Pre-committed Consumption and Family Insurance after Job Loss

Helene Onshuus*

October 2025

Abstract

How does family composition affect consumption insurance? I study the consumption expenditure response to job loss in Norway from 1997 to 2011, and present two empirical facts: Among job losers without children, the consumption expenditure response does not differ between couples and singles. Second, single parent job losers reduce consumption expenditure much less than married or cohabiting job losers with children. Both findings run contrary to workhorse models of family insurance. I resolve the puzzle in a model of household consumption with pre-committed consumption and show that constrained households have higher marginal propensities to consume than unconstrained households.

1 Introduction

Can families insure each other against income loss? The question has been asked both by researchers investigating the added worker effect and by macroeconomists assessing the degree of self-insurance. Although most recent studies of the added worker effect have found the effect of income loss on spousal labor supply to be modest, family insurance through spousal labor supply continues to feature prominently in macroeconomic theory. Among recent papers contributing to our understanding of the role of family insurance are Heathcote, Storesletten, and Violante (2010) and Blundell, Pistaferri, and Saporta-Eksten (2016, 2018). These papers construct and estimate comprehensive macroeconomic models where they find spousal labor supply to be an important self-insurance mechanism. However, there exists very little direct empirical evidence on whether family composition matters for consumption self-insurance once a bad income shock actually hits.

The raising of own children is an important motive for household formation, and since children consume without bringing in much additional income, the presence of children tends

*Oslo Metropolitan University and BI Norwegian Business School (helene.onshuus@oslomet.no). I am thankful to Kjetil Storesletten, Gisle Natvik, Andreas Fagereng, and colleagues and seminar participants at BI Norwegian Business School and Oslo Metropolitan University for excellent comments and discussions. I am grateful for funding from the Research Council of Norway (project #315008 and project #287720). Data for this paper is provided through Statistics Norway.

to place additional strain on the household's financial situation. Previous research has examined the role of children in shaping the consumption profile of parents (e.g. Browning and Ejrnæs (2009) or Love (2010)) and to a certain extent how children affect marginal consumption decisions (Blundell, Pistaferri, and Saporta-Eksten (2018). But there is little direct empirical evidence on how children affect consumption insurance after income loss.

In this study, I provide novel evidence on how the consumption expenditure response to unemployment varies between job losers with different household composition, splitting the sample into *couple/single* and *has children/childless*. I use a difference-in-difference design with staggered implementation around the time an individual registers as unemployed, and a dataset sourced from Norwegian administrative records covering the universe of individuals in Norway. The dataset combines precisely measured income and wealth data from the Norwegian Tax Authority, registered unemployment spells, marital status and family ties from the population registry, as well as demographic characteristics from various sources. Absent observed consumption expenditure, I compute a measure of imputed total consumption expenditure using the household budget constraint.

I uncover two novel empirical facts: First, I find that there is no difference in the consumption expenditure response between couples and singles without children once I take the household's liquid (financial) wealth position into account. This finding suggests that couples rely on self-insurance through precautionary saving rather than self-insurance through spousal labor supply. Second, I find that couples with children exhibit a much larger consumption expenditure response than singles with children, contrary to what workhorse models of self-insurance through spousal labor supply would suggest. Since I control for liquid wealth, the low consumption response by single parents cannot be explained by differences in the *ability* to self-insure. Overall, households with children reduce consumption less after job loss than does households without children.

I argue that the empirical findings of this paper can be explained by consumption constraints in the form of pre-committed consumption, or a lower bound on the consumption of certain goods. Some types of consumption simply cannot be reduced past a certain limit. Housing is an obvious example of a commitment good. One can increase the quality of one's current housing by investing more in it or by varying the consumption of energy or other goods that are complementary to housing. Over time, households can relocate in order to decrease expenditure, but in the short run there is a lower bound of housing expenditure that is determined by past choices. Second, children can be considered a commitment good for parents: Children tend to increase household consumption, and parents may be reluctant to reduce child-related expenses during a temporary income loss, for example if those expenditures are considered investments into the child's future economic outcomes.

In the model, pre-committed consumption depends on family type through two mechanisms: Economies of scale in household consumption and cost-sharing within the household. Access to economies of scale in household consumption is an important motive for household formation. Couples can share most housing-related expenses, while the good is more or less a within-family public good for the couple. Children can also be regarded as a within-family public good for the parents, and parents can split child-related expenses without loss of utility. Economies of scale in housing and couples' cost-sharing of housing and child-related expenses implies that singles have a higher likelihood of becoming constrained by the lower bound on consumption after a negative shock than couples.

In the second part of the paper I construct a toy model to formalize the idea of how a lower bound on the consumption of certain goods can decrease the marginal propensity to consume for constrained households. In the model, households allocate their lifetime resources between two consumption goods and a warm-glow bequest. One of the consumption goods is subject to a lower bound (what I call "pre-committed consumption") which depends on family type. I show how a lower bound on certain consumption categories decreases the marginal propensity to consume for constrained households, since (given decreasing marginal utility of consumption) it is always more painful to concentrate a reduction of consumption expenditure in one consumption category.

The mechanism I propose (pre-committed consumption with economies of scale in the commitment level) suggests a novel explanation for why the estimates of the added worker effect vary across studies: The need for increased spousal labor supply to self-insure against income loss also depends on the probability of being constrained by the pre-committed level of consumption.

This paper relates to the literature exploring the consequences of unemployment, and particularly those concerned with the consumption expenditure response after unemployment. Among the seminal studies in this literature are Gruber (1997) and Browning and Crossley (2001), who estimate the elasticity of consumption expenditure to marginal increases in unemployment benefits. This paper is more similar in approach to some recent studies that have used various versions of difference-in-difference models to estimate the magnitude of the consumption expenditure drop around the time of unemployment (Ganong and Noel, 2019; Landais and Spinnewijn, 2021; Gerard and Naritomi, 2021; Andersen, Jensen, Johannesen, Kreiner, Leth-Petersen, and Sheridan, 2023; Fagereng, Onshuus, and Torstensen, 2024).

Most importantly, this paper relates to the research on the role of family for consumption insurance. Most of these rely on estimated structural models, such as Heathcote, Storesletten, and Violante (2010), Blundell, Pistaferri, and Saporta-Eksten (2016, 2018), and Wu and Krueger (2021). Autor, Kostøl, Mogstad, and Setzler (2019) use exogenous disability income shocks to show that spousal labor supply can provide consumption insurance. Finally, this paper relates to the literature on the role children have in shaping consumption behavior and wealth accumulation over the life-cycle (Browning and Ejrnæs (2009); Love (2010)).

The empirical results I present in this paper have important implications for economic theory. My findings provide a more nuanced and granular picture of marginal consumption expenditure decisions in households of varying family composition. In times of demographic change, family-type heterogeneity in the pass-though of income shocks to consumption expenditure matters for the propagation of aggregate shocks, the effectiveness of monetary and fiscal policy, and similar macroeconomic research agendas. Further, the ability of households to self-insure is crucial for optimal social insurance policy design. The theoretical contribution of this paper, which is to demonstrate how a lower bound on certain consumption categories can decrease the marginal propensity to consume for constrained households, is an important piece of the puzzle in the attempt of researchers to understand how large income shocks transmit to consumption expenditure decisions.

The paper proceeds as follows: Section 2 presents the institutional setting, the data used in the analysis, and the sample selection criteria. Section 3 presents the empirical method, and Section 4 reports the empirical results. In Section 5, I present a simple model that is able to account for the pattern observed in the data, and in Section 6, I conclude.

2 Institutional setting, data and sample selection

2.1 Institutional setting

The Norwegian welfare state provides partial insurance against joblessness with mandatory participation, but the unemployment insurance scheme is both less generous than other welfare benefits, and less generous than other OECD countries (Addison and Teixeira 2003). Those eligible for unemployment insurance benefits (UIB) receive pre-tax benefits equal to 62.4% of previous pre-tax earnings, and the benefits are capped at a level which in 2014 was approximately at the 43rd percentile of the earnings distribution. UIB is paid for a maximum of 104 weeks.

Norway taxes labor income individually, even for married couples. Thus, changes in the labor income—including UIB—of one spouse, does not affect the marginal tax rate of the other. There is an exception when one spouse earns close to zero, but this matters little in my sample where almost all spouses, male and female, earn far more than zero.

2.2 Data

I use a dataset sourced from Norwegian Administrative records, covering the universe of Norwegian residents. The dataset combines wealth and income data from the Norwegian Tax Authorities with a complete registry of unemployment registrations from the Norwegian Labor and Welfare Administration, a registry of marital status, spouses and children, housing transactions from the Norwegian Mapping Authorities, and estimated value of owned housing from Fagereng, Holm, and Torstensen (2018). Each individual has a unique identification number, which is used to link information from the various sources.

The Norwegian Labor and Welfare Administration administers applications for UIB, and administers the payment of these benefits. I use a registry of all individuals who have registered as unemployed with the Norwegian Labor and Welfare Administration, where start and end dates of all unemployment spells are recorded. The Norwegian tax authorities maintain a registry with information on marital status (including cohabitation) and children. The dataset includes the identification number of the spouse, allowing me to construct household-level measures of all variables. Unmarried, but cohabiting couples are not generally registered, unless they have a child together, own housing together, or were previously married. I make no distinction between cohabiting and married couples, as cohabitation is quite common, and is generally regarded to have an equal standing as marriage.

Information on income and wealth is collected for tax purposes, and reported annually. The data is of high quality: Both wealth and income are reported by third-parties (employers, financial institutions, the Labor and Welfare Administration, etc.), and there is no top-coding, and virtually no measurement error. Income is reported by income source (labor earnings, capital income, business income, various benefits including UIB, various pensions, and other government transfers), and assets by asset class (deposits, debt, listed and non-listed stocks, and mutual funds, as well as real assets such as housing). While income is measured as the accumulated sum within the calendar year, assets are reported as a snapshot of the asset position on December 31st. Wealth data was first collected in 1993, and the dataset used in this paper contains annual information on wealth and income from 1993 to 2015.

Information of mutual funds is not available before 1995. Housing wealth is systematically mismeasured in the Norwegian wealth data before 2010. To compensate, I use estimated values for housing wealth from Fagereng, Holm, and Torstensen (2018), where they use a machine learning algorithm to estimate market values for each housing unit in their dataset.

Consumption expenditure ($Exp_{i,t}$) is not directly observed in the data, but can be backed out from the household budget constraint, following Fagereng and Halvorsen (2017) and Eika, Mogstad, and Vestad (2020). To do so, I calculate each household's net change in wealth from one year to the next, adjust for the return earned on each asset class, and subtract this measure of "net active saving" from disposable income:

$$Exp_{i,t} = Y_{i,t} - \sum_k (a^k_{i,t} - a^k_{i,t-1}(1 + \hat{r}^k_{t-1})) + (D_{i,t} - D_{i,t-1}(1 + r^D_{i,t})), \tag{1}$$

Households are indexed by i, and calendar year by t. The variables are defined as follows: Y is income after taxes and transfers (plus inheritance, lottery winnings, and registered inter-vivo transfers), a^k is asset holdings in asset class k (deposits, stocks, mutual funds), r^k is the return on asset k, D is debt, and r^D is interest paid on debt. The return on stocks and mutual funds is unobservable, and I follow Fagereng and Halvorsen (2017) and assume that the return on stocks equals the return on the Oslo Stock Exchange benchmark index (OSEBX), and that the return on mutual funds equals a linear combination of the OSEBX and S&P500. Further details are described in the Supplemental Appendix. Since rental payments will appear as spending, but the consumption of owner-occupied housing does not, there is a risk of bias in the diff-in-diff if households change homeownership status. One proposed fix is to add the value of owner-occupied housing as in (Eika, Mogstad, and Vestad, 2020). This fix does on the other hand imply that the consumption expenditure measure does not reflect actual spending decisions. Instead I keep the spending measure as described by Equation 1, and perform a robustness check where I add the value of owner-occupied housing. In the Supplemental Appendix I show that all main conclusions are unchanged.

Imputed measures of total consumption expenditure are generally regarded to be reliable (see Baker, Kueng, Meyer, and Pagel (2022)), but there is a risk that some consumption expenditure go undetected if the household receives informal, unregistered transfers from friends and family. If there is an increase in private, unobserved transfers after job loss, the estimated consumption expenditure response may be somewhat overstated, which is particularly problematic if the probability of receiving a transfer depends on family type. Norway had an inheritance tax until 2014, and inter-vivo transfers from parents to children were supposed to be recorded, reducing the risk of bias. Further, If there is a family-type bias in who receives informal, unregistered transfers, it seems likely that those more likely to receive are those who are more likely to be worse off. Thus, I would expect that single parents are the most likely to receive private transfers, and couples without children are the least likely. If that is the case, the consumption expenditure response I estimate below would be overstated for single parents, and perhaps some singles without children, but the estimated response is less likely to be overstated for couples without children. This sort of bias, does not alter the conclusions of this paper, but would rather imply that we should think of the estimated differences between couples and singles as a lower bound.

2.3 Sample selection

Table 1: Summary statistics

	Singles		Couples	
	Has children	No children	Has children	No children
Demographics:				
Age	41	38	41	51
Share women	0.48	0.37	0.48	0.53
Children below 18	1.4	0.0	2.0	0.0
Income:				
Earnings	62.2	59.1	63.0	61.0
Spousal earnings			57.7	57.0
Income a.t.	50.3	44.1	48.1	44.3
Spousal income a.t.			47.2	44.7
Wealth:				
Deposits	11.5	18.8	25.8	34.7
Debt	110.2	75.0	195.4	125.6
Stocks	1.0	1.6	3.0	4.5
Mutual funds	1.5	2.0	3.7	4.6
Share homeowner	0.42	0.26	0.78	0.82
Housing wealth	277.7	235.9	504.9	547.1
Net wealth	182.5	184.8	344.2	469.0
Consumption exp.:				
Total consumption exp.	48.9	42.3	87.2	81.5
N	3476	16927	18584	8850

Notes: Each column represents a subset of the full sample. All monetary variables are in 1000 USD, adjusted for inflation with 2014 as baseline year (NOK/USD = 6.3019). Age and the number of children is measured at the beginning of the year of job loss, all other variables are measured two years prior to separation. All wealth variables and consumption expenditure is measured at the household level.

An individual is identified as unemployed if he or she is registered as an unemployed

job-seeker with the Labor and Welfare Administration and receives UIB payments.[1] For the sample, I select individuals that are identified as unemployed between 1997 and 2011, excluding very short spells, temporary furloughs, and those registered as receiving occupational rehabilitation or as being unfit for work on medical grounds. I restrict the analysis to job losers between 25 and 57 years old, to avoid bias arising from transitions from unemployment into early retirement (widely available from 62). If someone has several spells during the sample period, I keep only the first. I require that they are eligible for unemployment benefits, that they have earned at least the basic amount[2] in all five years leading up to the unemployment spell, and that they do not have any unemployment spells in these last five years. These conditions ensure a stable labor market attachment in the years leading up to the unemployment spell, and is crucial for proper estimation of counterfactual outcomes and for the assessment of the pre-trends of earnings and consumption expenditure. As a consequence, individuals selected to the sample has, on average, higher income and higher wealth than the typical unemployed person in Norway.

I define four subgroups of the sample based on pre-job loss family type characteristics: couple/single, and has/does not have children. An individual is labeled as "has children" if they are registered as the parent of a child at the beginning of the year in which they register as unemployed, otherwise they are labeled as childless. An individual is labeled as married or cohabiting if he or she is registered as such at the beginning of the year when they register as unemployed. Otherwise, they are labeled as single. I do not allow them to switch groups, and balance the panel by dropping from the sample everyone who switch marital status during the event window. This is not an innocuous choice, as the leaving-out of those switching from e.g. married to single may exclude those who experience a particularly bad shock. In the Supplemental Appendix, I perform a robustness check using the unbalanced sample, where I only leave out household-year observations where the individual's marital status is not identical to the marital status at the beginning of the year of job loss. The main conclusions are unchanged.

The sample consists of 20,403 single job losers (3,476 of whom have children), and 27,434 married or cohabiting job losers (18,584 of whom have children). The summary statistics of each group are reported in Table 1. There is no age difference between singles and couples for households with children, but for households without children, the age difference is 38 versus 51 years. Married or cohabiting parents in the sample have on average 2.0 children, while single parents in the sample have on average 1.4 children. All groups have roughly 50% women, except childless singles where the share of women is somewhat lower. The average level of earnings of the job loser (measured two years prior to onset of unemployment) is slightly above 60,000 USD in all groups except childless singles who have an average of 59,100 USD.

The wealth position and the balance sheet composition of households differ greatly between the four subsamples. All wealth variables are measured on the household level, with the assumption that all couples pool their disposable wealth. The homeownership rate is highest among couples. While roughly 80% of couples own housing, the number is 42% for

[1] For those interested, the Supplemental Appendix provides further details about the data and the construction of unemployment spells.

[2] Norway has no minimum wage, but the basic amount is a number set by parliament each year that is used to calculate a wide array of welfare benefits. In 2014, the basic amount was equal to 14,023 USD.

single parents and 26% for childless singles. Deposits dominate on the financial asset side of the balance sheet in all subsamples, but the level differs significantly. While childless couples have 34,700 USD in deposits on average (and 9,100 USD in risky financial assets), couples with children have an average of 25,800 USD in deposits (6,700 USD in risky financial assets), childless singles have on average 18,800 USD in deposits (3,600 USD in risky financial assets), and single parents have on average 11,500 USD in deposits (2,500 USD in risky financial assets). Since wealth is measured on the household level, the per-adult liquid wealth position is more equally distributed between couples and singles, but single parents still have significantly less liquid wealth available. The per-adult debt burden is also highest among single parents.

3 Empirical framework

The object I want to estimate is the conditional expectation of the drop in earnings and consumption expenditure each year in the event window. Let $E[y^U_{i,t}|j_i, \mathbf{Z}_{it}]$ be the expected outcome if unemployed, conditional on family type j and control variables Z (age and financial wealth prior to job loss), and $E[y^N_{i,t}|j_i, \mathbf{Z}_{it}]$ the conditional expectation of the outcome if not unemployed. The average treatment effect of the treated (ATT) is given by $E[y^U_{i,t} - y^N_{i,t}|j_i, \mathbf{Z}_{it}]$.

To estimate the conditional expectation of earnings and total consumption expenditure, I exploit the fact that unemployment hits individuals at different times, and implement a difference-in-difference regression design with staggered implementation. I use the following estimation equation:

$$y_{i,t} = \alpha_i + \sum_{k=-5}^{4} \beta_1^k U_i^k + \sum_{k=-5}^{4} \beta_2^k U_i^k \cdot S_i + \sum_{k=-5}^{4} \beta_3^k U_i^k \cdot Children_i + \sum_{k=-5}^{4} \beta_4^k U_i^k \cdot S_i \cdot Children_i + \tau_t^{S_i, Children_i} + \boldsymbol{\gamma'} \mathbf{Z_{it}} + \varepsilon_{i,t} \quad (2)$$

where i indexes the individual, t indexes calendar year, and k indicates time relative to unemployment onset. The dynamic path of the outcome variable in the event window will be given by the coefficients on the relative time indicators U_i^k, which is a dummy variable taking the value 1 for observations k years after the year of unemployment onset, and zero otherwise.

I interact the relative time-dummy with a marital status dummy S_i, which takes the value 1 if the individual is *not* registered as married or cohabiting at the beginning of the year of job separation. In order to capture any differences in the consumption expenditure response of households with or without children, I also interact the relative time-dummy with a dummy variable indicating whether or not the individual is a parent at the time of job loss ($Children_i = 1$ if the individual has children at the beginning of the year $k = 0$), and a triple interaction where the relative time dummy is interacted with both the marital status dummy and the children-dummy.

The estimation equation includes household-fixed effects and time-fixed effects, which are allowed to differ between couples and singles, with and without children. I also include a

set of control variables Z_{it}. From economic theory and previous empirical research, we know that the level of consumption expenditure tends to be hump-shaped over the life cycle. To capture the possibility that the specific growth rate of consumption expenditure at any point in time may also depend on family type, I include as controls family-type specific age profiles. Other studies have shown that the distance to liquidity constraints matters greatly for the marginal propensity to consume out of income shocks, both on impact and in subsequent years. I use liquid financial wealth (deposits + stocks + mutual funds) as a proxy for the distance to liquidity constraints. To capture how liquid wealth may affect the dynamic path of consumption expenditure of households of different family composition differently, I interact financial wealth (measured one year before unemployment onset) with all relative-time dummies, and interact these with the marital status dummy and the children-dummy in the same way as described above.

As the recent literature on identification in difference-in-difference estimation models with two-way fixed effects and staggered implementation has shown, it is necessary to take steps to ensure proper identification. I follow Dobkin, Finkelstein, Kluender, and Notowidigdo (2018) and Andersen, Jensen, Johannesen, Kreiner, Leth-Petersen, and Sheridan (2023) and leave out two reference categories among the dummies indicating year before the onset of unemployment ($k = -5$ and $k = -2$). This identification strategy is proposed in Borusyak, Jaravel, and Spiess (2024) for settings where it is reasonable to assume that only a linear trend is not identified after controlling for confounders. Identification is ensured as long as the dependent variable in years $k = -2$ and $k = -5$ is uncorrelated with the timing of the event, after controlling for two-way fixed effects, financial wealth, and family-type specific age profiles. Given that grace periods in Norway are typically 3 months, it seems a reasonable assumption that the household has no certain knowledge of the upcoming job loss two years before. All standard errors are clustered at the household level.

4 Empirical results

The estimated ATTs are reported in Figure 1 (income after taxes and transfers) and Figure 2 (consumption expenditure). The figures report the dynamic paths of earnings and consumption expenditure. Each panel reports the outcomes of ordinary least squares (OLS) estimation of Equation 2 for one outcome variable.

The estimation results are reported in Figure 2, where each panel plots the conditional expectation of the change in the outcome variable in the event window. Panel (a) plots the dynamic path of individual income after taxes and transfers around the time of unemployment onset. There is no pre-trend in after-tax income for any of the four subgroups between 5 and 2 years prior to job loss. In the year before unemployment onset, however, there is a slight decline in income after taxes and transfers for all four subgroups, likely reflecting poor results in businesses preceding the decision to lay off employees, or lack of end-of-year bonus payments to employees already notified about the upcoming lay-off.

All groups see their earnings decline significantly in the year of unemployment onset, and the decline continues into year 1, for all groups except childless singles. This additional dip in earnings is also found in (Fagereng, Onshuus, and Torstensen, 2024), and is a product of time aggregation in the data, caused by unemployment spells lasting into the next calendar year. I

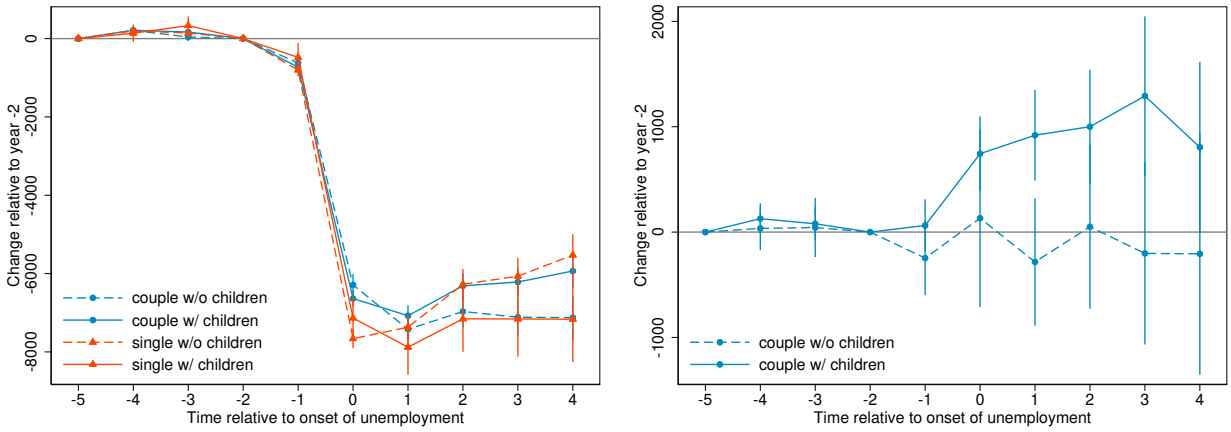

(a) Individual income after taxes and transfers **(b)** Spousal income after taxes and transfers

Notes: All monetary variables are in USD, adjusted for inflation with 2014 as base year. Year zero is the year of unemployment onset. Family type (couple/single, has/does not have children) is determined based on status on January 1st of year zero, and individuals do not move between groups. Balanced panel.

Figure 1: After-tax income around the time of job separation

find a weak recovery of earnings starting in the second year after unemployment registration, but earnings remain significantly below pre-job loss levels even after UIB exhaustion after 104 weeks.

The magnitude of the earnings loss is not equally distributed across household types. Couples with children and singles without children seem to face the least severe earnings losses, and after four years their income is around 6,000 USD below the pre-job loss trend. Singles with children and couples without children face both the largest and the most long-lasting earnings loss. After four years, the average income in these groups is still more than 7,000 USD below trend, while couples (both with and without children) remain around 6,000 USD below trend.

The extant literature on family labor supply has argued that spousal labor supply should rise as a response to income shocks, in order to mitigate the impact on household income. Indeed, as reported in Panel (b) of Figure 1, I find a small increase in spousal income after taxes and transfers, but it is concentrated among couples with children. Couples without children exhibit no change in spousal income after taxes and transfers. Panel (a) of Figure 2 reports the change in household income after taxes and transfers. The figure shows that even after the increase in spousal income, the job loss represents a significant drop in household income.

Panel (b) of Figure 2 reports the dynamic path of consumption expenditure around the time of unemployment for each family-type subgroup. The consumption expenditure response of single job losers with children stand out from the rest. While the other family-type subgroups exhibit a sharp reduction in spending in year 0 with subsequent gradual recovery, single job losers with children reduce consumption expenditure gradually, and continue to reduce consumption expenditure compared to the trend in all five years from the year of unemployment onset to four years after.

Importantly, there is no meaningful difference between couples and singles without children. (Blundell, Pistaferri, and Saporta-Eksten, 2016) discuss how couples trade off self-insurance through spousal labor supply and self-insurance through precautionary saving, and argue that one reason studies of the Added Worker Effect tend to find modest effects is that they do not take household wealth into account. My findings support this conclusion in part but suggests that most family insurance comes from couples' enhanced ability to self-insurance through precautionary saving: There is no difference in the consumption expenditure response of a married or cohabiting job loser and a single job loser (without children) if they have the same level of financial wealth.

Couples with children exhibit a somewhat smaller reduction in consumption expenditure compared to couples without children. Carneiro, Garcia, Salvanes, and Tominey (2021) discuss how fluctuations in parental income during childhood affects intergenerational mobility, and conclude that a steady flow of investments in children maximizes their adult outcomes. The observation that households with children reduce spending less than households without children is in line with this theory, if parents seek to minimize the effects of the income loss on child investments. Theories of childhood investments do not, however, explain the difference in consumption expenditure responses between single parents and married or cohabiting parents, since couples should be at least as able to smooth consumption expenditure as singles are for the same level of financial wealth.

As there are significant differences in the magnitude of the earnings loss between couples and singles, with and without children, I also calculate the cumulative consumption expenditure response, $\Phi_{j,s}$. The cumulative consumption expenditure response is constructed as the estimated cumulative drop in consumption expenditure, normalized by the estimated cumulative earnings loss:

$$\Phi_{j,s} = \frac{\sum_{k=-1}^{s} E[Exp_{it}^{U} - Exp_{it}^{N} | j_i, X_i]}{\sum_{k=-1}^{s} E[Y_{it}^{U} - Y_{it}^{N} | j_i, X_i]}, \tag{3}$$

where j denotes family type and s denotes horizon ($s = 0, ..., 4$). These coefficients are plotted in Panel (c) of Figure 2. Also here singles with children stand out from the rest, particularly in the first year of unemployment. Over time, the cumulative consumption expenditure responses converge, although the pattern remains. The numbers of the cumulative consumption expenditure response cannot be interpreted as marginal propensities to consume, and since I control for financial wealth in the regression they do not represent average consumption expenditure responses either. Instead, the numbers tells us that, for a household with average financial wealth and average age, the consumption expenditure response of single job losers with children is four times that of single job losers without children.

Since I do not observe cohabitation where the couple does not have a child together, and does not own housing together, it is possible that there are individuals among those registered as single who actually cohabit with someone. If that is the case, there will be some contamination from couples to singles in the estimation of the consumption expenditure and earnings path. Thus, the estimated consumption expenditure response of singles with children can be read as an upper bound of the true consumption expenditure response, and the estimated difference between couples and singles with children can be read as a lower bound on the true difference.

The results I report here indicate that both marital status and the presence of children matter for the degree of consumption smoothing. Even though some households use spousal labor supply to mitigate the income loss, the empirically observed pattern suggests a need for supplementary explanations as to why singles with children exhibit a lower average consumption expenditure response than any other family-type group. One plausible supplementary mechanism is pre-committed consumption.

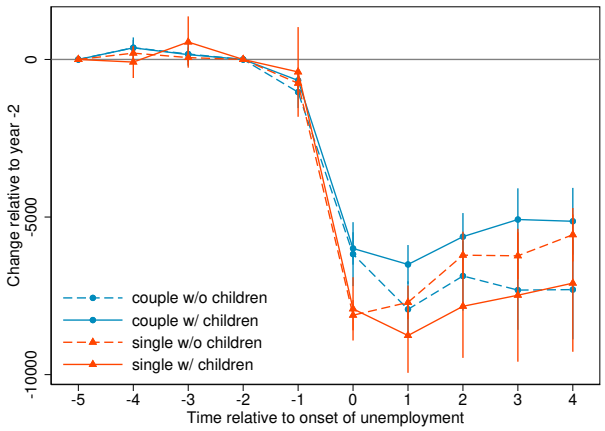

(a) Household income after taxes and transfers

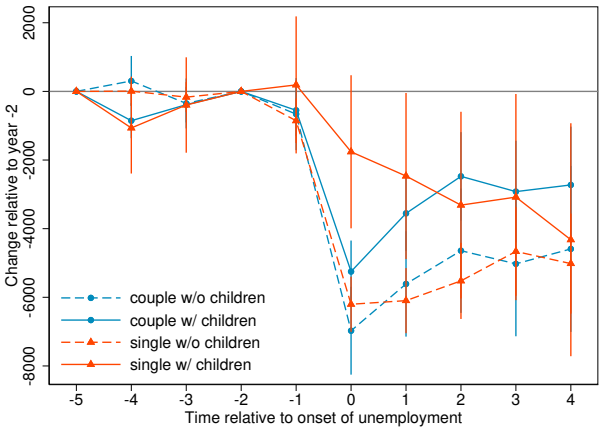

(b) Total consumption expenditure

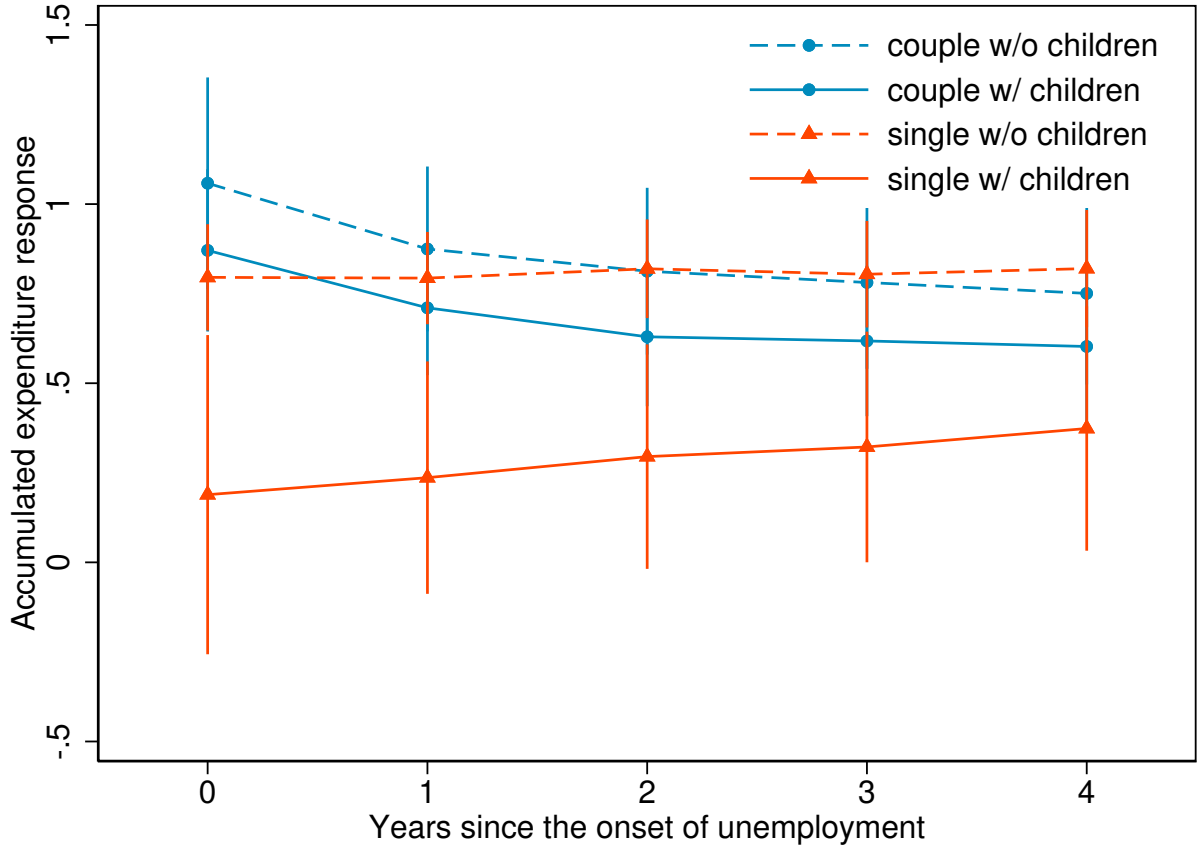

(c) Cumulative consumption exp. response

Notes: All monetary variables are in USD, adjusted for inflation with 2014 as base year. Consumption expenditure is measured at the household level. Year zero is the year of unemployment onset. Family type (couple/single, has/does not have children) is determined based on status on January 1st of year zero, and individuals do not move between groups. Balanced panel. Total consumption expenditure is imputed using the hh. budget equation. Panel (c) plots the cumulative drop in consumption exp. normalized by the cumulative drop in household income after taxes and transfers.

Figure 2: The consumption expenditure response to job loss

5 Model

In this section, I argue that the differences in consumption behavior can be explained by varying degrees of pre-committed consumption, and I construct a toy model to illustrate the mechanism.

The mechanism is simple: Some types of consumption cannot be reduced past a certain limit. When a negative income shock pushes the household towards this lower bound, the marginal propensity to consume falls, because the household is forced to concentrate the spending drop in the freely adjustable consumption category. Since it is more painful for the household (given decreasing marginal utility of consumption) to concentrate the reduction of spending in one consumption category, constrained households become more willing to reduce saving in order to mitigate the spending reduction.

One way to think about a lower bound on certain types of consumption is this: Households have made consumption choices in the past which restrict their behavior today. Consumption of certain goods may be elastic in the long run, but inelastic in the short run. One can think of housing, which is lumpy and carries a transaction cost, and which demands a certain level of expenditure each month in rent, local taxes or municipal fees, as well as basic energy usage and necessary maintenance. Renters (at least in Norway) cannot leave their contracts on short notice, and housing takes time to sell. On top of that, households may have low willingness to move if the children go to school nearby.

In the model, the constraint takes the form of an exogenously given lower bound on one of the consumption goods. Importantly, households are not necessarily bound by the constraint at every point in time. One might think that there are certain housing related expenditures that are voluntary, such as paying for the removal of snow, renovations, "excessive" energy use, etc., or one may think that some homeowners (and some renters) have the option to rent out (or sublet) part of the house. Other goods can also have features resembling a lower bound on consumption: After school activities for children often carry a fixed monthly or bi-annual fee, although parents are free to choose how much to spend on new equipment. Households can vary their electricity use, but there exists a minimum level necessary for cooking, lighting and heating. More generally, the lower bound captures the idea that there are some expenses households are willing to or able to reduce, and others they are not willing to or not able to drop. A lower bound can for example capture the unwillingness of parents to let an income shock affect child-related spending.

I make two assumptions: The commitment level depends on the number of people in the household, through economies of scale (or a demographic shifter). The commitment level also depends on how many adults there are in the household to share the expenses required to satisfy the lower bound. These two assumptions mean that the presence of children increase the commitment level, while the presence of a spouse decreases the commitment level (one can think of children as a commitment good which is also a public good for the parents). I show how these two assumptions about consumption commitments can reduce the marginal propensity to consume (MPC) of single parents compared to other households, consistent with the empirical results.

5.1 Model set-up

In the interest of brevity, I abstract from uncertainty, labor supply, and intra-household bargaining. The model has only one time-period, but in the absence of uncertainty all forms of saving for future consumption or precautionary motives is nested in the warm-glow bequest function. The household maximizes lifetime utility subject to a budget constraint stating that total spending on the two consumption goods and the bequest cannot exceed lifetime resources, which is assumed exogenous and fixed.

$$\max U(a_i, z_i, b_i) = \alpha \frac{(a_i)^{1-\sigma}}{1-\sigma} + (1-\alpha)\frac{(z_i)^{1-\sigma}}{1-\sigma} + \theta \frac{(b_i)^{1-\sigma}}{1-\sigma} \tag{4}$$

$$\text{s.t.} \quad a_i + z_i + b_i = Y \tag{5}$$

The utility function U is the per-capita, or per-adult, lifetime utility. The agent draws utility from their consumption of each of the two consumption goods. Good a_i is freely adjustable, while good z_i is subject to a fixed lower bound, $z_i \geq z \cdot \phi_i$. The agent is free to increase consumption of good z_i.

The commitment level depends on a demographic shifter, ϕ_i. To formalize ideas, I assume that the demographic shifter varies across family types:

$$\phi_i = \frac{e(N_i^{adult}, N_i^{children})}{N_i^{adult}}, \tag{6}$$

where e_i is a function of the number of adults (N^{adult}) and the number of children ($N^{children}$) in the household. As functional form for $e(\cdot)$, one can for example use the OECD or the OECD-modified equivalence scales, or the square root of family members. All three functions yield a significantly higher commitment level for single parents than for any other family type.[3]

5.2 The constrained and unconstrained marginal propensity to consume

Let the household be hit by an unexpected income shock. The MPC will depend on whether or not the household is constrained by the lower bound on good z. The unconstrained (denoted by U) and constrained (denoted by C) MPCs are given by

$$MPC^U = \frac{\alpha^{\frac{1}{\sigma}} + (1-\alpha)^{\frac{1}{\sigma}}}{\alpha^{\frac{1}{\sigma}} + (1-\alpha)^{\frac{1}{\sigma}} + \theta^{\frac{1}{\sigma}}} > MPC^C = \frac{\alpha^{\frac{1}{\sigma}}}{\alpha^{\frac{1}{\sigma}} + \theta^{\frac{1}{\sigma}}} \tag{7}$$

The MPC is always higher for unconstrained households. The intuition is straightforward: The household prefers to reduce consumption of both goods proportionally after

[3] If one uses the OECD-modified equivalence scale (square root function), and assume that all households with children have two children each, the formula gives the following numbers: ϕ_i(couple without children) = 0.75 (0.7), ϕ_i(couple with children) = 1.05 (1), ϕ_i(single without children) = 1 (1), and ϕ_i(single with children) = 1.6 (1.7). For some smoking-gun evidence, the Survey of Consumer Expenditure in Norway reports that single parents share of total expenditures allocated to housing-related expenses is almost 20% higher (6 pp.) than that of couples with children. Both the OECD-modiefied and the square root function is by construction (or nature) normalized to equal one for singles without children.

the income shock. Since marginal utility is decreasing, it is more painful to concentrate the reduction in consumption in the adjustable good, and a household that faces a large enough income shock to be pushed into the constraint, faces a trade-off between two choices (a and b) instead of three (a, b and z). Thus, the constrained household maintains a suboptimally high level of consumption (and spending) today, at the cost of reduced saving. Since single households with children have a higher commitment level (relative to household income), they are more likely to be constrained.

An important take-away from this exercise is that a lower MPC does not necessarily reflect better self-insurance or lower welfare consequences in a life-cycle perspective. If the constraint keeps consumption of the commitment good suboptimally high, welfare is in fact reduced.

6 Conclusion

In this paper I have presented novel evidence on how the consumption expenditure response to unemployment varies between households with different family composition. I find that single parent job losers stand out with significantly lower consumption expenditure responses to unemployment than both childless singles, and couples with and without children. Couples with children exhibit somewhat increased spousal income after taxes and transfers after unemployment onset, which may indicate that they use spousal labor supply as a self-insurance mechanism. Still, couples with children reduce consumption expenditure more following job loss than single parents.

I show that these findings can be rationalized in a model of consumption and saving behavior where a household allocates their disposable income between consumption of a freely adjustable consumption good, consumption of a good subject to a lower bound ("pre-committed consumption") and a warm-glow bequest. The model replicates the empirical results if the pre-commitment level is allowed to depend on family type according to the number of family members and the number of income-earners. I show that a lower bound on the consumption of certain goods can decrease the marginal propensity to consume of constrained households.

These results have important implications for economic theory. They suggest a novel explanation for the mixed empirical results on the Added Worker Effect: The need for self-insurance through spousal labor supply also depends on the probability of being constrained by the lower bound on the commitment good. Second, researchers seeking to understand or model consumption insurance should be aware of how consumption commitments affect the marginal propensity to consume.

References

Addison, J. T., and P. Teixeira (2003): "The economics of employment protection," *Journal of Labor research*, 24(1), 85–128.

Andersen, A. L., A. S. Jensen, N. Johannesen, C. T. Kreiner, S. Leth-Petersen,

and A. Sheridan (2023): "How do households respond to job loss? lessons from multiple high-frequency datasets," *American Economic Journal: Applied Economics*, 15(4), 1–29.

Autor, D., A. Kostøl, M. Mogstad, and B. Setzler (2019): "Disability benefits, consumption insurance, and household labor supply," *American Economic Review*, 109(7), 2613–2654.

Baker, S. R., L. Kueng, S. Meyer, and M. Pagel (2022): "Consumption imputation errors in administrative data," *The Review of Financial Studies*, 35(6), 3021–3059.

Blundell, R., L. Pistaferri, and I. Saporta-Eksten (2016): "Consumption inequality and family labor supply," *American Economic Review*, 106(2), 387–435.

——— (2018): "Children, time allocation, and consumption insurance," *Journal of Political Economy*, 126(S1), S73–S115.

Borusyak, K., X. Jaravel, and J. Spiess (2024): "Revisiting event-study designs: robust and efficient estimation," *Review of Economic Studies*, p. rdae007.

Browning, M., and T. F. Crossley (2001): "Unemployment insurance benefit levels and consumption changes," *Journal of public Economics*, 80(1), 1–23.

Browning, M., and M. Ejrnæs (2009): "Consumption and children," *The Review of Economics and Statistics*, 91(1), 93–111.

Carneiro, P., I. L. Garcia, K. G. Salvanes, and E. Tominey (2021): "Intergenerational mobility and the timing of parental income," *Journal of Political Economy*, 129(3), 757–788.

Dobkin, C., A. Finkelstein, R. Kluender, and M. J. Notowidigdo (2018): "The economic consequences of hospital admissions," *American Economic Review*, 108(2), 308–52.

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

Fagereng, A., and E. Halvorsen (2017): "Imputing consumption from Norwegian income and wealth registry data," *Journal of Economic and Social Measurement*, 42(1), 67–100.

Fagereng, A., M. B. Holm, and K. N. Torstensen (2018): "Housing wealth in Norway, 1993–2015," *Journal of Economic and Social Measurement*, pp. 1–17.

Fagereng, A., H. Onshuus, and K. N. Torstensen (2024): "The consumption expenditure response to unemployment: Evidence from Norwegian households," *Journal of Monetary Economics*, p. 103578.

Ganong, P., and P. Noel (2019): "Consumer spending during unemployment: Positive and normative implications," *American Economic Review*, 109(7), 2383–2424.

Gerard, F., and J. Naritomi (2021): "Job displacement insurance and (the lack of) consumption-smoothing," *American Economic Review*, 111(3), 899–942.

Gruber, J. (1997): "The Consumption Smoothing Benefits of Unemployment Insurance," *The American Economic Review*, 87(1), 192–205.

Heathcote, J., K. Storesletten, and G. L. Violante (2010): "The macroeconomic implications of rising wage inequality in the United States," *Journal of political economy*, 118(4), 681–722.

Landais, C., and J. Spinnewijn (2021): "The value of unemployment insurance," *The Review of Economic Studies*, 88(6), 3041–3085.

Love, D. A. (2010): "The effects of marital status and children on savings and portfolio choice," *The Review of Financial Studies*, 23(1), 385–432.

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