} = (1-\alpha)\psi_{F\leftarrow s} + \alpha\psi_{M\leftarrow s}, \qquad s \in \{F, M\}.$$

The observed own–cross gap is therefore

$$\begin{aligned}
\widetilde{\psi}_{F\leftarrow F} - \widetilde{\psi}_{F\leftarrow M} = &\left(\psi_{F\leftarrow F} - \psi_{F\leftarrow M}\right) \\
&- \alpha\Big[\left(\psi_{F\leftarrow F} - \psi_{F\leftarrow M}\right) + \left(\psi_{M\leftarrow M} - \psi_{M\leftarrow F}\right)\Big].
\end{aligned}$$

The bound requires own-shock dominance: $\psi_{F\leftarrow F} \geq \psi_{F\leftarrow M}$ and $\psi_{M\leftarrow M} \geq \psi_{M\leftarrow F}$. This condition generally holds in the full-commitment and limited-commitment models we outline in Section 6. Since $\alpha \geq 0$, the second term is nonnegative, implying

$$\widetilde{\psi}_{F\leftarrow F} - \widetilde{\psi}_{F\leftarrow M} \leq \psi_{F\leftarrow F} - \psi_{F\leftarrow M}.$$

A similar result holds for the pass-through gap in male consumption. Fixed expenditure spillovers attenuate the signed own–cross gap for both spouses. We note the result is local and may require additional assumptions to aggregate across households.

F Solution and Estimation of Household Models

F.1 Estimation Details

We solve the three-period model by backward induction, beginning with the terminal-period allocation and then solving the participation and allocation problems in each preceding period. For each candidate parameter vector we simulate $N = 10{,}000$ couples, drawing with replacement from the empirical joint distribution of assets and permanent incomes among newly formed couples.

Conditional on these initial states, permanent income shocks are drawn from $\eta_{s,t} \sim N(0, \widehat{\sigma}^2_{\eta_s})$, $s \in \{F, M\}$, where $\widehat{\sigma}^2_{\eta_F}$ and $\widehat{\sigma}^2_{\eta_M}$ are the gender-specific permanent-shock variances estimated from the income data. Initial match quality is zero, and match-quality shocks are drawn from $N(0, \sigma^2_q)$. The initial Pareto weight is drawn uniformly from the weights that satisfy both participation constraints, and draws for which that set is empty are rejected. We then apply the model's policy functions to the simulated histories and construct the same divorce hazards and pass-through elasticities as in the data.

Pre-stage: Formation Parameters Divorce hazards, initial assets and income, and permanent-income variances are estimated for newly formed couples aged 25 to 45, whose initial conditions we are able to observe during our sample window. The pass-through targets remain the main-sample BPP-IV estimates, which are more precisely estimated in the larger sample.

First Stage: Public-good Taste Weights We first estimate α_{gF} and α_{gM} from the static allocation of private and joint consumption, following Lise and Yamada (2019). The first-order condition for private consumption of spouse s is

$$\omega_s u_s'(C_s) \frac{\partial C_s}{\partial c_s} = \lambda,$$

where $\omega_F = \mu$ and $\omega_M = 1 - \mu$. The first-order condition for the joint good is

$$\sum_{s \in \{F,M\}} \omega_s u_s'(C_s) \frac{\partial C_s}{\partial g} = \lambda.$$

Substituting the private-good conditions into the joint-good condition yields

$$\frac{\alpha_{gF}}{\alpha_{cF}} \left(\frac{g}{c_F} \right)^{\phi - 1} + \frac{\alpha_{gM}}{\alpha_{cM}} \left(\frac{g}{c_M} \right)^{\phi - 1} = 1. \tag{12}$$

Thus, the Pareto weight and the marginal utility of consumption cancel from the estimating equation. We fix $\phi = 0.1$, normalize $\alpha_{cs} + \alpha_{gs} = 1$ for each spouse, and estimate the two public-good weights from the sample analogue of equation (12). This procedure gives $\widehat{\alpha}_{gF} = 0.67$ and $\widehat{\alpha}_{gM} = 0.62$. Table F.1 reports the corresponding weights at other values of ϕ. These first-stage estimates are held fixed in both the single-budget limited-commitment model and the separate-budgets model.[61]

Table F.1: Public-Good Weights at Alternative Values of ϕ

ϕ	Women α_{g1}	Men α_{g2}
0.05	0.683	0.632
0.10	0.671	0.622
0.25	0.631	0.586
0.50	0.549	0.514

Note: The table reports first-stage public-good taste weights estimated from the intratemporal first-order condition at each value of ϕ, using equation (12).

Second Stage: Simulated Method of Moments Holding the first-stage estimates fixed, we estimate

$$\vartheta = (\gamma_F, \gamma_M, \delta_F, \delta_M, \sigma_q)'$$

separately in each model by weighted simulated method of moments. Let $\widehat{m}$ denote the vector of empirical moments and $m^S(\vartheta)$ its simulated counterpart. The estimator is

$$\widehat{\vartheta} = \arg\min_{\vartheta} \left[\widehat{m} - m^S(\vartheta)\right]' W \left[\widehat{m} - m^S(\vartheta)\right],$$

where W is a diagonal matrix of the inverse squares of the targeted moments. The moment vector contains the own- and cross-spouse pass-through elasticities for personal consumption, pass-through to joint consumption from female and male shocks,[62] and observed separation rates at $t = 2$ and $t = 3$.

Table F.2 reports parameters that are pinned or taken from the first stage. Tables F.3 and F.5 report second-stage estimates and targeted moments for the single-budget limited-commitment model and the separate-budgets model.

[61] The first-order condition used to estimate these weights does not hold under separate budgets. We keep the weights for comparability. In the estimated models, the shares of personal and joint consumption are similar under single-budget and separate-budget limited commitment.

[62] The permanent shocks are observed in the simulation. The model pass-through elasticities are regression coefficients of consumption growth on the permanent-shock innovations, scaled by each spouse's pre-period income share.

F.2 Additional Tables

Table F.2: Estimation and Simulation Parameters

Object	Value or source	Source
Discount factor, β	0.97	Fixed
Gross return, $1 + r$	$1/\beta$	Fixed
Asset shares at separation, κ_F, κ_M	$0.50, 0.50$	Fixed
Borrowing constraint, $\underline{a}$	0	Fixed
CES parameter, ϕ	0.10	Fixed
Public-good weights, α_{gF}, α_{gM}	$0.67, 0.62$	First stage
Permanent-shock variances, $\sigma^2_{\eta_F}, \sigma^2_{\eta_M}$	$0.03, 0.02$	Pre-stage
Mean formation log income, y^0_F, y^0_M	$-1.08, -0.76$	Pre-stage
Mean formation log assets, a^0_F, a^0_M	$-2.80, -2.56$	Pre-stage

Note: The table reports parameters fixed externally or estimated before the second-stage estimation. Fixed parameters are either standard values taken from the literature (β) or set by the authors. Pre-stage estimates are formation parameters estimated on newly formed couples aged 25 to 45. Income and assets are normalized by household income, so mean household income in the sample is 1. First stage denotes the public-good weights from the intratemporal condition in equation (12) at $\phi = 0.1$.

Table F.3: Single-Budget Limited-Commitment Estimates and Targeted Moments

Panel A. Second-stage parameters		
Parameter	Meaning	Sample
γ_F	Risk aversion, women	2.25
γ_M	Risk aversion, men	2.78
δ_F	Divorce cost, women	≈ 0
δ_M	Divorce cost, men	≈ 0
σ_q	Match-quality shock s.d.	0.71
Panel B. Targeted moments		
Moment	Data	Sample
Divorce hazard, $t = 2$	0.089	0.086
Divorce hazard, $t = 3$	0.065	0.115
Personal own, women	1.44	0.94
Personal cross, women	0.10	0.77
Personal own, men	0.88	0.75
Personal cross, men	0.07	0.49
Joint, women	0.49	0.73
Joint, men	0.35	0.76

Note: Panel A reports estimated second-stage parameters of the single-budget limited-commitment model for the sample specification ($\phi = 0.1$). Panel B compares the empirical moments targeted in estimation with their model counterparts. Personal "own" denotes the pass-through of a spouse's personal consumption to their own income shock; "cross" denotes the pass-through to the partner's income shock. "Joint" denotes the pass-through to the joint good.

Table F.4: Alternative Single-Budget Limited-Commitment Specifications

	Direct							Implied	
	Data	Main	ϕ=0.25	ϕ=0.5	κ_g=0.85	κ_g=0.75	$a \geq -0.5$	Data	Model
Divorce hazard, $t = 2$	0.089	0.086	0.085	0.083	0.082	0.067	0.086	0.089	0.079
Divorce hazard, $t = 3$	0.065	0.115	0.114	0.116	0.113	0.083	0.116	0.065	0.109
Personal own, women	1.44	0.94	0.98	1.08	1.02	1.17	0.91	0.96	0.89
Personal cross, women	0.10	0.77	0.80	0.79	0.69	0.68	0.78	0.23	0.69
Personal own, men	0.88	0.75	0.73	0.76	0.87	0.91	0.75	0.70	0.78
Personal cross, men	0.07	0.49	0.45	0.39	0.50	0.36	0.48	0.01	0.56
Joint, women	0.49	0.73	0.72	0.74	0.78	0.78	0.70	0.56	0.74
Joint, men	0.35	0.76	0.77	0.77	0.79	0.82	0.76	0.39	0.74

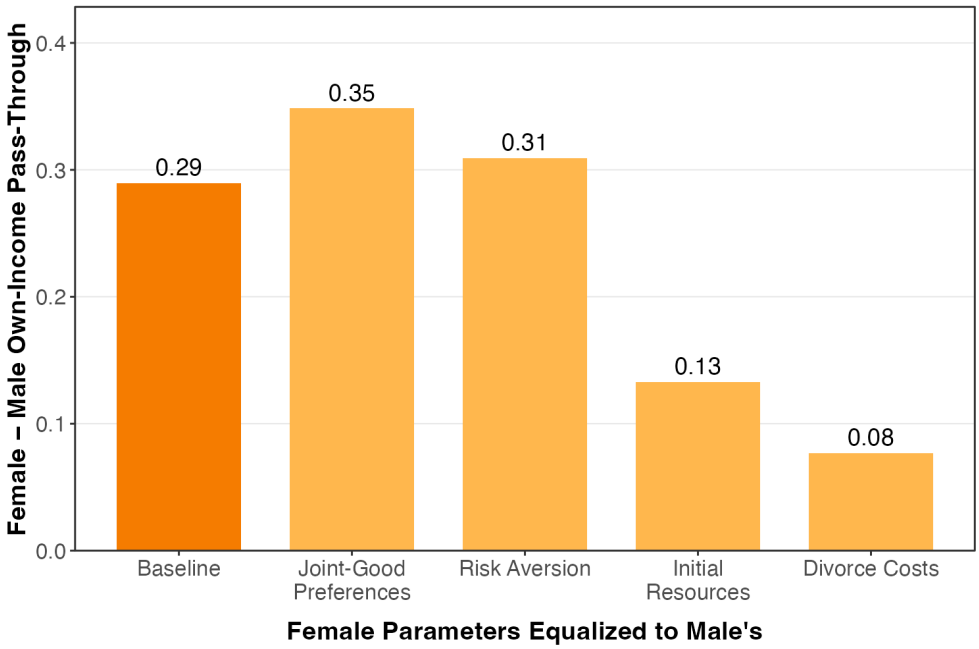
Note: Each column is a single-budget limited-commitment specification estimated on couples aged 25–45. Direct specifications target pass-through into direct personal expenditure and share the data column. Implied targets clothing-imputed personal consumption and the corresponding joint residual. Main is the baseline specification ($\phi = 0.1$, $\kappa_g = 1$, $\underline{a} = 0$). The column $a \geq -0.5$ lowers the asset floor to $\underline{a} = -0.5$, in units of household income. The ϕ columns use public-good weights from the intratemporal first-order condition. κ_g allows for rivalry in the joint good by scaling the effective consumption of each partner ($\kappa_g g$). Personal "own" is the pass-through of a spouse's personal consumption to their own income shock; "cross" is the pass-through to the partner's income shock. "Joint" is the pass-through to the joint good.

Table F.5: Separate-Budget Limited-Commitment Estimates and Targeted Moments

Panel A. Second-stage parameters		
Parameter	Meaning	Sample
γ_F	Risk aversion, women	2.45
γ_M	Risk aversion, men	2.85
δ_F	Divorce cost, women	2.17
δ_M	Divorce cost, men	≈ 0
σ_q	Match-quality shock s.d.	0.90
Panel B. Targeted moments		
Moment	Data	Sample
Divorce hazard, $t = 2$	0.089	0.066
Divorce hazard, $t = 3$	0.065	0.088
Personal own, women	1.44	1.17
Personal cross, women	0.10	0.15
Personal own, men	0.88	0.88
Personal cross, men	0.07	0.11
Joint, women	0.49	0.91
Joint, men	0.35	0.84

Note: Panel A reports estimated second-stage parameters of the non-pooled limited-commitment model for the sample specification ($\phi = 0.1$). Panel B compares the empirical moments targeted in estimation with their model counterparts. Personal "own" denotes the pass-through of a spouse's personal consumption to their own income shock; "cross" denotes the pass-through to the partner's income shock. "Joint" denotes the pass-through to the joint good.

Figure F.1: Gender Gap in Own-Income Pass-Through under Alternative Parameter Values



Note: Bars report the gap between the woman's and the man's own-income pass-through to personal consumption in the separate-budgets limited-commitment model. The base specification is the estimated model. Each other bar gives the woman the man's value of one object, holding the rest fixed: the public-good weight, risk aversion, initial income and assets, or the divorce cost. The changes are not cumulative.

Table F.6: Income-Floor Gains at Alternative Thresholds and Replacement Rates

	Changing q, constant τ			Changing τ, constant q		
	q_5	q_{10}	q_{20}	$\tau = 0.1$	$\tau = 0.25$	$\tau = 0.5$
Panel A. Women						
LC+SB	0.012	0.08	0.47	0.19	0.47	0.89
FC	0.005	0.02	0.14	0.06	0.14	0.28
Panel B. Men						
LC+SB	0.007	0.05	0.46	0.18	0.46	0.84
FC	−0.000	0.02	0.14	0.06	0.14	0.28

Note: The table shows the median gain in lifetime consumption, in percent, for all spouses. The income floor is financed by a tax on income above the 80th percentile, with the tax rate set so that the present value of taxes equals the present value of transfers. Changing q holds $\tau = 0.25$. Changing τ holds q at the 20th percentile. LC+SB is the estimated limited-commitment model with separate budgets. FC imposes full commitment with a single budget, using the estimated preferences and parameters.