

CEBI WORKING PAPER SERIES

Working Paper 12/20

HOW DO HOUSEHOLDS RESPOND TO JOB
LOSS? LESSONS FROM HIGH-FREQUENCY
DATA SETS

Ager Lau Andersen
Amalie Sofie Jensen
Niels Johannesen
Claus Thustrup Kreiner
Søren Leth-Petersen
Adam Sheridan

ISSN 2596-44TX

CEBI

Department of Economics
University of Copenhagen
www.cebi.ku.dk

How Do Households Respond to Job Loss? Lessons from Multiple High-Frequency Data Sets*

Asger Lau Andersen[†], Amalie Sofie Jensen[‡], Niels Johannesen[†],
Claus Thustrup Kreiner[†], Søren Leth-Petersen[†] and Adam Sheridan[†]

April 2020

Abstract

How do households respond to job loss, and which self-insurance channels are most important? By linking customer data from the largest bank in Denmark with information from government administrative registers, we quantify a broad range of responses to job loss in a unified empirical framework. Two response margins stand out: during the first 24 months after job loss, households reduce spending by 30% of the income loss while reduced saving in liquid assets accounts for 50%. Other response margins highlighted in the literature - spousal labor supply, private transfers, home equity extraction, mortgage refinancing, and consumer credit - are less important.

*Acknowledgments: We thank Sumit Agarwal, Peter Ganong, Alessandro Martinello, Andrea Weber, as well as seminar participants at the New Consumption Data workshop in Copenhagen, Danmarks Nationalbank, and the CEPR Fourth European Workshop on Household Finance at Lund, Sweden for helpful comments and discussions. The activities of the Center for Economic Behavior and Inequality (CEBI) are financed by a grant from the Danish National Research Foundation. Financial support from the Economic Policy Research Network (EPRN), the Candys foundation, the Danish Council for Independent Research, and the Carlsberg Foundation is also gratefully acknowledged. This research was facilitated by Adam Sheridan's Industrial PhD project, jointly financed by Danske Bank, which provided data access, and Innovation Fund Denmark. Danske Bank did not review the conclusions of this paper before circulation and the opinions expressed are those of the authors alone and do not represent the views of Danske Bank. The use of the data for this project was approved by the Danish Data Protection Agency, Statistics Denmark, and the Ministry of Industry, Business and Financial Affairs.

[†]Center for Economic Behavior and Inequality (CEBI), Department of Economics, University of Copenhagen.

[‡]Department of Economics, Princeton University.

People who lose their jobs typically experience a large and persistent drop in income (Jacobson, LaLonde, and Sullivan 1993; Davis and Wachter 2011; Kawano and Lalumia 2015; Flaaen, Shapiro, and Sorkin 2019; Seim 2019). Several studies show that household consumption also drops at the onset of unemployment (e.g. Gruber 1997; Browning and Crossley 2001, 2009; Hendren 2017; Ganong and Noel 2019; Landais and Spinnewijn 2019). But compared to the drop in disposable income, the impact on consumption is typically moderate. For example, recent evidence based on bank transaction data for U.S. households shows that the drop in spending is one third of the drop in income at the onset of unemployment (Ganong and Noel 2019).

This suggests that households self-insure against job loss to a significant extent, but how they do this remains an open question.¹ Any drop in income that is not matched by a drop in spending must be met by an increase in funds from other sources, a reduction in saving, or an increase in borrowing. Various response margins within each of these categories have been proposed by different strands of literature. Households can raise money inflows from other sources through an expansion of *spousal labor supply* (Lundberg 1985; Cullen and Gruber 2000; Stephens 2002; Hardoy and Schøne 2014; Halla, Schmieder, and Weber 2020) or through increases in *private transfers* from family and friends (Altonji, Hayashi, and Kotlikoff 1997; McGarry 2016). They can reduce *debt repayments* and/or increase *borrowing* by taking up alternative mortgage products or tapping into home equity (Hurst and Stafford 2004; Cocco 2013), or by borrowing more through unsecured lines of credit (Sullivan 2008). Finally, households may reduce saving by running down their buffer-stock of *liquid assets* (Carroll 1997; Basten, Fagereng, and Telle 2016).

In this paper, we examine the empirical relevance of each of these self-insurance responses, as well as the effect on spending, in a unified empirical framework. Existing studies typically analyze a single response margin, with samples, data, and methods varying across studies. In contrast, our study provides a comprehensive assessment of the relative importance of each margin by analyzing all responses for the same sample of

¹We use the term self-insurance to refer to all responses that weaken the contemporaneous impact of shocks to income on household consumption. These include responses that counteract the drop in income, such as higher spousal labor supply, as well as pure consumption-smoothing responses that transfer liquidity across time, i.e., borrowing or saving.

households, applying the same definition of job loss, and using the same research design.

We do this by linking transaction-level data from the largest bank in Denmark with administrative data from multiple government registers. The combined data, covering the period January 2009 to December 2016, has several advantages: First, the monthly frequency of the data enables sharp empirical identification in an event study design. Second, the detailed transaction data allows us to construct precise and comprehensive measures of monthly spending and saving at the micro level.² Third, the data includes detailed demographic information, including the identity of spouses and cohabiting partners. This enables us to identify spousal labor supply responses and to study outcomes at the household level, thus eliminating potential measurement error stemming from intra-household effects. Fourth, the possibility of linking the transaction data to other administrative data enables us to address concerns about completeness and representativeness, which are key issues when using transaction data from a single provider (Baker 2018; Ganong and Noel 2019). In particular, we study a subsample of households who do not bank elsewhere to assess the influence of transactions not captured in our data (completeness), and we use detailed socio-demographic information for all individuals in the population to assess the impact of selection (representativeness).

We find a significant and persistent income loss for the average person affected by job loss, despite extensive social insurance.³ The effect on monthly household disposable income is a drop of 30% on impact and close to 10% two years after the job loss. Household spending also drops significantly, but far less than income: Two years after job loss the

²A small but growing number of studies use transaction data from banks or financial aggregators to obtain high-frequency spending measures. Most closely related to our study are Ganong and Noel (2019) and Gerard and Naritomi (2019) who use event study designs to study spending responses to unemployment in the U.S. and Brazil, respectively. The focus in these studies, and in most other studies using transaction data (Kueng 2018; Baker 2018; Gelman et al. 2014; Olafsson and Pagel 2018; Baker and Yannelis 2017; Gelman et al. 2015), is on analyzing excess sensitivity of consumption to income changes. We focus on analyzing the overall importance of the different response margins of households hit by unemployment shocks, which is made possible by the combination of transaction data and government administrative data. Other studies impute spending from annual information about income and changes in assets and liabilities (e.g., Browning and Leth-Petersen 2003; Landais and Spinnewijn 2019) or use self-reported measures of spending (e.g., Parker and Souleles 2019; Kreiner, Lassen, and Leth-Petersen 2019).

³Since our focus is on the overall importance of the different response margins, the analysis concentrates on the responses to job loss for the average person. This reflects a mixture of responses for households experiencing short spells and long spells of unemployment.

cumulative spending drop amounts to 30% of the cumulative income loss, leaving a gap of 70% that reflects the effects of household self-insurance. We find that this gap is filled by lower accumulation of liquid assets (~50%), increases in private transfers and other inflows (~10%), higher spousal labor supply (~5%), and lower net debt repayments (~5%). The combined effect of these responses closely matches the income loss in each month following job loss, suggesting that the analysis captures all relevant margins. We conclude that reduced saving in liquid assets is the single most important way that households self-insure against job loss.

The paper proceeds as follows: Sections 1-3 present background information on the institutional setting, data, and empirical methods, while section 4 presents the main results of our analysis. Section 5 presents additional analyses showing, among other things, that our main findings are robust to changes addressing concerns about completeness and representativeness, and that income and spending responses in the Danish data are similar to recent findings for U.S. households. Section 6 concludes.

1 The Danish Institutional Setting

Labor market: The Danish labor market is characterized by flexible hiring and firing rules for employers combined with high income security for employees (Andersen and Svarer 2007). Dismissing workers is low-cost for employers compared to many other countries (OECD 2013). The notice period is typically 3 to 6 months for white-collar workers but shorter for blue-collar workers (Scheuer and Hansen 2011). This means that many laid-off workers have a few months to prepare for the impending drop in wage income.

The unemployment insurance system is partly funded by worker's contributions and partly by the government. Members of the insurance system receive benefits worth 90% of the pre-unemployment wage up to a cap of around \$3,000 per month. Because of this cap, actual compensation rates are considerably lower for many wage earners.⁴ Benefits are

⁴In 2010, 91% of all wage earners in the age group studied in this paper had wage income exceeding the cap. 34% had wage income exceeding twice the size of the cap.

taxed the same way as labor income. The maximum duration of UI benefits is two years. This provides high income security compared to many other countries, including the U.S. where the maximum duration is typically six months. Unemployed workers who are ineligible for UI benefits may receive a means-tested basic social transfer of around \$1,700 per month, with a supplement for families with children. Other government transfers, such as housing support and child benefits, are also income-dependent and may help reduce the income drop after job loss.

Financial markets: Households in Denmark buy financial services from two main types of financial institutions: retail banks and specialized mortgage banks. Retail banks offer a wide range of financial services, including deposit accounts and various credit facilities. Mortgage banks only offer mortgage loans financed by covered bonds and they offer both fixed and adjustable-rate mortgages, with and without interest-only payments, and with a duration of up to 30 years. At origination, mortgage borrowers always face the current rate in the covered bond market. The maximum allowed loan-to-value ratio for mortgage loans is 80%.⁵ Fixed-rate mortgages can be refinanced at a fairly low cost (Andersen et al. 2015). Mortgage debt is full recourse in Denmark, and defaults are very rare (Kreiner, Leth-Petersen, and Willerslew-Olsen 2020).

Payment system: The payments landscape in Denmark limits the problem of “invisible” cash transactions when using bank transaction data to measure spending. Card usage is higher in Denmark than in any other European country and checks are no longer in use (Danmarks Nationalbank 2017). Almost all bill payments are made electronically, with over 95% of Danish households paying bills by direct debit (Danish Competition and Consumer Authority 2014). Only 16% of the value of point-of-sale retail transactions is in cash, compared to 39% for the U.S. (Danmarks Nationalbank 2017, Greene and Stavins 2018).

⁵Homeowners can go beyond the 80% limit by taking out additional collateralized loans from retail banks, but these are more expensive.

2 Data Construction

We link monthly information about individuals from six administrative data sources using a unique personal identity number assigned to all Danes at birth or first residence.⁶ The combined data allows us to track individuals and their spouses from January 2009 to December 2016. This section describes the data and the construction of key variables. Appendix A contains further details about variable definitions.

Employment and layoffs: We identify employment, job separations, periods of unemployment, and the individual’s main employer using population-wide monthly payroll records collected by the Danish Tax Agency. Employers have to report wages for each employee to the tax agency and government agencies must report income transfers. Evasion is minimal (Kleven et al. 2011; Alstadsæter, Johannesen, and Zucman 2019). The records are used for tax collection and for computation of official employment statistics. Each record contains information about the gross amount paid, the month in which the amount was earned, a unique employer ID and sector code (for salary payments), and a transfer program code (for income transfers).

Disposable income, spending, saving, and non-mortgage debt repayments: We use transaction and account records from the largest bank in Denmark (“Danske Bank”, henceforth just referred to as “the bank”). The data is similar to the JP Morgan Chase data used by Ganong and Noel (2019) in their recent study of spending through unemployment spells in the U.S. More than one third of the Danish adult population are in our data. The records contain information on all deposit and loan account balances, as well as detailed information about all transactions in each account.

We adopt a broad definition of household disposable income, equal to all external inflows to the household’s bank accounts. To construct this measure, we focus on specific types of account inflows: First, direct deposits, which will include all labor, pension, and government transfer income. Second, person-to-person transfers that originate from outside the household, which include transfers from extended family, friends, and other

⁶Technically, the data providers send the data to Statistics Denmark who de-identify and store it on secure servers with remote access for researchers. Card et al. (2010) highlight the Danish micro data and data infrastructure as a blueprint for data construction.

external inflows. Third, cash deposits into accounts. We then break household disposable income down into salary income for each household member, income transfers from the government, and other. To do this, we combine the transaction data with the payroll data from the tax authority, as described above, allowing us to identify which income payments are from employers or government agencies (see Appendix C for details).

For spending, we focus on three types of payments – debit or credit card, in-store mobile, and bill – and cash withdrawals from ATMs. These categories account for almost 80% of all outflows in a given month for the average household (see Appendix B). For card and in-store mobile payments, we can categorize the type of spending using the four-digit Merchant Category Code (MCC) of the recipient business. MCCs are an international standard for classifying merchants by the type of goods and services they provide. For bill payments, we know the identity of the creditor for each transaction. The bank maintains a grouping of creditors into categories that correspond to the MCC grouping and we use this to categorize bill payments into the same groups as for card and mobile payments. To construct our baseline measure of monthly expenditure, we sum outgoing transactions by each of the payment methods and all cash withdrawals from ATMs. We use the categorization of spending to remove tax and debt payments, as well as fees paid to the bank. Appendix Figures A2 and A3 show a high level of correspondence between selected components of our spending measure and corresponding aggregate time series from official sources. This suggests that our spending data does well in terms of accuracy and representativeness of the broader population.

We measure net repayments on non-mortgage loans as the change in end-of-month balances on loan accounts. Positive values correspond to net repayment, negative to net borrowing. Saving in liquid assets is the sum of the change in end-of-month balances on deposit accounts and the net outflow across all accounts stemming from financial securities trades. By focusing on flows in and out of the portfolio, instead of changes in the value of the portfolio, we isolate active saving responses from movements due to capital gains and losses.

Household structure: The population register provided by Statistics Denmark con-

tains annual demographic information about the entire Danish population. The data includes information about age and gender of all individuals and, importantly, the personal ID numbers of spouses (including cohabiting partners) in each calendar year. This enables us to study outcomes at the household level rather than at the individual level where measurement can be biased by invisible intra-household effects if, for example, a spouse purchases more of the consumption goods of the household when unemployed than when employed.⁷ The identity of spouses is also needed to identify spousal labor supply responses to the unemployment shocks.

Bank relationships: The Danish Tax Agency collects end-of-year information about all interest-bearing loans and deposits held in Danish banks by Danish residents. The data is third-party reported by financial institutions, and it contains account-level information about balances as well as a unique identifier for the reporting institution. With this data, we can address a key concern when working with transaction data from a single provider, namely whether the available data provides a complete picture of the activities of households who may also transact through other banks or intermediaries.

Mortgage loans: We use a loan-level data set collected from Danish mortgage banks by the Danish Ministry of Business and Growth and the Danish central bank. It provides an end-of-year snapshot of all active mortgage loans to private individuals in Denmark. It contains detailed information about the date of origin, time to maturity, original and outstanding balance, and interest rate on each loan. It also describes the type of loan, including whether it is a fixed- or adjustable-rate loan and whether it is an interest-only loan. Combining the end-of-year snapshot in a given year with that of the previous year, we can detect whether there were any changes to an individual's portfolio of mortgage loans during the calendar year. We use the information on dates-of-origin for the new loan(s) to determine exactly when this change happened, and thus construct a high-frequency data set with information about mortgage loans held at the end of each month (see Appendix D for details).

⁷We find that over 30% of actual couples are not linked to each other in the bank data, where a link is inferred from the existence of a joint account or a household identifier based on self-reporting relationships. Without information on household structure from the population register, these individuals would be treated as separate households.

Mass layoffs: In a sensitivity analysis, we use firm-level information about mass-layoffs reported by firms to the Ministry of Employment to isolate involuntary job losses. The data and results are described in Appendix E.

3 Sample Selection and Research Design

We define an unemployment shock as a situation where the salary payments from the individual's main employer cease and total gross wage income drops below 1,000 DKK (\$190, January 2010 price level). The first month where these conditions are met is defined as the month of the job loss. To focus on transitions into unemployment, we require that the individual receives unemployment benefits or social insurance at some point between months -1 and 3 relative to the month of the job loss, and that he or she does not receive early retirement, sickness or parental leave benefits in this time period. To focus on shocks, rather than recurring events, we restrict attention to individuals who have gross wage income of at least 10,000 DKK (\$1,920) for at least 18 consecutive months before the job loss and do not return to the same employer within three months after the job loss.

The observation window for the event analysis is 18 months before to 24 months after the month of the job loss. The unit of analysis is individual-by-month, but outcome variables are generally measured at the household level by summing over the adult members.

Our analysis sample consists of individuals born between 1950 and 1979 who experienced an unemployment shock between July 2009 and December 2015.⁸ We focus on stable households by requiring that the individual either stays single or has the same spouse in all of the months in which they enter the analysis. We exclude individuals if they or their spouse bought or sold real estate, or if they worked at the same firm as their spouse prior to the job loss. The former restriction is imposed because housing trades are associated with massive financial transactions, making it difficult to isolate

⁸The payroll data covers January 2008 to March 2016. Since the definition of an event requires 18 months of data pre-event and 3 months post-event, this means that the unemployment shock must happen between July 2009 and December 2015 to satisfy all criteria.

the saving- and spending responses to the unemployment event. The latter restriction is imposed because correlated income shocks stemming from the same employer prevent us from cleanly examining the spousal labor supply effect of job loss.

Finally, to produce our main results, we limit the sample to households who are active customers at the bank. Following previous literature, we define an active customer as a person with at least five spending outflows in each month of the observation window (Ganong and Noel 2019). For couples, we require that both partners are active customers.

Using outcomes based on account and transaction data of customers in one financial institution raises concerns about whether the sample is representative of the full population, and whether one captures the complete set of relevant transactions (Baker 2018). Our combined data makes it possible to address these concerns. Table 1 provides summary statistics for individuals in different samples, measured six months before the month of job loss. Column (1) shows that our gross sample consists of 66,844 individuals before restricting it to active customers in the bank. Introducing this restriction reduces the sample to 10,002 individuals, as shown in column (2). The active customers are on average slightly better educated, more likely to be single, work in the public sector, and reside in the capital region than individuals in the gross sample, and they also earn slightly higher incomes. But, overall, the two samples are quite similar in their socio-economic characteristics.

The active customers hold a non-trivial share of their deposits (5%) and non-mortgage loans (11%) with other retail banks. Column (3) shows statistics for a subsample of *exclusive* customers, defined as active customers who do not have deposits or loans at other retail banks at any time during the observation window. In section 5, we present robustness analysis showing that our results are unchanged if we instead use this subsample, alleviating concerns about lack of completeness. The same conclusion holds if we instead impose representativeness by reweighting observations in the sample of active customers to match the socio-economic background characteristics of the gross sample shown in column (1).

We estimate the dynamic effects of job loss using a standard event study model:

Table 1: Sample selection and summary statistics

	(1)	(2)	(3)
	Gross sample	Active customers (baseline sample)	Exclusive customers
No. of individuals	66,844	10,002	5,224
	----- Sample means -----		
Female	0.43	0.47	0.48
Age	46.2	46.6	46.1
Couple	0.67	0.59	0.52
Capital region	0.33	0.44	0.42
Higher education	0.23	0.28	0.27
Primary sector	0.01	0.01	0.01
Manufacturing	0.19	0.15	0.15
Homeowner	0.65	0.63	0.59
Annual gross income for person who lost job (DKK)	371,621	394,499	375,019
Share of hsh. bank deposits held at other banks	0.71	0.05	0.00
Share of hsh. retail bank loans held at other banks	0.71	0.11	0.00

Column (1) shows statistics for the gross sample with no requirements on customer status at Danske Bank. Column (2) shows statistics for the baseline sample of active customers, i.e., individuals who have at least five outgoing spending transactions in each month of the event observation window and whose partner (if any) satisfies the same criterion. Column (3) is for the sample of exclusive customers, i.e., active customers who have no deposits or loans at other retail banks and whose partner (if any) satisfies the same criterion. All variables measured in month -6 relative to the month of job loss, except the following: Annual gross income, measured over the calendar year in which month -6 occurs; shares of household loans and deposits held at other banks, measured at end of calendar year before month -6. Appendix Table A1 provides additional summary statistics for each of the three samples.

$$y_{it} = \gamma_t + \delta_i + \sum_h \beta_h \cdot \mathbb{1}[e_{it} = h] + \epsilon_{it}, \quad (1)$$

where i indexes individuals, t indexes calendar months, y_{it} is the outcome of interest, γ_t is a year-by-calendar month fixed effect, δ_i is an individual fixed effect, and e_{it} is event time, defined as distance in months to the unemployment event, with negative values indicating that individual i has not yet lost his/her job in month t . The coefficients of interest are the β_h , which summarize the dynamics of the outcome variable around the time of the job loss. Each coefficient expresses the difference in the outcome in event month h relative to the pre-event level. We include observations up to 18 months before and 24 months after the month of job loss. The identifying assumption is that the time to/since job loss is uncorrelated with omitted factors affecting the outcome, conditional on calendar month and individual fixed effects. To ensure identification in this type of model, we need two reference categories (e.g. Dobkin et al. 2018) so we leave out the dummy variables

for $h = -18$ and $h = -6$. To limit the influence of extreme outliers, we censor the outcome variables at the 2.5 and 97.5 percentiles within each event month. We cluster standard errors at the individual level to allow for arbitrary forms of heteroskedasticity and autocorrelation across observations for the same individual.

4 Main Results


Figure 1 plots estimation results from equation 1. It shows the direct impact of job loss on monthly income (markers) and response margins (bars) for the average household on a time line centered around the month of job loss.⁹ To facilitate comparisons across income classes, we normalize nominal outcomes by measuring them relative to the household's pre-event disposable income, defined as the average disposable income in months -18 to -3.

Job loss has a large effect on the affected person's after-tax salary income (black dots in Figure 1). Salary payouts are higher than normal in the two months before job loss, due to sizable severance payments for some individuals, but then drop sharply at layoff. The average drop corresponds to about half of the households' pre-event disposable income, reflecting that most households also have income from other sources, including salary income earned by the spouse. Salary payouts recover steadily in the following months, as some of the laid-off individuals return to employment, but they never catch up to the pre-displacement level within the two-year observation window of our analysis. In month 24, the gap remains almost half its initial size. This is in line with previous findings of persistent income losses following the transition into unemployment (Jacobson, LaLonde, and Sullivan 1993; Davis and Wachter 2011; Kawano and Lalumia 2015; Flaaen, Shapiro, and Sorkin 2019; Seim 2019). To obtain an expression of the total cumulative impact on income over the full observation window, we sum the estimates from month -5 to 24.¹⁰ The total effect on after-tax salary over these months amounts to a loss of seven months

⁹The key estimates underlying the figure, including their standard errors, are reported in Appendix Table A2. Appendix Figure A2 shows the responses separately for each outcome.

¹⁰We include the estimates for months -5 to -1 to include effects taking place immediately before the month of job loss, including severance payments. The cumulative estimates are reported in Appendix Table A2.

Figure 1: Income, spending, and self-insurance responses to job loss



The figure shows estimation results from the event study model (1) of the effects of job loss on a range of outcomes. All outcomes are measured relative to the household’s average disposable income between event months -18 and -3 and winsorized at the 2.5 and 97.5 percentiles within each event month. Disposable income is defined as all external inflows to the household’s bank accounts. Estimates for the direct effect on income are illustrated by series of shaped markers. The series labeled “After-tax salary” shows coefficient estimates from a regression with after-tax salary payments for the household member who lost his/her job as the outcome. The series labeled “After-tax salary + gov. transfers” is the sum of these coefficients and the corresponding ones from a regression with income from government transfers as the outcome. Estimates for behavioral responses are shown in bars. We estimate coefficients for each outcome in separate regressions and illustrate the sums of these coefficients by the height of the stacked bars. In calculating these sums, each component is signed so that a positive value indicates a change that contributes to compensating for the loss of income. The series labeled “Borrowing and reduced debt repayments” shows the sums of coefficients for two separate outcomes: non-mortgage loan net repayments and mortgage loan repayments. Figure A4 and Table A2 in the Appendix show, respectively, full dynamics and selected coefficient estimates with standard errors for each separate outcome.

of pre-event household disposable income.

Social insurance provides significant income compensation. In Figure 1, the drop in after-tax income becomes much smaller when we include income transfers from the government (red dots). Over the full observation window, we estimate that these transfers compensate for two thirds of the salary loss for the average household. Thus, the total direct effect on income, cumulated over months -5 to 24, is a loss equivalent to about 2½ months of pre-event household disposable income.

The bars in Figure 1 show how households respond to compensate for this income loss. The joint effect of the compensating responses is illustrated by the height of the stacked bars. The fact that the estimated responses match the direct income loss almost perfectly – the height of the stacked bars is very close to the red dot in every month – suggests that we capture all relevant response margins.

Starting with household *spending*, we find a clear negative effect in all 24 months following job loss (black bars). Over the entire period, we estimate that the reduction in spending corresponds to 30% of the direct income loss. This aligns with the finding in existing literature that the spending response to job loss, while significant, is substantially smaller than the corresponding effect on income.¹¹ Figure A5 in the Appendix shows how the strength of this response varies across expenditure categories: In line with theory, we find that households maintain spending on consumption commitments (Chetty and Szeidl 2007, 2016), as proxied by utility bills, but cut down substantially on discretionary goods, as proxied by restaurant and bar spending. In between these extremes, the percentage drop in grocery spending is about the same size as the drop in overall spending. This suggests that the overall spending drop reflects an actual reduction of consumption and not merely self-insurance through postponement of luxury goods or durables purchases (Browning and Crossley 2000, 2009).

What are the most important ways that households self-insure against the income loss

¹¹Gerard and Naritomi (2019) study the effects of job displacement in Brazil and find that displaced workers *increase* spending at the time of layoff, despite experiencing a drop in the longer term. This is because most displaced workers in their sample receive a large lump-sum payment, in the form of government-mandated severance pay, when they are laid off. As shown in Appendix Figure A7, we find a similar result if we limit the sample to individuals who receive significant severance pay.

associated with job loss? The *spousal labor supply* effect provides little self-insurance for the average household (blue bars). The cumulative change in spousal salary payouts over the full period covers 6% of the lost income of the laid-off person. The small increase in spousal labor earnings is entirely along the intensive margin, with no significant effect on spouses' employment rate (Appendix Figure A6).

The effect on *private transfers and other inflows* is somewhat stronger (yellow bars). The cumulative increase in these inflows over the full observation window is 0.3 months of pre-displacement disposable income, corresponding to 12% of the direct income loss. This reflects informal insurance through gifts and loans from extended family and friends but can also capture inflows stemming from sales of real assets or consumer durables.

We find only a modest impact of job loss on *borrowing and debt repayments* (white bars). This effect is strongly concentrated in month 1 after displacement where we observe a sizable increase in non-mortgage borrowing. For mortgage loans, we find a statistically significant – but economically modest – decrease in average monthly debt repayments. This is driven by a small share of the households who convert their mortgage loans to loan types with lower debt service costs, whereas we find no impact on home equity extraction through mortgage refinancing (Appendix Figure A8). Over the full period, these borrowing adjustments compensate for less than 5% of the direct income loss.

Saving in liquid assets (gray bars) is the single-most important self-insurance response margin. It accounts for 49% of the cumulated direct income loss. This is significantly larger than for any other response, economically as well as statistically.¹² The accumulation of liquid assets spikes upward just before the job loss, mirroring the increase in income from severance pay, and then drops significantly at the onset of unemployment. The effect continues to be large throughout the period. Note that this does not reflect that households continue to decumulate assets. In fact, net saving is roughly zero for the average household from month five onwards. This should be compared to a counterfactual of *positive* net saving – that is, accumulation of liquid assets – in the absence of job loss (Appendix Figure A9).

¹²The null that this share is numerically equal to the corresponding share for private transfers and other inflows (the second-largest response) has a p-value of 0.016 against a two-sided alternative.

5 Additional Analyses

Representativeness and completeness

Table 2 explores how our key estimates change as we alter sample selection criteria and estimation methods to address concerns about completeness and representativeness. All columns report estimates of cumulative effects from month -5 to 24 relative to the month of job loss. Odd-numbered columns show these effects measured in months of pre-event household disposable income, while even-numbered columns express them relative to the cumulated direct income loss. Columns (1)-(2) provide the estimates for our baseline sample. Columns (3)-(4) show the corresponding estimates for the subsample of households that are exclusive customers at the bank. Since these households do not bank elsewhere, results for this subsample should be free of any problems related to lack of completeness in the transaction data. Finally, columns (5)-(6) provide estimates from regressions where observations are reweighted to make our sample of active bank customers match the demographic characteristics of the gross sample drawn from the full population, thus addressing concerns about representativeness.¹³

Across all columns, the estimates of the cumulated direct income loss (including the effect on government income transfers) are in the range of 2-3 months of household disposable income. The combined effect of the behavioral responses that we consider is also stable across columns and always close to the estimated direct income loss. There is some variation when it comes to the relative importance of each response but the overall conclusions are robust: Household spending drops by about 25-40% of the direct income loss, suggesting substantial self-insurance. The most important self-insurance response for all samples is reduced saving in liquid assets, which accounts for about 40-50% of the direct income loss. The compensating effects from spousal labor supply responses, borrowing, and loan repayments are small and/or insignificant. Finally, the estimates for

¹³We construct the weights as the inverse predicted probabilities from a probit of active customer status on the demographic characteristics reported in Table 1. The dependent variable in the probit model is a dummy for belonging to the sample of active customers. The independent variables are dummy variables for age (five-year intervals), sex, couple, capital region residence, higher education, sector of employment before layoff (seven categories), and homeownership.

Table 2: Robustness analysis

	Baseline		Exclusive customers		Weighted regressions	
	(1)	(2)	(3)	(4)	(5)	(6)
----- Cumulative effects, months -5 to 24 -----						
	Rel. to disp. inc. before job loss	Percent of direct income loss	Rel. to disp. inc. before job loss	Percent of direct income loss	Rel. to disp. inc. before job loss	Percent of direct income loss
<i>Direct income effects</i>						
[1] Salary, affected person	-6.93 (0.14)		-7.94 (0.21)		-6.47 (0.15)	
[2] Gov. transfers, household	4.56 (0.08)		5.15 (0.11)		4.32 (0.08)	
[3] Direct income loss (-[1] - [2])	2.37 (0.139)	100.0% (0.0%)	2.79 (0.20)	100.0% (0.0%)	2.15 (0.15)	100.0% (0.0%)
<i>Response margins</i>						
[4] Salary, spouse	0.15 (0.07)	6.2% (2.9%)	0.22 (0.09)	7.9% (3.5%)	0.16 (0.08)	7.5% (4.0%)
[5] Private transfers and other inflows	0.29 (0.15)	12.1% (6.4%)	0.32 (0.21)	11.5% (7.9%)	0.25 (0.15)	11.4% (7.2%)
[6] Spending	-0.72 (0.15)	-30.3% (6.8%)	-1.10 (0.22)	-39.6% (8.8%)	-0.57 (0.16)	-26.4% (7.8%)
[7] Net saving in liquid assets	-1.16 (0.30)	-49.2% (12.4%)	-1.07 (0.44)	-38.3% (15.8%)	-1.12 (0.31)	-52.0% (14.4%)
[8] Non-mortgage loan net repaym.	-0.05 (0.10)	-2.0% (4.5%)	0.06 (0.14)	2.0% (5.2%)	-0.04 (0.11)	-1.8% (5.4%)
[9] Mortgage loan repayments	-0.06 (0.01)	-2.7% (0.6%)	-0.05 (0.02)	-1.8% (0.7%)	-0.06 (0.01)	-2.9% (0.7%)
[10] Total ([4] + [5] - [6] - [7] - [8] - [9])	2.42 (0.27)	102.4% (10.6%)	2.71 (0.39)	97.0% (12.9%)	2.20 (0.29)	102.0% (12.2%)
Number of individuals	10,002	10,002	5,224	5,224	10,002	10,002

The table reports results from robustness checks of the results illustrated in Figure 1 and reported in Table A2. Columns (1)-(2) show results for our baseline sample, reproduced from columns (5)-(6) of Table A2; columns (3)-(4) for the subsample of exclusive customers, as defined in section 3; and columns (5)-(6) for regressions where observations are reweighted so that our sample of active customers matches the characteristics of the gross sample shown in column (1) of Table 1. All estimates are based on regressions where the reported outcomes are measured relative to the household's average disposable income in the pre-event months. Odd-numbered columns report the sum of coefficients for event months -5 to 24 from such regressions. Even-numbered columns report the ratios between these sums and the corresponding sum for the direct income loss shown in row [3]. Standard errors (in parentheses) are estimated by bootstrapping with 500 replications. The bootstrapping procedure is carried out with resampling of individuals, rather than individual observations, to account for heteroskedasticity and autocorrelation within observations for the same individual. Graphical illustrations of the results in columns (3)-(4) and (5)-(6), paralleling those shown in Figure 1, are shown in Appendix Figures A10 and A11.

private transfers and other inflows suggest that increases in such inflows compensate for 11-12% of the direct income loss across samples.

Mass layoffs

A potential concern is that some individuals in our sample resigned voluntarily, in which case we may not cleanly identify the effects of involuntary job loss. In an attempt to focus on actual job displacements, we conduct a sensitivity analysis in which we limit the sample to individuals who lose their jobs concurrently with mass layoffs at their firms, identified from the firm mass-layoff data collected by the Ministry of Employment. The results align with our main findings (see Appendix E for details).

Comparing spending responses to U.S. estimates

Denmark stands out from most other countries in the duration of UI benefits, as described in section 1. To assess whether we can expect our results to extend to other countries, we compare them to high-frequency estimates of dynamic income and spending responses for U.S. households reported in Ganong and Noel (2019). For this exercise, we follow their analysis by estimating the percentage change in household income and spending for a “survivor” sample in which individuals exit the estimation sample once they are re-employed.¹⁴

Figure 2 shows such estimates for our sample. We find results very similar to those for U.S. households in the first six months after layoff where UI benefits are available in both countries. After the pre-event spike, household income drops sharply in the Danish sample and then stabilizes around 20% below its pre-event level, similar to the 20-25% drop that Ganong and Noel (2019) find for U.S. households. Household spending is 5-10% below its pre-displacement level in our sample, which is nearly identical to the U.S. case. The similarity in results for these key outcomes suggests that the overall level of

¹⁴Ganong and Noel 2019 condition on staying unemployed because their aim is to explore excess sensitivity in spending responses at benefit exhaustion. In our main analysis, we follow individuals after job loss unconditional on continued unemployment in order to measure the overall impact of job loss for the average person.

self-insurance is comparable in the two countries, at least within the six-months horizon.

Responses are different in the two countries after the first six months, at which point UI benefits are exhausted in the U.S. but continue in Denmark. Households in our sample experience slow gradual declines in income and spending, whereas U.S. households see a sharp drop in income and cut spending drastically to 20% below pre-displacement level. Thus, this simple comparison suggests that households in these two countries behave similarly when they are exposed to similar conditions, and differently when they are not. Moreover, the initial similarity and subsequent divergence in results indicate that the longer duration of UI benefits in Denmark significantly reduces the consumption loss from unemployment. But the key observation for our purposes is that there is an almost proportional relationship between the drop in income and the drop in spending following job loss both across countries and within countries over time. The ratio between the drop in spending and the drop in income hovers around 30% in our sample and around 28% for U.S. households (see Appendix Figure A13).¹⁵ This suggests that the aggregate self-insurance response is around 70% of the drop in income irrespective of its size. Insofar as each self-insurance response is also proportional to the income change then our result about the relative importance of the different types of self-insurance responses also apply more generally, independent of differences in the generosity of the UI benefits system.


6 Concluding Remarks

The self-insurance responses to unemployment shocks studied in this paper fall into two categories: Reduced *saving in liquid assets* and increased *borrowing* are pure consumption-smoothing responses, which bring forward consumption without changing overall consumption possibilities. In contrast, increases in *spousal labor supply* and *private transfers* from family or friends mitigate the impact of the shock on the consumption possibilities of the household. Our analysis shows that the first type of responses is quantitatively far more important than the second type. Thus, apart from UI benefits, households have


¹⁵Ganong and Noel (2019) do not report direct estimates for the ratio of spending changes to income changes. We calculate such ratios from their estimates of relative spending changes, relative income changes, and initial (pre-layoff) mean values of spending and income.

Figure 2: Relative effects of job loss on household income and spending, conditional on staying unemployed

(a) Income (salary+gov. income transfers)



(b) Spending



The figure plots the estimated effects of job loss on household income and spending, measured relative to their own sample averages in event months -18 to -2. The sample is a dynamic sample of individuals that stay unemployed. In event months -18 to 0, this includes everyone in the baseline sample. For event month $t > 0$, it includes those who have not returned to employment at any point between month 0 and month t . Employment status is defined as having gross wage income above 10,000 DKK (at January 2010 price level). Vertical lines indicate 95% confidence intervals. Standard errors are estimated allowing for clustering at the level of the individual.

little insurance against the overall consumption loss caused by unemployment.

Households in our sample smooth consumption mainly through adjustments to net saving in liquid assets. The modest effect on borrowing and debt repayments may seem surprising when considering the substantial liberalization and innovation in credit markets over recent decades. It is, however, consistent with evidence in DeFusco and Mondragon (2020) showing that the unemployed have high latent demand for mortgage refinancing but are constrained by their employment status. Under this interpretation of our results, households are willing, but unable, to borrow during unemployment, leaving reduced saving in liquid assets as the important margin of adjustment.

References

- Alstadsæter, Annette, Niels Johannesen, and Gabriel Zucman. 2019. “Tax Evasion and Inequality.” *American Economic Review* 109 (6):2073–2103.
- Altonji, Joseph G, Fumio Hayashi, and Laurence J Kotlikoff. 1997. “Parental altruism and inter vivos transfers: Theory and evidence.” *Journal of Political Economy* 105 (6):1121–1166.
- Andersen, Steffen, John Y. Campbell, Kasper Meisner Nielsen, and Tarun Ramadorai. 2015. “Inattention and Inertia in Household Finance: Evidence from the Danish Mortgage Market.” Working Paper 21386, National Bureau of Economic Research.
- Andersen, Torben M. and Michael Svarer. 2007. “Flexicurity – Labour Market Performance in Denmark.” *CEPR Economic Studies* 53 (3):389–429.
- Baker, Scott R. 2018. “Debt and the Response to Household Income Shocks: Validation and Application of Linked Financial Account Data.” *Journal of Political Economy* 126 (4):1504–1557.
- Baker, Scott R and Constantine Yannelis. 2017. “Income changes and consumption: Evidence from the 2013 federal government shutdown.” *Review of Economic Dynamics* 23:99–124.
- Basten, Christoph, Adreas Fagereng, and Kjetil Telle. 2016. “Saving and Portfolio Allocation Before and After Job Loss.” *Journal of Money, Credit and Banking* 48 (2-3):293–324.
- Browning, Martin and Thomas F Crossley. 2000. “Luxuries are easier to postpone: A proof.” *Journal of Political Economy* 108 (5):1022–1026.
- . 2001. “Unemployment insurance benefit levels and consumption changes.” *Journal of Public Economics* 80 (1):1–23.
- . 2009. “Shocks, stocks, and socks: Smoothing consumption over a temporary income loss.” *Journal of the European Economic Association* 7 (6):1169–1192.

- Browning, Martin and Søren Leth-Petersen. 2003. “Imputing consumption from income and wealth information.” *Economic Journal* 113 (488):282–301.
- Card, David, Raj Chetty, Martin S Feldstein, and Emmanuel Saez. 2010. “Expanding access to administrative data for research in the United States.” Tech. rep., White paper for the initiative "Future Research in the Social, Behavioral & Economic Sciences", U.S. National Science Foundation.
- Carroll, Christopher D. 1997. “Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis.” *Quarterly Journal of Economics* 112 (1):1–55.
- Chetty, Raj and Adam Szeidl. 2007. “Consumption commitments and risk preferences.” *Quarterly Journal of Economics* 122 (2):831–877.
- . 2016. “Consumption commitments and habit formation.” *Econometrica* 84 (2):855–890.
- Cocco, João F. 2013. “Evidence on the Benefits of Alternative Mortgage Products.” *Journal of Finance* 68 (4):1663–1690.
- Cullen, Julie Berry and Jonathan Gruber. 2000. “Does unemployment insurance crowd out spousal labor supply?” *Journal of labor Economics* 18 (3):546–572.
- Danish Competition and Consumer Authority. 2014. “Betalingservice.” Tech. rep., Danish Competition and Consumer Authority.
- Danmarks Nationalbank. 2017. “Danish households opt out of cash payments.” *Analysis no. 24* .
- Davis, Steven J. and Till von Wachter. 2011. “Recessions and the costs of job loss.” *Brookings Papers on Economic Activity* 41 (2):1–72.
- DeFusco, Anthony and John Mondragon. 2020. “No job, no money, no refi: Frictions to refinancing in a recession.” *Journal of Finance* (forthcoming).

- Dobkin, Carlos, Amy Finkelstein, Raymond Kluender, and Matthew J. Notowidigdo. 2018. “The Economic Consequences of Hospital Admissions.” *American Economic Review* 108 (2):308–52.
- Flaaen, Aaron, Matthew D. Shapiro, and Isaac Sorkin. 2019. “Reconsidering the consequences of worker displacements: firm versus worker perspective.” *American Economic Journal: Macroeconomics* 11 (2):193–227.
- Ganong, Peter and Pascal Noel. 2019. “Consumer spending during unemployment: positive and normative implications.” *American Economic Review* 109 (7):2383–2424.
- Gelman, Michael, Shachar Kariv, Matthew D Shapiro, Dan Silverman, and Steven Tadelis. 2014. “Harnessing naturally occurring data to measure the response of spending to income.” *Science* 345 (6193):212–215.
- . 2015. “How individuals smooth spending: Evidence from the 2013 government shutdown using account data.” Tech. rep., National Bureau of Economic Research.
- Gerard, François and Joana Naritomi. 2019. “job displacement insurance and (the lack of) consumption-smoothing.” *NBER Working Paper No. 25749*. .
- Greene, Claire and Joanna Stavins. 2018. “The 2016 and 2017 surveys of consumer payment choice: summary results.” Tech. rep., Federal Reserve Bank of Boston.
- Gruber, Jonathan. 1997. “The consumption smoothing benefits of unemployment insurance.” *American Economic Review* 87 (1):192–205.
- Halla, Martin, Julia Schmieder, and Andrea Weber. 2020. “Job displacement, family dynamics and spousal labor supply.” *American Economic Journal: Applied Economics* (forthcoming).
- Hardoy, Inés and Pål Schøne. 2014. “Displacement and household adaptation: insured by the spouse or the state?” *Journal of Population Economics* 27 (3):683–703.
- Hendren, Nathaniel. 2017. “Knowledge of Future Job Loss and Implications for Unemployment Insurance.” *American Economic Review* 107 (7):1778–1823.

- Hurst, Erik and Frank P Stafford. 2004. "Home is where the equity is: Mortgage refinancing and household consumption." *Journal of Money, Credit, and Banking* 36 (6):985–1014.
- Jacobson, Louis S, Robert J LaLonde, and Daniel G Sullivan. 1993. "Earnings losses of displaced workers." *American Economic Review* 83 (4):685–709.
- Kawano, Laura and Sara Lalumia. 2015. "How Income Changes During Unemployment. Evidence from Tax Return Data." *Journal of Human Resources* 52 (2):418–456.
- Kleven, Henrik Jacobsen, Martin B Knudsen, Claus Thustrup Kreiner, Søren Pedersen, and Emmanuel Saez. 2011. "Unwilling or unable to cheat? Evidence from a tax audit experiment in Denmark." *Econometrica* 79 (3):651–692.
- Kreiner, Claus Thustrup, David Dreyer Lassen, and Søren Leth-Petersen. 2019. "Liquidity Constraint Tightness and Consumer Responses to Fiscal Stimulus Policy." *American Economic Journal: Economic Policy* 11 (1):351–379.
- Kreiner, Claus Thustrup, Søren Leth-Petersen, and Louise Willerslew-Olsen. 2020. "Financial Trouble Across Generations: Evidence from the Universe of Personal Loans in Denmark." *Economic Journal* 130 (625):233–262.
- Kueng, Lorenz. 2018. "Excess Sensitivity of High-Income Consumers." *Quarterly Journal of Economics* 133 (4):1693–1751.
- Landais, Camille and Johannes Spinnewijn. 2019. "The Value of Unemployment Insurance." *Working paper. London School of Economics* .
- Lundberg, Shelly. 1985. "The added worker effect." *Journal of Labor Economics* 3 (1):11–37.
- McGarry, Kathleen. 2016. "Dynamic aspects of family transfers." *Journal of Public Economics* 137:1–13.
- OECD. 2013. "Protecting jobs, enhancing flexibility: A new look at employment protection legislation." OECD Publishing.

- Olafsson, Arna and Michaela Pagel. 2018. “The liquid hand-to-mouth: Evidence from personal finance management software.” *Review of Financial Studies* 31 (11):4398–4446.
- Parker, Jonathan A. and Nicholas S. Souleles. 2019. “Reported Effects versus Revealed-Preference Estimates: Evidence from the Propensity to Spend Tax Rebates.” *American Economic Review: Insights* 1 (3):273–290.
- Sabelhaus, John, David Johnson, Stephen Ash, David Swanson, Thesia Garner, John Greenlees, and Steve Henderson. 2013. “Is the Consumer Expenditure Survey representative by income?” Tech. rep., National Bureau of Economic Research.
- Scheuer, Steen and Kristian Rune Hansen. 2011. “Funktionærloven: En lønmodtagerlov for de fleste?” *Samfundsøkonomen* (2):24–31.
- Seim, David. 2019. “On the incidence and effects of job displacement: Evidence from Sweden.” *Labour Economics* 57:131–145.
- Stephens, Melvin, Jr. 2002. “Worker displacement and the added worker effect.” *Journal of Labor Economics* 20 (3):504–537.
- Sullivan, James X. 2008. “Borrowing during unemployment unsecured debt as a safety net.” *Journal of Human Resources* 43 (2):383–412.

Appendix

A Variable Definitions

Income variables

Disposable income: Sum of monthly incoming transactions to bank accounts owned by the affected person and/or the spouse. The following transaction types are included: Direct deposits, person-to-person transfers from outside the household, and cash deposits. We categorize these inflows into the following subcategories (see Appendix C for details):

Salary income, affected person: Sum of monthly salary payment inflows to bank accounts owned by the person affected by job loss. Salary payments to joint accounts are attributed to the spouse if they come from the spouse's employer, and otherwise to the affected person.

Salary income, spouse: Sum of monthly salary payment inflows to bank accounts owned by the spouse. Salary payments to joint accounts are attributed to the spouse if they come from the spouse's employer, and otherwise to the affected person.

Government income transfers: Sum of monthly government income transfer inflows to bank accounts owned by the affected person and/or the spouse.

Private transfers and other inflows: The residual of disposable income minus salary (affected person and spouse) and government income transfers.

Spending

Total spending: The sum of outgoing transactions from the household's bank accounts using either of the payment methods card, mobile phone, and bill, plus cash withdrawals. Outflows categorized as tax or debt payments are excluded. We aggregate to the household level by summing spending for the person affected by job loss and the spouse, if any. Outflows from joint accounts are split evenly between the

two account owners before summing to avoid double-counting. See Appendix B for further details.

Utilities: The value of the subset of transactions in total spending measure with MCC “4900”, “4812”, “4814”, “4821”, “4899”, or bill payment label “utilities”, “elec”, “gas”, “water”, “heating”, “internet”, “cable TV”, “telephone”, or “TV license”

Restaurant and bar spending: The value of the subset of transactions in total spending measure with MCC “5813”, “5462”, “5811”, “5812”, or “5814”.

Groceries: The value of the subset of transactions in total spending measure with MCC “5411”, “5422”, “5441”, “5499”, or “5921”, or bill payment label “groceries”.

Net saving and debt repayments

Net saving in liquid assets: The sum of (1) the change in end-of-month balances on deposit accounts at the bank owned by the affected person or the spouse, and (2) outflows minus inflows to all accounts from transactions with type code “securities trade”.

Net repayments on non-mortgage loans: The change in end-of-month balances on loan accounts (with amounts owed coded as negative balances) owned by the affected person or the spouse.

Mortgage loan repayments: Average monthly mortgage payments with current mortgage loans. Calculated by determining the average monthly payment over a full calendar year for each loan, then summing across all mortgage loans that the household had at the end of the current month. See Appendix D for further details.

B Measuring Household Spending from Transaction Data

The basis for our measure of spending is raw transaction data from the bank's records. Each record holds information about the time and type of transaction and the amount transacted. For card transactions and bill payments, information about the type of recipient is provided in the form of Merchant Category Codes (MCCs). Before making it available to us, the bank aggregates the raw transaction to daily totals within each combination of customer account, transaction type and recipient category.

We include three types of outgoing transactions in our spending measure: Card payments (including payments initiated via mobile phone applications), bill payments and cash withdrawals. For the average person in our gross sample of active bank customers (defined as individuals who have at least five spending transactions in each month of the calendar year), these transaction types account for 80% of all transactions leaving the household in a given month. The remaining outflows include transactions that do not reflect consumption, for example fee payments to the bank and financial security purchases, and uncategorized bank transfers where the purpose is unobservable.

In a next step, we use recipient MCCs to exclude tax and debt repayments, which are not considered as spending. We then sum all the remaining outgoing transactions to construct a monthly spending measure for each individual person. For couples, household-level variables are constructed by aggregating spending for the two spouses. We split outgoing transactions from joint accounts evenly between the account owners to avoid double-counting.

Figure A1 shows the development in average quarterly spending per household for the gross sample of active bank customers, broken down by transaction type. The share of spending done by card or mobile phone transactions rises over the analysis period, from 43% in 2009Q1 to 57% in 2016Q4. Conversely, the share of cash spending gradually falls from 16% to 9% over the same years, while bills account for about 35-40% of total spending throughout the period. Cash payments account for 14% of total spending via

cash or cards in 2016. In comparison, a 2017 household survey by the Danish central bank found a value-weighted cash payment share of 16% of total cash and card payments (Danmarks Nationalbank 2017).

Figure A2 compares the time series development in monthly card spending per person in our gross sample of active bank customers vs. the full population. For the latter, we use aggregate statistics on card transactions and the number of adults in the population published by Statistics Denmark. The two series follow each other extremely closely, suggesting that our spending measure is accurate in timing and that the sample of active customers does well in terms of representing trends in the broader population.

Figure A3a shows average levels of total annual household spending across income groups for households in our sample of active bank customers and compares them to estimates from Statistics Denmark’s consumer expenditure survey. Figure A3b shows the same comparison for annual grocery spending. When averaging over households in all income groups, we find that the measures of both total and grocery spending are very similar across the two data sources, suggesting a high degree of completeness in our transaction data. Looking across groups, we see a slightly steeper income gradient in our measures of both total and grocery spending than in the survey-based measures, perhaps because of disproportional underreporting at the top of the income distribution in the latter (Sabelhaus et al. 2013). But, overall, there is a strong correspondence between the two data sources in this dimension, suggesting that our transaction data measures capture cross-sectional variation in spending well.

C Categorizing Inflows to Bank Accounts

We measure household disposable income as the sum of direct deposits, person-to-person transfers and cash deposits flowing into the household’s bank accounts, excluding transactions between the household’s own accounts. We break this measure down into salary payments for the person affected by job loss, salary payments for the spouse (if any), government income transfers, and other. There are two main steps in this process, which

we describe in detail below: first, we construct a mapping from employer IDs to the IDs of the bank branches that they use to pay out salary, and similarly for the government agencies that pay out income transfers. Second, we look at each individual's incoming transactions and use this mapping to identify payments coming from employers and government agencies. For example, if person A works for company B, which uses bank branch C for its salary payments, then we interpret all payments from bank branch C going into person A's account as salary payments from company B.

We start by linking employers and government agencies to the registration number of the bank account(s) that they use to pay out salary and income transfers, respectively. A registration number is a four-digit Danish national bank code. Each number is associated with a unique bank *branch*, but branches may have more than one registration number (e.g., one for business customers' accounts and one for personal customers' accounts). There are more than 3,000 unique numbers across all banks. Some large customers have their own unique number. For example, all payments from the central government come from accounts with the same unique registration number, which is used solely for this purpose.

We link each employer and government agency to a registration number in the following way: First, for each employer/agency j and each month t , we use the payroll data from the Danish Tax Agency to identify all individuals in our sample of bank customers who appear on the employer's/agency's payroll. Second, for each bank registration number, we use the transaction data to compute the share of individuals in that group who received a payment from an account with that registration number in month t . We record the registration number with the highest share as the one associated with payments from employer/agency j in month t .¹⁶

In the transaction data, we code an incoming direct deposit as a salary payment if the sender's bank registration number is associated with an employer that the recipient works for according to the payroll data. That is, if person A appears on the payroll of employer B in month t and receives a payment from an account with a registration

¹⁶If the central government registration number is the top rank and j is not a central government agency, we use the second-ranked registration number

number that has been linked to employer B through the mapping procedure described above, then we conclude that this is a salary payment from employer B to person A. In addition, we also interpret transactions with certain type codes (e.g. “salary transfers”) as salary payments.

We code an incoming transaction as a government income transfer if either of the following conditions is satisfied: i) The sender’s bank registration number is linked to a government agency that the person received money from in that month according to the payroll data; ii) The bank registration number is the one used by the central government, and the person does not work for the central government (in which case we code it as a salary payment).

D Constructing Monthly Mortgage Data from Annual Snapshots

The mortgage data set provides an end-of-year snapshot of all active mortgage loans to private individuals in Denmark in each year from 2009 to 2015. We use this to construct monthly measures of the type of mortgage loans in the household’s portfolio and the average monthly payment on each of these loans.

We start by mapping the household’s portfolio of active loans in each month during the year by comparing the end-of-year snapshot with that of the previous year. Each loan has a unique ID that allows us to track it over time. If a loan appears in both snapshots, we conclude that it must have been part of the household’s portfolio all year. If it appears only in the most recent snapshot, we use information about the date of origination to infer when it entered the portfolio. For loans that disappear from the household’s portfolio during the year (e.g. because of refinancing), we assume that the date of termination coincides with the origination date for the household’s new loan. Cases where a loan disappears from the portfolio without being replaced by a new one are rare but do occur in our sample. In such cases, we use data on total interest payments on mortgage loans

from annual tax returns to infer when the loan was terminated.¹⁷

Once we have the full mapping of the household’s loan portfolio in each month, we use the detailed information about each loan to characterize this portfolio. First, we use information about the loan type to infer whether the household holds any adjustable-rate loans or interest-only loans.¹⁸ Second, we combine information about the loan’s current remaining balance, time to maturity, interest rate and amortization profile (interest-only vs. amortizing) to impute the average monthly payment over the full calendar year.

E Mass Layoffs

We obtain information about mass layoff events from the Ministry of Employment. All firms with more than 20 employees are obliged to report to the Ministry if they plan to lay off workers on a large scale. The exact definition of “large scale” depends on firm size, ranging from 10 workers for firms with 20-100 employees to 30 workers for firms with more than 300 employees. Our data contains information about the ID of the firm, the extent of the planned layoffs, and the date of reporting to the Ministry of Employment. Through the firm ID, we can link it to the payroll data from the Tax Agency and thus construct a subsample of individuals who have been laid off shortly after their employer reported a planned mass layoff. The report must be submitted before workers are given notice of their impending layoff. Since we do not observe when this happens for the individual worker, we include all cases where the date of reporting is within months -7 to -1 relative to the observed month of layoff. This leaves us with a subsample of 1,156 individuals.

Figure A12 shows event graphs for income and spending for the mass layoff subsample

¹⁷More precisely, we use the information about the loan’s interest rate and remaining balance to calculate how much interest would have been paid on the loan over the full year. We then compare that number to information from the tax return data on how much interest on mortgage loans the household actually paid. If the former number is twice as large as the latter, we conclude that the loan was terminated after the first six months of the year.

¹⁸A loan can change amortization profile during the year, i.e. from amortizing to interest-only, or vice versa. The end-of-year snapshots provide information about the date of the most recent such change, allowing us to infer the loan’s profile in any given month during the year. A loan cannot change from fixed to adjustable rate, or vice versa. Households need to prepay their existing loan and take out a new one if they wish to switch between these loan types.

vs. the full sample of active customers. The estimated development in salary income for the person affected by job loss – as well as the ensuing increase in government transfers – are nearly identical for two samples, suggesting that there is no significant difference in the size and persistence of the shock. Spending responses also look highly similar across the two samples, although the small number of observations in the mass layoff sample implies that confidence intervals are substantially wider. Combined, these results suggest that the total amount of self-insurance is about the same in the mass layoff subsample as in the full sample.

Table A3 shows results for cumulated effects over the full observation window for both samples and for all outcomes. The cumulative direct income loss is almost the same in the two samples. This suggests that the presence of voluntary resignations (e.g., individuals who deliberately take time off between jobs) in our baseline sample is no cause for concern, since the cumulative income loss would most likely be smaller in such cases. There are some differences when it comes to the relative importance of the behavioral responses to this income loss, especially for private transfers and other inflows where we find a negative but insignificant effect in the mass layoff sample. In general, the cumulative responses are imprecisely estimated in this sample due to the limited number of observations. But the point estimates suggest that our main findings are robust: First, household spending drops by 30-45% of the direct income loss, suggesting substantial self-insurance. Second, the compensating effects from spousal labor supply, borrowing and loan repayments are small and/or insignificant. Third, the most important self-insurance response is reduced saving in liquid assets, which accounts for 50-65% of the direct income loss.

F Tables and Figures

Table A1: Sample selection and summary statistics, extended

	(1)	(2)	(3)
	Gross sample	Active customers (baseline sample)	Exclusive customers
No. of individuals	66,844	10,002	5,224
	----- Sample means -----		
Female	0.43	0.47	0.48
Age	46.16	46.59	46.13
Couple	0.67	0.59	0.52
Capital region	0.33	0.44	0.42
Higher education	0.23	0.28	0.27
Primary sector	0.01	0.01	0.01
Manufacturing	0.19	0.15	0.15
Construction	0.07	0.06	0.06
Trade & transport	0.26	0.26	0.26
Other services	0.20	0.22	0.20
Public Sector	0.23	0.28	0.28
Arts & entertainment	0.03	0.04	0.03
Homeowner	0.65	0.63	0.59
Annual gross income for person who lost job (DKK)	371,621	394,499	375,019
Spouse employed	0.79	0.84	0.85
Share of hsh. bank deposits held at other banks	0.71	0.05	0.00
Share of hsh. retail bank loans held at other banks	0.71	0.11	0.00
Household inflows to Danske Bank accounts	11,974	40,033	36,339
- salary, affected person	5,527	19,450	18,827
- salary, spouse	2,836	9,885	8,510
- government income transfers	859	2,630	2,431
- private transfers and other inflows	2,752	8,068	6,571
Household spending from Danske Bank accounts	7,401	25,920	24,448
Household loan balances at Danske Bank	49,590	174,754	165,132
Household liquid assets at Danske Bank	62,717	201,778	192,459
Household mortgage payments, all mortgage banks	2,779	3,115	2,781
Household mortgage debt, all mortgage banks	696,257	747,244	654,545

The table is an extended version of Table 1 in the main text. Column (1) shows statistics for the gross sample with no requirements on customer status at Danske Bank. Column (2) shows statistics for the baseline sample of active customers, i.e., individuals who have at least five outgoing spending transactions in each month of the event observation window and whose partner (if any) satisfies the same criterion. Column (3) is for the sample of exclusive customers, i.e., active customers who have no deposits or loans at other retail banks and whose partner (if any) satisfies the same criterion. All variables are measured in month -6 relative to the month of job loss, except the following: Annual gross income, measured over the calendar year in which month -6 occurs; shares of household loans and deposits held at other banks, household mortgage debt at all mortgage banks, all measured at the end of the calendar year before month -6.

Table A2: The dynamic effects of job loss on income, saving and spending

	(1)	(2)	(3)	(4)	(5)	(6)
	Month 0	Month 6	Month 12	Month 24	Cumulative, months -5 to 24	Cumulative, months -5 to 24
<i>Direct income effects</i>	----- Relative to monthly disposable income before job loss -----					Percent of direct income loss
[1] Salary, affected person	-0.512 (0.004)	-0.305 (0.005)	-0.243 (0.006)	-0.210 (0.008)	-6.926 (0.142)	
[2] Government transfers, household	0.219 (0.003)	0.206 (0.003)	0.158 (0.003)	0.127 (0.005)	4.561 (0.078)	
[3] Direct income loss (-[1] - [2])	0.293 (0.004)	0.098 (0.005)	0.086 (0.006)	0.083 (0.008)	2.365 (0.139)	100.0% (0.0%)
<i>Response margins</i>						
[4] Salary, spouse	0.010 (0.002)	0.006 (0.002)	0.004 (0.003)	0.004 (0.004)	0.146 (0.069)	6.2% (2.9%)
[5] Private transfers and other inflows	0.036 (0.004)	0.014 (0.005)	0.010 (0.006)	0.009 (0.009)	0.286 (0.148)	12.1% (6.4%)
[6] Spending	-0.019 (0.004)	-0.030 (0.005)	-0.022 (0.006)	-0.028 (0.009)	-0.717 (0.154)	-30.3% (6.8%)
[7] Net saving in liquid assets	-0.195 (0.008)	-0.045 (0.009)	-0.055 (0.012)	-0.064 (0.017)	-1.163 (0.298)	-49.2% (12.4%)
[8] Non-mortgage loan net repayments	-0.032 (0.003)	-0.003 (0.003)	-0.001 (0.004)	-0.002 (0.006)	-0.048 (0.104)	-2.0% (4.5%)
[9] Mortgage loan repayments	-0.001 (0.000)	-0.002 (0.000)	-0.003 (0.001)	-0.003 (0.001)	-0.063 (0.013)	-2.7% (0.6%)
[10] Total ([4] + [5] - [6] - [7] - [8] - [9])	0.292 (0.007)	0.100 (0.009)	0.094 (0.011)	0.110 (0.016)	2.423 (0.274)	102.4% (10.6%)


All estimates are based on regression estimates from estimation of equation (1). The reported outcomes are measured relative to the household's average income in the pre-event months. Rows in normal font show coefficient estimates from single regressions with the indicated outcomes. Rows in bold font show combination of coefficients from multiple regressions, as indicated in parenthesis. Columns (1) to (4) report coefficients on the indicator variables representing months 0, 6, 12 and 24 after the unemployment event, respectively. Column (5) reports the sum of coefficient values for event months -5 to 24. Column (6) reports the ratio between the sum shown in the same row, column (5) and the corresponding sum shown in row [3], column (5). Standard errors (in parentheses) are estimated by bootstrapping with 500 replications. The bootstrapping procedure is carried out with resampling of individuals, rather than individual observations, to account for the panel nature of the data set.

Table A3: Cumulative effects of job loss, baseline vs. mass layoff sample

	Baseline		Mass layoffs	
	(1)	(2)	(3)	(4)
----- Cumulative effects, months -5 to 24 -----				
	Rel. to disp. inc. before job loss	Percent of direct income loss	Rel. to disp. inc. before job loss	Percent of direct income loss
<i>Direct income effects</i>				
[1] Salary, affected person	-6.93 (0.14)		-6.87 (0.42)	
[2] Gov. transfers, household	4.56 (0.08)		4.77 (0.24)	
[3] Direct income loss (-[1] - [2])	2.37 (0.139)	100.0% (0.0%)	2.10 (0.39)	100.0% (0.0%)
<i>Response margins</i>				
[4] Salary, spouse	0.15 (0.07)	6.2% (2.9%)	0.06 (0.20)	3.0% (11.4%)
[5] Private transfers and other inflows	0.29 (0.15)	12.1% (6.4%)	-0.34 (0.45)	-16.0% (27.1%)
[6] Spending	-0.72 (0.15)	-30.3% (6.8%)	-0.92 (0.45)	-43.9% (27.5%)
[7] Net saving in liquid assets	-1.16 (0.30)	-49.2% (12.4%)	-1.36 (0.93)	-64.9% (53.2%)
[8] Non-mortgage loan net repaym.	-0.05 (0.10)	-2.0% (4.5%)	-0.03 (0.30)	-1.5% (17.4%)
[9] Mortgage loan repayments	-0.06 (0.01)	-2.7% (0.6%)	-0.01 (0.04)	-0.7% (2.2%)
[10] Total ([4] + [5] - [6] - [7] - [8] - [9])	2.42 (0.27)	102.4% (10.6%)	2.06 (0.86)	98.1% (47.1%)
Number of individuals	10,002	10,002	1,156	1,156

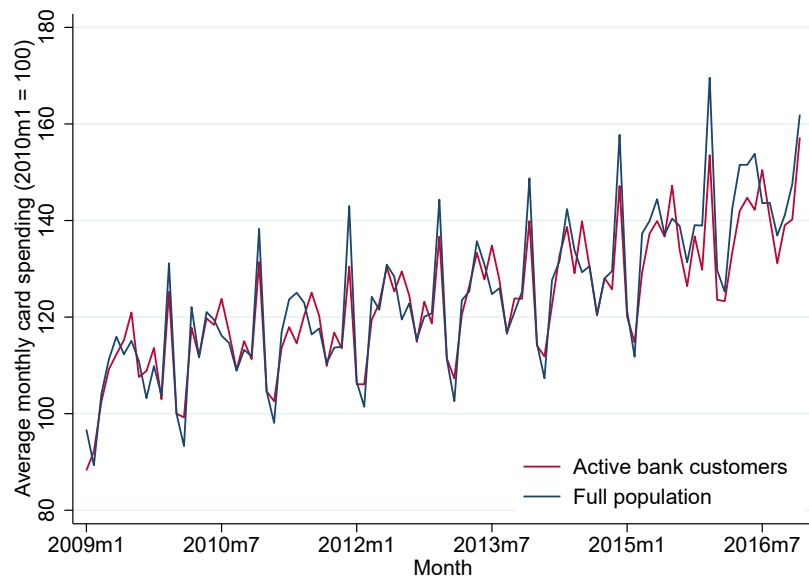
The table reports results from estimation of equation (1) for our baseline sample vs. a subsample of individuals who lost their job concurrently with a mass layoff at their employer. Columns (1)-(2) show baseline results, reproduced from columns (5)-(6) of Table A2. Columns (3)-(4) report parallel results for the mass layoff subsample. All estimates are based on regressions where the reported outcomes are measured relative to the household's average disposable income in the pre-event months. Odd-numbered columns report the sum of coefficients for event months -5 to 24 from such regressions. Even-numbered columns report the ratios between these sums and the corresponding sum for the direct income loss shown in row [3]. Standard errors (in parentheses) are estimated by bootstrapping with 500 replications. The bootstrapping procedure is carried out with resampling of individuals, rather than individual observations, to account for heteroskedasticity and autocorrelation within observations for the same individual.

Figure A1: Average spending for active customers, by payment method and quarter



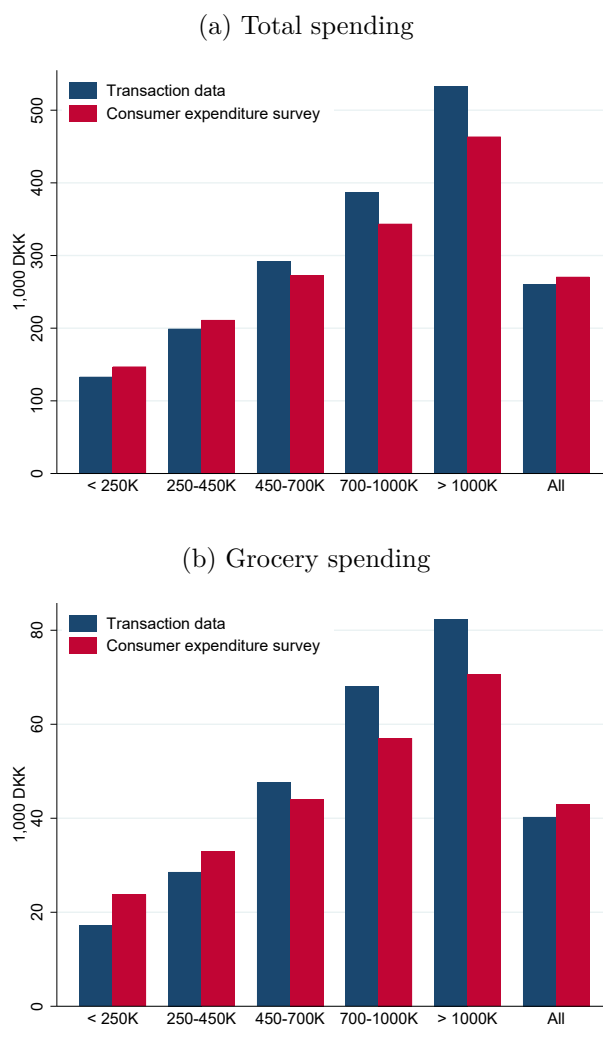
The figure shows the breakdown of the spending measure on categories of outflows for the complete sample of active customers. Card payments include payments via cellular phone.

Figure A2: Card spending per person, transaction data vs. aggregate data from Statistics Denmark



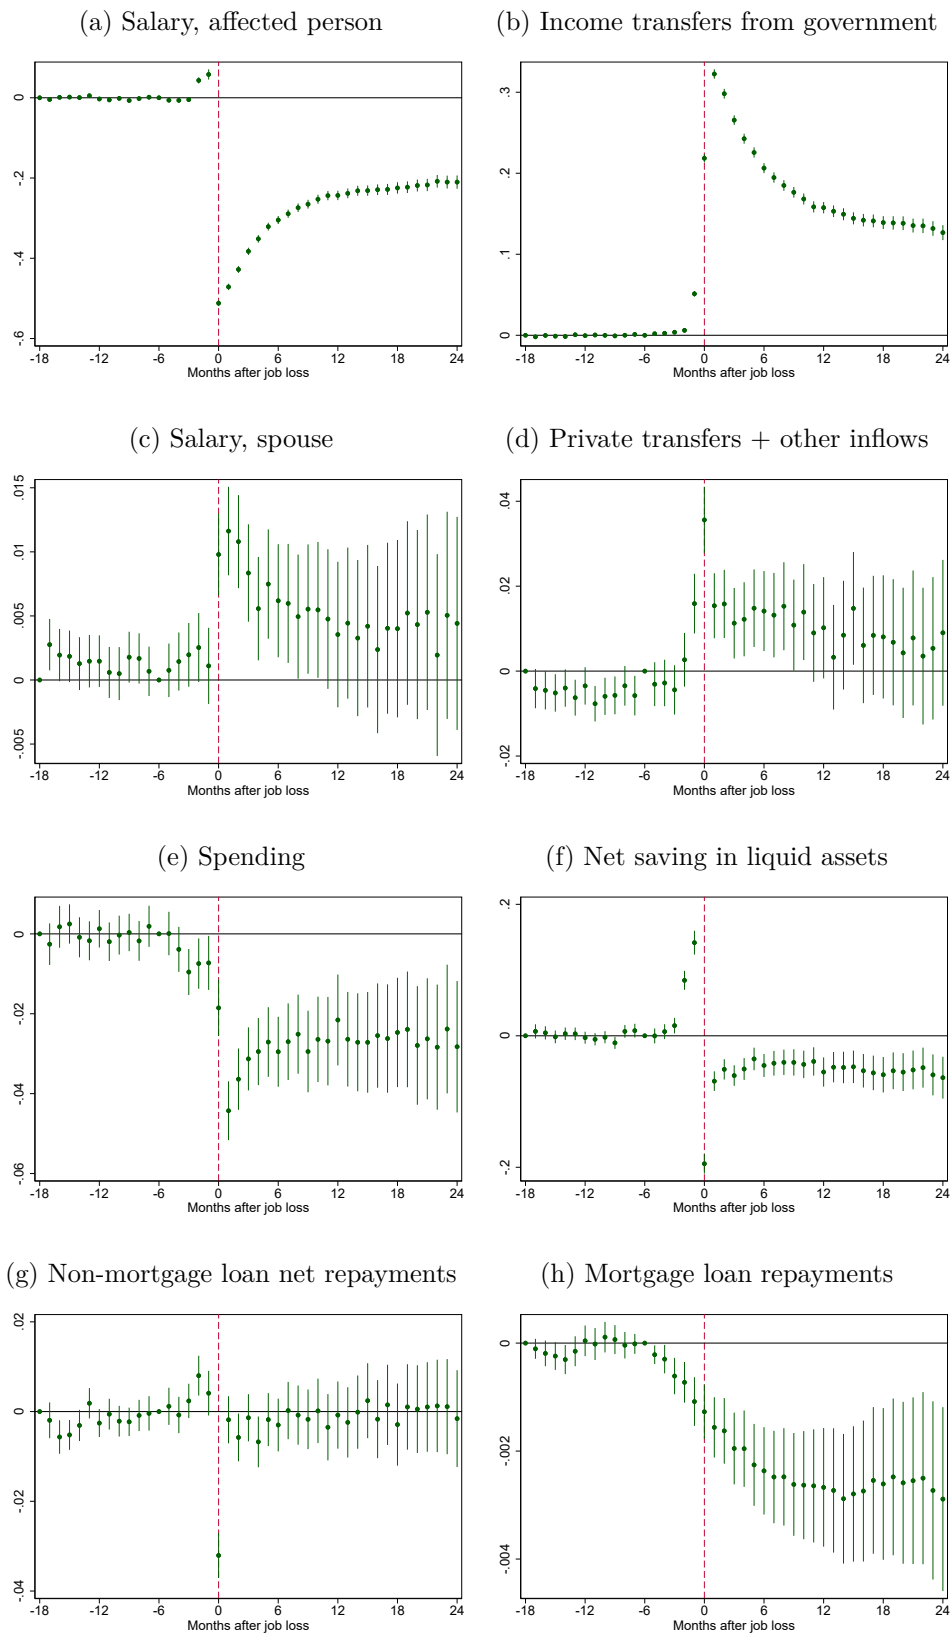
The figure compares the development in card spending per person for the complete sample of active bank customers vs. the full population. The red line shows average monthly card spending per person for the complete sample of active bank customers, based on transaction data. The blue line shows per capita card spending for the full population, calculated from official statistics published by Statistics Denmark. To construct the latter series, we have divided total aggregate card spending in each month by the number of persons in the population aged 18 or older.

Figure A3: Annual household spending by income group, 2009-16, transaction data vs. consumer expenditure survey



The figure shows averages of total annual household spending and grocery spending across income groups. Income groups are defined by total annual household income in DKK. Average spending levels are computed within each group and each year and then averaged across years. Blue columns are based on the bank transaction data for the complete sample of active customer households, i.e., households in which we observe at least five spending transactions in each month of the calendar year for each adult person. Red columns are based on the Statistics Denmark's consumer expenditure survey (Forbrugsundersøgelsen). For this data source, grocery spending is computed as the sum of spending on food, beverages, and tobacco.


Figure A4: Dynamic responses to job loss, individual outcomes




The figure shows estimation results from the event study model (1) of the effects of job loss on various outcomes. All outcomes are measured relative to the household's average disposable income between event months -18 and -3 and winsorized at the 2.5 and 97.5 percentiles within each event month. Disposable income is defined as all external inflows to the household's bank accounts. Vertical lines represent 95% confidence intervals. Standard errors are estimated allowing for clustering at the level of the individual.

Figure A5: Dynamic responses for subcomponents of spending


(a) Utilities



(b) Restaurants and bars

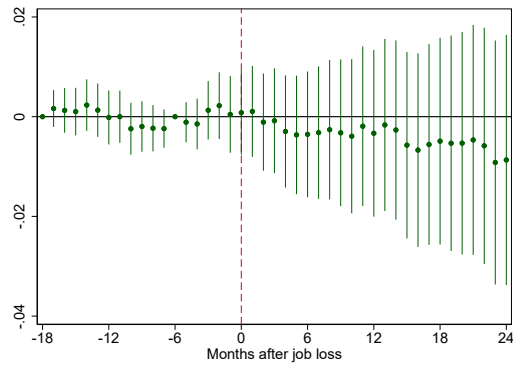


(c) Groceries



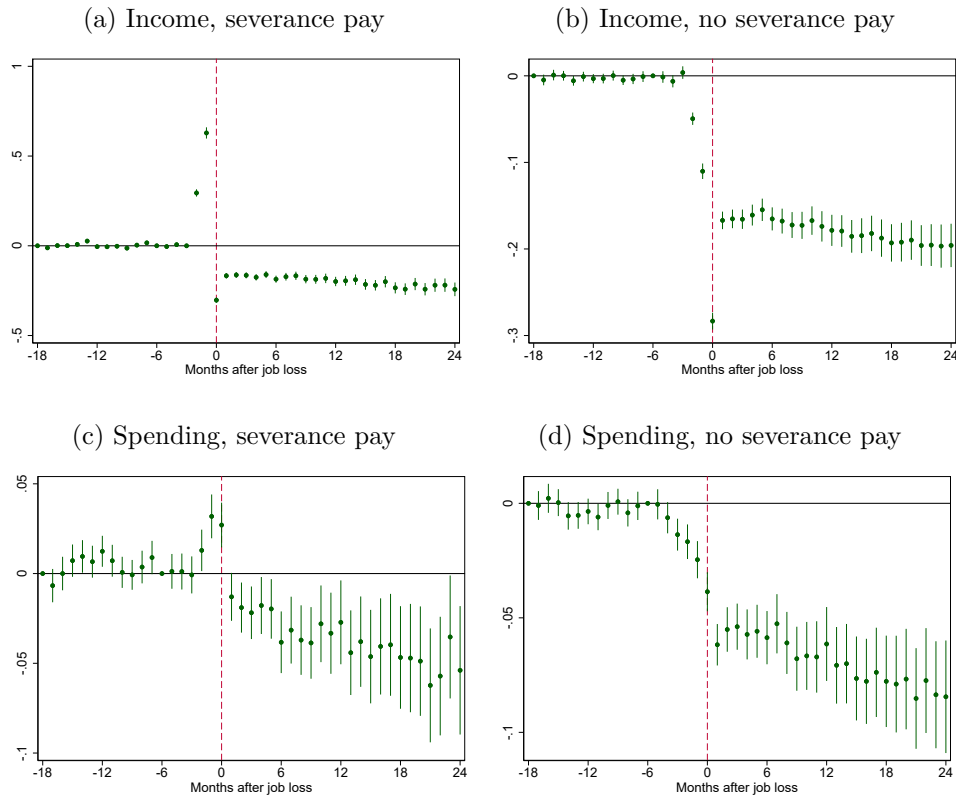
The figure shows estimation results from the event study model (1) of the effects of job loss on selected categories of household spending. Spending categories are defined by Merchant Category Codes, as described in Appendices A and B. All outcomes are measured relative to the household's average disposable income between event months -18 and -3 and winsorized at the 2.5 and 97.5 percentiles within each event month. Disposable income is defined as all external inflows to the household's bank accounts. Vertical lines represent 95% confidence intervals. Standard errors are estimated allowing for clustering at the level of the individual.

Figure A6: Spouse employment (couples only)



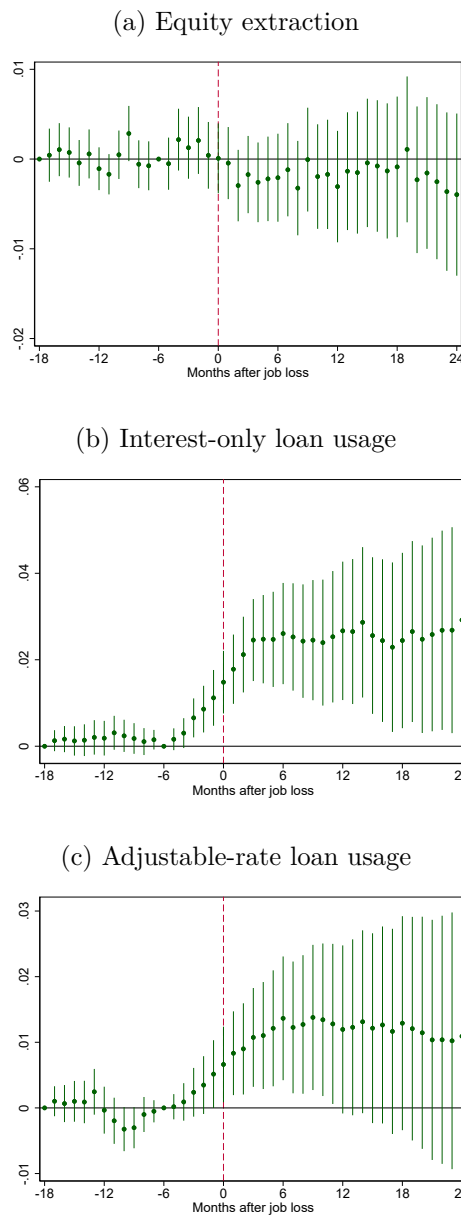
The figure shows estimation results from the event study model (1) of the effects of job loss on spouses' employment rates. The dependent variable is a dummy variable equal to 1 if the spouse of the person experiencing job loss appears on the payroll of at least one employer in the given month. Individuals with no spouse are excluded. Vertical lines represent 95% confidence intervals. Standard errors are estimated allowing for clustering at the level of the individual.

Figure A7: Income and spending, severance pay vs. no severance pay (survivor sample)




The figure shows estimation results from the event study model (1) of the effects of job loss on income (salary and government transfers) and spending. All outcomes are measured relative to the household's average disposable income between event months -18 and -3 and winsorized at the 2.5 and 97.5 percentiles within each event month. The sample is a dynamic sample of individuals that stay unemployed. In event months -18 to 0, this includes everyone in the baseline sample. For event month $t > 0$, it includes those who have not returned to employment at any point between month 0 and month t . Employment status is defined as having gross wage income above 10,000 DKK (at January 2010 price level). Panels a and c show results for the subsample of individuals whose income in month -1 or -2 is at least 50% above the average in the pre-event months. Panels b and d show corresponding results for the subsample who do not satisfy this criterion. Vertical lines represent 95% confidence intervals. Standard errors are estimated allowing for clustering at the level of the individual.

Figure A8: Mortgage loan responses (mortgagors only)




The figure shows estimation results from the event study model (1) of the effects of job loss on mortgage loan outcomes. Panel a shows results from a regression where the dependent variable is a dummy for equity extraction. The dummy takes the value 1 if the household replaced an existing mortgage loan with a new one with principal exceeding $(B+20,000)/0.95$, where B is the outstanding balance on the existing loan in DKK. This criterion takes into account that refinancing involves a fixed fee plus a rate loss that is proportional to the principal. Panels b and c show results from regressions where the dependent variables are dummies for whether the household has at least one interest-only loan and adjustable-rate loan, respectively. The sample consists of individuals from households with at least one mortgage loan. Vertical lines represent 95% confidence intervals. Standard errors are estimated allowing for clustering at the level of the individual.

Figure A9: Net saving in liquid assets, levels



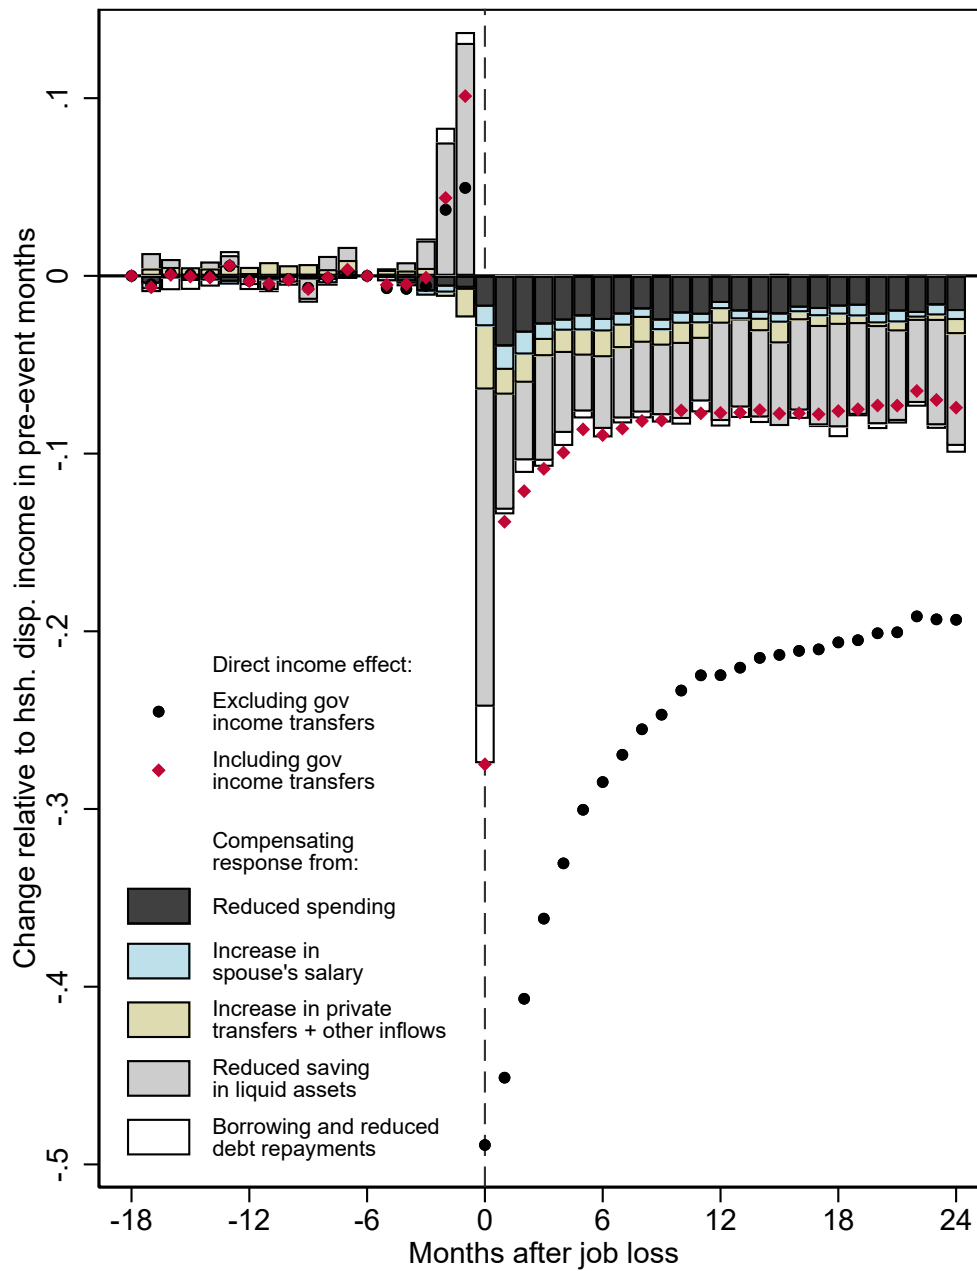
The figure shows average predicted values from the event study model (1) with net saving in liquid assets as the outcome. The dependent variable is measured relative to the household's average disposable income between event months -18 and -3 and winsorized at the 2.5 and 97.5 percentiles within each event month. Each estimate is the average predicted value when the event time variable takes the value indicated on the horizontal axis and control variables are evaluated at their actual values. Vertical lines represent 95% confidence intervals. Standard errors are estimated allowing for clustering at the level of the individual.

Figure A10: Income, spending and self-insurance responses to job loss, exclusive bank customers



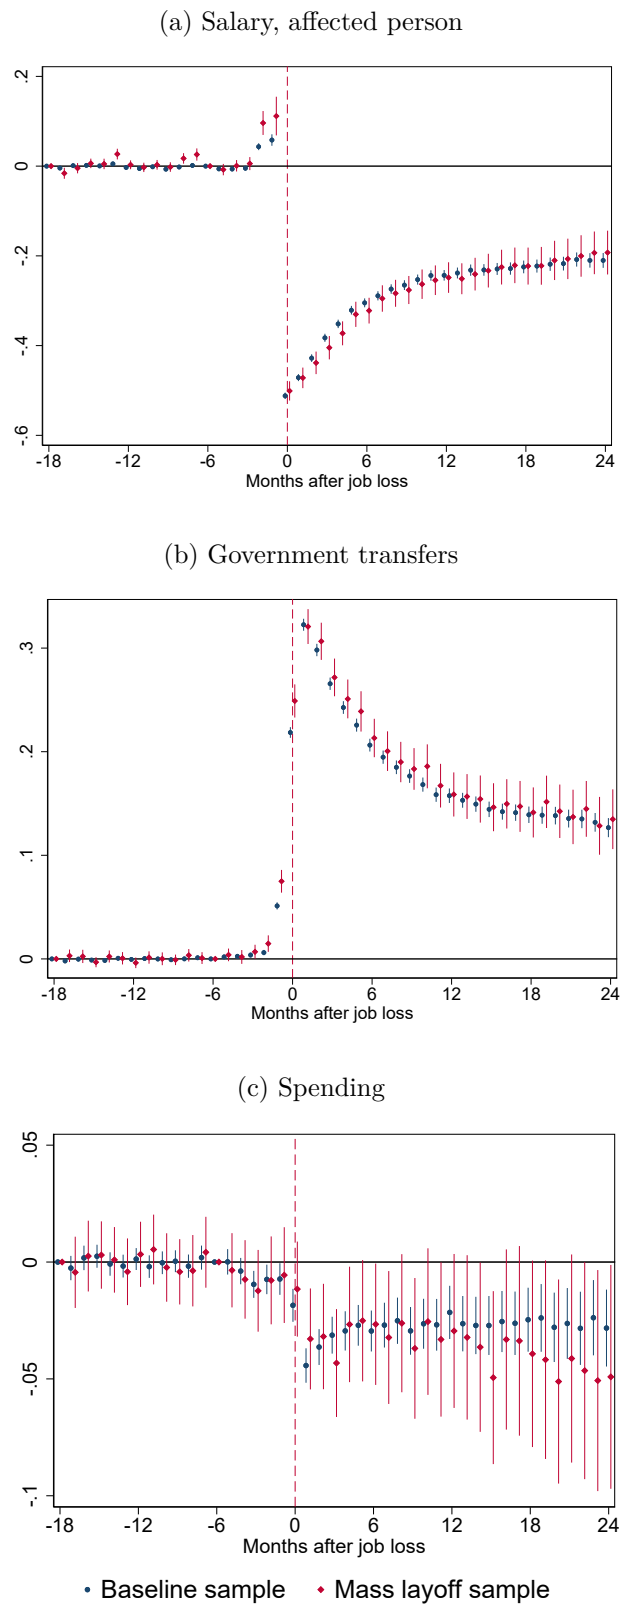
The figure shows estimation results from the event study model (1) of the effects of job loss on a range of outcomes, estimated on the subsample who belong to households in which all adult members are exclusive customers of the bank. See notes to Figure 1 in the main text for further details.

Figure A11: Income, spending and self-insurance responses to job loss, weighted regressions



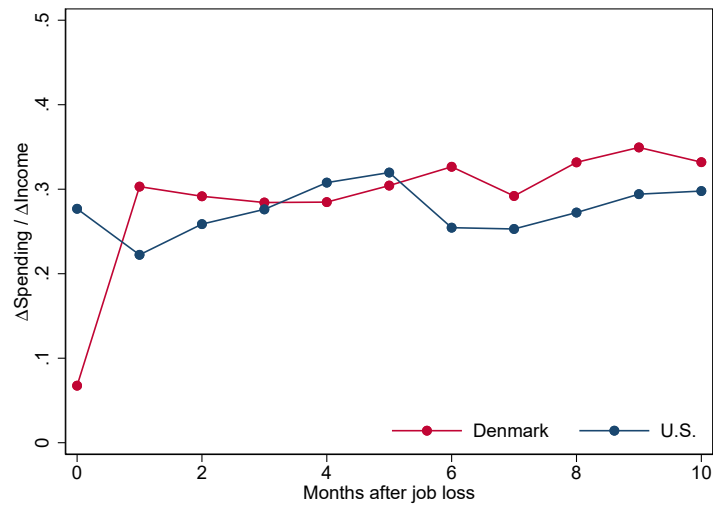
The figure shows estimation results from the event study model (1) where observations are weighted so that the sample of active customers matches the characteristics of the gross sample shown in column (1) of Table 1 in the main text. The weights are the inverse predicted probabilities from a probit of active customer status on the demographic characteristics reported in that table. See notes to Figure 1 in the main text for further details.

Figure A12: Impact of job loss on income and spending: Baseline sample vs. mass layoff sample



The figure shows estimation results from the event study model (1) of the effects of job loss on income and spending. Blue markers show estimates for the baseline sample of active customers. Red markers show estimates for the subsample of individuals who were laid off concurrently with a mass layoff at their employer. All outcomes are measured relative to the household's average disposable income between event months -18 and -3 and winsorized at the 2.5 and 97.5 percentiles within each event month. Vertical lines represent 95% confidence intervals. Standard errors are estimated allowing for clustering at the level of the individual.

Figure A13: Spending response relative to direct effect on income conditional on staying unemployed, Denmark vs. U.S.



The figure shows estimates of the ratio between household spending and income responses to unemployment. For Denmark, we construct the series for the “survivor sample” of individuals who stay unemployed until the indicated month by dividing the estimate for the effect on spending in that month by the estimate for the effect on household income from salary and government income transfers. For the U.S., we back out estimates based on results reported in Ganong and Noel (2019): Each point estimate is calculated as the relative effect on spending divided by the relative effect on income, multiplied by the ratio of mean pre-onset spending to mean pre-onset income.