

How Does Unemployment Affect Consumer Spending?

Peter Ganong and Pascal Noel*

November 4, 2015

Abstract

We study the spending of unemployed individuals using anonymized data on 235,000 checking accounts that received a direct deposit of unemployment insurance (UI) benefits. The account holders are similar to a representative sample of U.S. UI recipients in terms of income, spending, assets, and age.

Unemployment causes a large but short-lived drop in income, generating a need for liquidity. At onset of unemployment, monthly spending drops by 6%, and work-related expenses explain one-quarter of the drop. Spending declines by less than 1% with each additional month of UI receipt. When UI benefits are exhausted, spending falls sharply by 12%.

Unemployment is a good setting to test alternative models of consumption because the change in income is large. We find that families do little self-insurance before or during unemployment, in the sense that spending is very sensitive to monthly income. We compare the spending data to three benchmark models; the drop in spending from UI onset through exhaustion fits the buffer stock model well, but spending falls much more than predicted by the permanent income model and much less than the hand-to-mouth model. We identify two failures of the buffer stock model relative to the data – it predicts higher assets at onset, and it predicts that spending will evolve smoothly around the largely predictable income drop at benefit exhaustion.

Keywords: Unemployment, Spending, Liquidity Constraints, Buffer Stock, Permanent Income Hypothesis

JEL Codes: E21, E24, J65

*We thank our advisors Larry Katz, David Laibson, and Jeff Liebman. We also thank John Beshears, John Campbell, Raj Chetty, Gabe Chodorow-Reich, David Cutler, Ed Glaeser, Nathan Hendren, Simon Jäger, Rohan Kekre, Annie Levenson, Mandy Pallais, Mikkel Plagborg-Møller, Martin Rotemberg, David Scharfstein, Andrei Shleifer, and Dan Shoag for helpful comments. Thanks to Nathan Palmer for sharing simulation code. Technical support was provided by Ista Zahn of the Research Technology Consulting team, at the Institute for Quantitative Social Science, Harvard University. Thanks to Ed Dullaghan, Wayne Vroman and Scott Schuh for sharing institutional knowledge related to the UI system and the payments system. This research was made possible by a data-use agreement between the authors and the JPMorgan Chase Institute (JPMCI), which has created anonymized data assets that are selectively available to be used for academic research. All statistics from JPMCI data, including medians, reflect cells with at least 10 observations. The opinions expressed are those of the authors alone and do not represent the views of JPMorgan Chase & Co. While working on this paper, Ganong and Noel were paid contractors of JPMCI. We gratefully acknowledge funding from the Washington Center for Equitable Growth, the Alfred P. Sloan Foundation Grant No. G-2011-6-22, “Pre-Doctoral Fellowship Program on the Economics of an Aging Workforce,” awarded to the National Bureau of Economic Research, and the National Institute on Aging Grant No. T32-AG000186 awarded to the National Bureau of Economic Research.

1 Introduction

Many Americans have little liquid assets, limited access to credit, and immediately spend a substantial fraction of tax rebates, suggesting that financial constraints would necessitate substantial spending reductions during unemployment.¹ However, some mainstream economic models assume that individuals are able to smooth short-term income fluctuations.² We analyze anonymized bank account data on the spending of families receiving unemployment insurance (UI) benefits to test between these competing views.

Bank account data offer a rich view of the financial lives of families who receive UI. We analyze anonymized data on monthly checking account inflows and outflows assembled by the JPMorgan Chase Institute (JPMCI) and identify UI receipt through direct deposit of benefits. We build a dataset with two key advantages for studying spending during unemployment relative to surveys used in prior work.³ First, monthly bank account data enables us to trace out high-frequency drops and rebounds in spending at unemployment onset, re-employment and UI benefit exhaustion. Second, we can estimate the role of work-related expenses and how much spending drops on necessities.

Recipients of UI benefits tend to be middle-class families and the JPMCI sample looks similar to external benchmarks. Most states require UI claimants to have earnings in four of the five quarters prior to separation, meaning that low-income workers are often ineligible for benefits. Summary statistics on account holders in the JPMCI data are similar to external benchmarks for total family income, spending, debt payments, checking account balances and age.⁴

The first half of our paper describes the economic lives of families receiving UI. We

¹Evidence for this view includes Parker et al. (2013), Shapiro and Slemrod (2009) and Angeletos et al. (2001).

²Shimer and Werning (2008) model optimal unemployment insurance under an assumption of perfect access to liquidity. Blundell et al. (2008) find in a model calibrated to annual US data that there is complete insurance of transitory shocks, except among families with permanently low income.

³Examples include Cochrane (1991), Gruber (1997), Browning and Crossley (2001), and Stephens (2001).

⁴For each comparison, we choose the sample in the JPMCI data that best matches an easily-accessible external benchmark. We compare the family income and age of UI recipients in the JPMCI data to UI recipients in the SIPP. We compare spending and debt payments of all JPMCI families to all families the Consumer Expenditure Survey and the Survey of Consumer Finances (SCF). We compare checking account balances of employed families in the JPMCI data to employed families in the SCF.

divide our empirical analysis into three sections: (1) the onset of UI, (2) spending for those re-employed while receiving UI and (3) spending for those who exhaust UI benefits.

Spending drops sharply at the onset of unemployment, and this drop is better explained by liquidity constraints than by a drop in permanent income or a drop in work-related expenses. We find that spending on nondurable goods and services drops by \$160 (6%) over the course of two months.⁵ Consistent with liquidity constraints, we show that states with lower UI benefits have a larger drop in spending at onset. It is unlikely that permanent income can explain the drop at onset because the average lifetime income loss for UI recipients in the JPMCI data is only 14% of one year’s income.⁶ Finally, we define work-related expenses as those spending categories which decline at retirement for a sample of retirees with substantial liquid assets. Our definition, which includes food away from home and transportation, closely mirrors prior work by Aguiar and Hurst (2013). Work-related expenses drop more than other expenditure categories at onset. We estimate that the excess drop in this category explains about one-quarter of the total drop in spending at onset.

For UI recipients who are able to find work prior to exhaustion, spending remains depressed after re-employment as they rebuild their financial buffer. Prior work studying short-term unemployment using annual spending data assumed that spending recovered fully upon re-employment (Chodorow-Reich and Karabarbounis 2015, Kolsrud et al. 2015). In fact, someone who is unemployed for three months has 6% lower spending during unemployment *and* 3% lower spending (relative to onset) after re-employment. Decreased spending after re-employment was misinterpreted as a drop during unemployment, leading researchers to overstate the drop in spending for the short-term unemployed by as much as factor of three. We provide new estimates of the spending drop for researchers calibrating optimal UI benefits in a Baily (1978)-Chetty (2006) framework and studying how unemployment affects output over the business cycle (e.g. Kaplan and Menzio (2015)).

⁵This drop in spending occurs both absolutely and relative to a control group of families with annual income between \$30,000 and \$80,000. All income and spending estimates in this paper are reported relative to this control group.

⁶Both in the JPMCI data, and in a representative sample using the SIPP, we find that the average UI recipient experiences a quick recovery in her labor income. Although a prior literature started by Jacobson et al. (1993b) which studied income paths of high-tenure workers separated in mass layoffs found large permanent income losses, there is less research about the experience of a typical UI recipient.

Comparing spending of high- and low-asset families after re-employment provides further evidence for the central role of liquidity in explaining spending behavior. Some consumption models predict that families target a specific ratio of wealth to permanent income (Carroll 1997). To smooth an income shock of a fixed size, low-asset families need to draw down a larger fraction of their assets. We find empirically that spending remains depressed after re-employment for these low-asset families as they rebuild their buffers, consistent with the prediction of target ratio models.

As UI benefit exhaustion approaches, families who remain unemployed barely cut spending, but then cut spending by 12% in the month after benefits are exhausted. Benefit exhaustion offers a particularly powerful research design for studying excess sensitivity of spending to income because the drop in income is predictable, it contains little news about a job-seeker’s future income prospects and does not change the opportunity for home production. When benefits are exhausted, the average family loses about \$1,000 of monthly income.⁷ In the same month, spending drops by \$260 (12%). Grocery spending drops from \$287 per month during UI receipt to \$250 per month immediately after exhaustion. Although we do not have data on what types of foods people buy, analysis of food diaries by Aguiar and Hurst (2005) suggests that there is a substantial change in food quality.

We take these empirical facts – the large spending drop at onset, the slow decline during UI receipt, and the even larger spending drop at exhaustion and compare them to predictions from three benchmark models of consumption: a permanent income consumer, a buffer stock consumer and a hand-to-mouth consumer.

Unemployment is a particularly good setting for testing alternative models of consumption because it causes such a large change in family income. A literature starting with Akerlof and Yellen (1985), Mankiw (1985) and Cochrane (1989) has argued that because ignoring small price changes has a second-order impact on utility, a rule of thumb such as setting spending changes equal to income changes may be “near-rational.” More recently, many researchers have documented evidence of an immediate increase in spending in re-

⁷Family income drops by less than the amount of lost benefits because some UI recipients find work at the time of benefit exhaustion.

sponse to tax rebates and similar one-time payments.⁸ Some authors have interpreted this as evidence of widespread liquidity constraints. Fuchs-Schuendeln and Hassan (2015) argue that the high sensitivity of income to tax rebates is not sufficient to reject the permanent income hypothesis, because the welfare cost is small of adopting rule of thumb behavior for tax rebates. They calculate that in 18 studies using micro evidence the cost of rule of thumb behavior is 5% or less of annual consumption. For someone who is unemployed and exhausts UI benefits, the comparable statistic is 20%. Because the stakes are higher for unemployment, near-rationality is less of a concern and the path of spending offers a more convincing test of alternative models of consumption.

We compare the path of spending during unemployment in the data to three benchmark models and find that the buffer stock model fits better than a permanent income model or a hand-to-mouth model. We calibrate a model of consumption and savings in the tradition of Deaton (1991), Aiyagari (1994), and Carroll (1997). In our model, the only income risk comes from unemployment. UI benefits expire after six months. The decline in spending from onset through exhaustion in the data is equal to the decline predicted by the buffer stock model when agents hold assets equal to 0.84 months of income at the start of unemployment.

However, the buffer stock model has two major failures – it predicts substantially more asset holdings at onset and it predicts that spending should be much smoother at benefit exhaustion. First, a key prediction of buffer stock models is that agents accumulate precautionary savings to self-insure against income risk. Our model, where unemployment is the only risk, predicts that agents should hold three times as much assets as they do in the data.⁹ Second, models of forward-looking agents with exponential time preferences predict that spending should evolve smoothly in the face of predictable income changes. Even an agent who has zero assets at onset will avoid spending all of her UI benefits in order to

⁸Examples of work estimating excess sensitivity using one-time payments include Souleles (1999), Hsieh (2003), Johnson et al. (2006), Parker et al. (2013), Baugh et al. (2013), and Kueng (2015). Baker and Yannelis (2015) and Gelman et al. (2015) examine a temporary loss of labor income due to the federal government shutdown.

⁹Models with realistic income processes predict asset holdings which are an order of magnitude larger (Gourinchas and Parker (2002), Laibson et al. (2015)).

smooth the income drop at exhaustion. Two channels which could contribute to this sudden drop at exhaustion are over-optimistic beliefs about UI duration which update suddenly at exhaustion (Spinnewijn 2015) and inattention (Reis 2006, Cochrane 1989, Kueng 2015).

To summarize, we find that families do relatively little self-insurance when unemployed as spending is quite sensitive to current monthly income. We built a new dataset to study the spending of unemployed families using anonymized bank account records from JPMCI. Using rich category-level expenditure data, we find that work-related expenses explain only a modest portion of the spending drop during unemployment. The overall path of spending for a seven-month unemployment spell is consistent with a buffer stock model where agents hold assets equal to less than one month of income at the onset of unemployment. Because unemployment is such a large shock to income, our finding that spending is highly sensitive to income overcomes the near-rationality critique applied to prior work. Finally, we document a puzzling drop in spending of 12% in the month UI benefits exhaust, suggesting that families do not prepare for benefit exhaustion.

The paper proceeds as follows. Section 2 describes the JPMCI data set and why it is suited for measuring how unemployment affects spending. Section 3 quantifies the drop in spending at the onset of unemployment and argues that liquidity constraints are a better explanation than permanent income loss or work-related expenses. Section 4 shows that families rebuild their liquid assets after re-employment, consistent with a target ratio. Section 5 shows that income and spending drop sharply at benefit exhaustion. Section 6 compares predictions from different consumption models to the data. Section 7 concludes.

2 Data and External Validity

We construct a dataset suitable for studying unemployment and spending using JPMCI data from October 2012 to May 2015.¹⁰ We rely primarily on transaction-level checking

¹⁰Following Aguiar and Hurst (2005), we use the word “spending” to describe a specific subset of checking account outflows in the JPMCI data. We reserve the word “consumption” for discussing models such as the permanent income consumer, the hand-to-mouth consumer, and the buffer stock consumer. In the context of evaluating these models, we assume that the spending in the data actually reflects monthly consumption. Also, for consistency with the prior literature, we use the letter c in equations to describe the spending variable.

account inflows, checking account outflows, and debit card spending, which have been categorized and aggregated to the monthly level. We also use four additional datasets from JPMCI: spending on Chase credit cards, credit bureau records for Chase credit card customers, estimates of annual income, and estimates of total liquid asset holdings.

Administrative spending data have four advantages over the survey datasets used to study spending during unemployment in prior work: comprehensiveness, sample size, detailed spending categories and monthly frequency.¹¹ Many researchers have used the Panel Study of Income Dynamics (PSID), but it suffers from an ambiguous reference period and until recently it only covered food expenditures.¹² Changes in food expenditures are difficult to interpret around unemployment because of transitions to home production and because food may be a necessity good (Shimer and Werning 2007). Another data source is surveys which ask unemployed people how much they have cut spending since their job separation (Browning and Crossley 2001, Hurd and Rohwedder (2010)). Relative to this prior work, the JPMCI data cover all types of spending and a sample size of 235,000 UI recipients enables us to study subsamples such as benefit exhaustees or low-asset families re-employed after three months. With debit and credit card expenditure categories, we can estimate the role of work-related expenses and understand whether someone is cutting necessity goods when unemployed. Finally, we use the monthly frequency of the data to test predictions from different consumption models about how spending should change at re-employment and at UI exhaustion.

Families in the JPMCI dataset look similar to external benchmarks on family income, spending in certain categories, liquid assets and age.¹³ This representativeness is a strength

¹¹One exception to the widespread use of survey data to study spending during unemployment is recent work by Kolsrud et al. (2015) which uses annual administrative data on income and asset holdings from Sweden to infer spending. The Kolsrud et al. (2015) data are superior to the JPMCI data in that they capture asset holdings across all banks, while the JPMCI data have the advantages of a monthly frequency and detailed expenditure categories. One example of a survey with some data on income and spending at a monthly frequency is Hannagan and Morduch (2015).

¹²Examples of papers studying the impact of unemployment on food expenditure in the PSID include Cochrane (1991), Gruber (1997), Stephens (2001), Chetty and Szeidl (2007), Saporta-Eksten (2014), Chodorow-Reich and Karabarbounis (2015) and Hendren (2015). The PSID asks about “usual” weekly expenditure on food at home and then about food away from home without prompting a frequency. Most analysts have interpreted this as referring to the prior year’s expenditure (Blundell et al. (2008), Chodorow-Reich and Karabarbounis (2015)).

¹³Following Baker (2014), we assume the sample unit to be analogous to a “consumer unit” in the Consumer

of the JPMCI spending data in comparison with spending data from personal finance websites.¹⁴ For example, Kueng (2015) reports that median after-tax family income in Alaska in the personal finance website dataset was about 50% higher than for a representative sample. However, these personal finance websites have strengths relative to the JPMCI dataset, such as better coverage of asset holdings and families with multiple checking accounts. Another strength of the JPMCI data is the availability of anonymized linked credit bureau records, where we can see all outstanding debts and delinquencies.

2.1 Finding UI Recipients and Building A Full View of Family Finances

We look at checking account transaction descriptions to determine whether someone received UI by direct deposit. Particular text descriptions are associated with electronic transfers from state UI agencies.¹⁵ The population-weighted average of state-level direct deposit adoption rates was 45% in 2012 (Saunders and McLaughlin (2013)).¹⁶

One challenge for measuring a family’s spending is that some families have multiple checking accounts and we take three steps to achieve the best possible coverage of a family’s spending. The McKinsey Consumer Financial Life Survey showed that 39% of banked families had multiple accounts in 2013 (Welander 2014). Of these families, 39% had an additional account at the their primary bank and 71% had an additional account at another bank. First, to address this concern, we study all of the checking accounts which each family has linked together (see Appendix A.1 for details). Second, we focus on families who use Chase as their primary bank. Most people “home” on a single credit or debit card for point-of-sale payments (Cohen and Rysman 2013, Shy 2013). Given that changing cards is

Expenditure Survey, a family in the Survey of Income and Program Participation, and a “primary economic unit” in the Survey of Consumer Finances. We refer to the sampling unit as a “family”, even though the family may have only one member.

¹⁴Recent work using data from these websites includes Baker and Yannelis (2015), Gelman et al. (2015), Baugh et al. (2013), Kuchler (2014), and Kueng (2015). From a representativeness perspective, the best data source is administrative datasets on income and asset holding which cover all citizens like those used by Kolsrud et al. (2015) and Kostøl and Mogstad (2015).

¹⁵Altogether, we found transaction descriptions associated with 32 states. The bank has branches in 21 of these states.

¹⁶In addition, as evidence for external validity, we estimate that the share of US families receiving UI via direct deposit is close to the share of families in the data. Across the US, an average of 2.9 million people received UI benefits each week in 2014. We estimate that in an average week in 2014, 1.0% of families in the US received UI benefits via direct deposit. In the bank data, the average monthly UI reciprocity rate in 2014 was 0.8%.

easier than changing checking accounts, we believe that the same “homing” behavior exists for checking accounts and study accounts with at least five monthly outflows.¹⁷ Finally, sometimes two customers will form a family unit without linking their accounts. In our robustness checks, we study *unlinked* checking accounts which appear to reflect the same family.

2.2 Estimating Income and Comparison to External Benchmarks

To construct economically meaningful measures of income, JPMCI has applied extensive logic to categorize checking account inflows into twenty-two groups. We organize inflows into four major groups: payroll paid using direct deposit (61% of inflows three months prior to onset of UI), government income (4%), transfers from outside savings and investment accounts (“dissaving”, 10%) and other income (4%).¹⁸ Together, these categories cover 79% of total inflows and we place the remainder of inflows – which are largely made up of paper checks – into a residual category (21%).

Subjects in the JPMCI dataset who receive direct deposit of their UI benefits have similar incomes to a representative sample of UI recipients, suggesting that our analysis will have external validity for all UI recipients. In the SIPP, we construct the distribution of family income in the 12 months prior to UI receipt. In the JPMCI data, we use checking account inflows (except dissaving), rescaled into pre-tax dollars. Median family income is \$61,000 using the SIPP and \$54,000 using checking account income. Figure 1 shows that the income distribution of families receiving UI in the JPMCI data is broadly similar to the distribution for families receiving UI in the SIPP.¹⁹

¹⁷How many payments make a checking account primary? We do not have the data to answer this question directly, but 94% of people with one checking account have at least five outflows per month according to the Survey of Consumer Payment Choice.

¹⁸Appendix A.2 provides additional detail on the types of inflows observed in the JPMCI data. Appendix Table 1 shows additional summary statistics for each category. Measurement error is widespread and we winsorize all inflow variables at the 95th percentile.

¹⁹The share of direct deposit labor income in inflows (67%) is a bit lower than our estimated external benchmarks (78%). We calculate the external benchmark estimate by multiplying labor income as a share of family income (91% prior to UI receipt in the SIPP) times fraction of payroll dollars distributed by direct deposit (86% in the SCF). Because some paper checks likely reflect transfers between different accounts, the true ratio of direct deposit labor income to total income likely exceeds 67%. Table 1 provides additional

Although bank account data may not provide a good window into spending for everyone, such data provide good coverage for UI recipients, because they tend to be in middle-class families. To be eligible for UI benefits, a claimant needs substantial work history in the prior year. Table 1 shows the impact of this requirement quantitatively using the SIPP. In the twelve months prior to unemployment, UI recipients had median monthly pre-tax family income of about \$5,100 and a poverty rate of only 8%. While UI recipients are poorer than all employed people, they are higher-income and older than the general pool of unemployed people. Finally, we find using the Survey of Consumer Finances (SCF) that only 5% of employed families lack a bank account, suggesting that the vast majority of UI recipients have a bank account.

2.3 Estimating Spending and Comparison to External Benchmarks

Much as with income, checking account outflows can be hard to interpret and JPMCI has categorized them into thirty different groups. We organize outflows under four broad headings: spending on goods and services consumed immediately (54% of outflows), consumer debt payments (17%), unclassifiable payments (23%) and saving (6%).²⁰

Most of our analysis in this paper focuses on spending on goods and services consumed immediately. Our definition of spending has three components: (1) debit and credit card spending (\$1472 monthly, 34% of total outflows), (2) cash withdrawals (\$590, 14%) and (3) bill payments (\$309, 7%). Note that this definition includes spending on Chase credit cards at the time goods are purchased, rather than when the credit card bill is paid, which may be months later.²¹ In addition, all credit and debit card transactions include a Merchant Category Code that enables us to test whether specific expenditure categories change in the way predicted by theories of home production (Aguiar and Hurst (2013)).

statistics on the income of UI recipients in the SIPP, with comparisons to the JPMCI data. See Rothstein and Valetta (2014) for additional details on income of UI recipients.

²⁰Appendix A.2 provides additional detail on the content of each of these categories. We winsorize all inflow variables at the 95th percentile.

²¹Mean monthly Chase credit card spend is \$208. Because our sample screen requires five outflows in *every* month, our sample is skewed toward frequent debit card users and away from frequent credit card users.

When we compare the JPMCI spending data to external benchmarks, we find under-coverage of total consumption using a “top-down” approach while we find better coverage of eight clearly-identified expenditure categories using a “bottom-up” approach. First, for the “top-down” approach, we focus on nondurable goods and services in the Consumer Expenditure Survey (CEX) and in the Bureau of Economic Analysis’ Personal Consumption Expenditures (PCE).²² We estimate that our spending measure is 94% of the CEX benchmark and 44% of the PCE benchmark. We believe that our true coverage of spending for UI recipients is somewhere between these two numbers: the CEX is too low because of under-reporting and PCE is too high because it includes the consumption of very wealthy people who are not relevant for our study. Second, using a “bottom-up” approach, we compare spending on food away from home, food at home, fuel and utilities in Table 2. Estimated spending by JPMCI customers is 96-131% of the CEX benchmark and 53-77% of the PCE benchmark. We similarly compare spending on mortgages, auto loans, credit card payments, and student loans; conditional on making a payment, mean outflows are 56-77% of what is reported in the SCF by families making the same payments.

2.4 External Validity – Geography, Age, and Checking Account Balances

The JPMCI sample also looks broadly representative of US families in terms of geography, checking account balances and age, lending additional support to our argument for external validity. Chase has physical branches in 23 states, including the five most populous states in the US: California, Texas, Florida, New York and Illinois. We compare the age distribution of UI recipients in the JPMCI data to the SIPP in Appendix Figure 1 and find that these two distributions are closely aligned. Because we only observe the age of the primary account

²²We exclude healthcare and pensions because employers often pay for these services directly. We exclude housing because we are unable to measure rent, which is typically paid using paper checks, and we exclude utilities because PCE combines housing and utility costs into a single category. For 2013, CEX estimated total mean monthly spending of \$4,258 and PCE estimated \$7,615. It is well known that CEX understates consumption expenditures. Passero et al. (2011) carefully crosswalk CEX and PCE expenditure categories and found the ratio of CEX to PCE was 0.60 across all categories and 0.77 across comparable categories. To ensure comparability with these external data sources, the statistics from the JPMCI data reported in this section are for all accounts with 5 monthly outflows, rather than just for UI recipients.

holder in the JPMCI data, we compare it to the age of the family head in the SIPP. UI recipients in the JPMCI data (mean age: 41.1) are slightly younger than UI recipients in the SIPP (mean age: 44.3).

To understand how representative the JPMCI sample is in terms of assets, we compared it to the SCF. We compared balances for employed families in the data to balances for employed families in “the checking account you use the most” in the SCF. Figure 1 shows that the distribution of balances is similar between the two samples. Table 3 reports summary statistics comparing the JPMCI and SCF samples. In the SCF, the median total liquid assets for an employed family is \$4,900 and the median balance in a family’s primary checking account is \$1,500. The difference in medians highlights a limitation of checking account data, which is that most liquid assets are held *outside* a family’s checking account. The median checking account balance in the data is \$1,460, suggesting that on this dimension, families in the JPMCI data are similar to a cross-section of US families. In the data, we see substantial inflows from outside accounts during unemployment so even though we are unable to measure total asset holdings reliably, we can measure the extent to which families draw down their assets or draw on funds from informal insurance networks during unemployment.

2.5 Comparison Groups

To eliminate seasonality, inflation, secular trends, and business cycle fluctuations, all results for income and spending are presented relative to a comparison group. In the JPMCI data, there is an upward secular trend in spending of five percent per year and in labor income of six percent per year. This increase is larger than can be explained by economic fundamentals during this period. We believe that this trend reflects secular growth in the use of debit cards, credit cards and ACH (Federal Reserve System 2013). We considered three different comparison groups to address this issue: families which (1) received UI in at least one month, (2) received direct deposit payroll in 21-31 of the 32 months in the sample and (3) had annual income estimates between \$30,000 and \$80,000. All three groups have similar means for checking account income and spending. More importantly, as shown in Appendix

Figure 2, all three groups have similar trends in spending. We chose the annual income estimate sample as our control group and adjust income and spending using this formula:

$$y_{it} = y_{it,raw} - \left(\bar{y}_t^{30K-80K} - \bar{y}^{30K-80K} \right)$$

where i is a family, t is a month, and $y_{it,raw}$ are the original data. We create an adjusted series y_{it} by subtracting a term equal to the mean for the control group in month t minus the grand mean for the control group across all months in the sample. This modification enables us to examine how income and spending of a family receiving UI change relative to a sample of similar families.

3 Onset of Unemployment

In Section 3.1, we show that income and spending fall immediately at the onset of unemployment. Labor income falls prior to UI receipt, because there is a delay between job separation and arrival of the first UI check. Workers typically cannot file for benefits until they have separated from their job. State UI websites suggest that if everything goes smoothly, a worker will wait three to four weeks between filing her claim and receiving her first benefit check.²³ In the two months before a worker first receives UI, her labor income paid by direct deposit falls by about \$400. Spending on nondurable goods and services falls by \$160, which is 6% of its pre-onset mean. The drop in spending does not reflect shifts to alternative payment channels. Families make up for lost income by drawing down their liquid assets rather than borrowing on their credit cards.

Browning and Crossley (2001) describe three reasons why spending may fall at the start of an unemployment spell – a temporary income loss, a permanent income loss and a decrease in work-related expenses – and Section 3.2 argues that the temporary income loss appears to be the most important explanation. First, in an attempt to isolate the role of temporary income, we show that spending drops more at onset in states where income drops more.

²³<https://labor.ny.gov/directdeposit/directdepositfaq.shtm#DD5>

Second, to understand the role of permanent income losses, we examine the path of family income in the wake of a UI spell and find that by 24 months after onset it has recovered to 95% of its pre-onset level and is on an upward trend. This finding may seem surprising in light of prior work by Jacobson et al. (1993b), but is largely attributable to the fact that we study all UI recipients whereas prior work has focused on high tenure workers who separate in mass layoffs. Finally, we use category-level spending changes at retirement to construct an estimate of work-related spending. We find that an excess drop in work-related expenses can explain 22-31% of the total drop in expenditure at onset.

3.1 Basic Facts About Onset

3.1.1 Spending Drops by \$160 (6%)

Labor income falls sharply at the start of an unemployment spell and UI benefits make up for much of the immediate drop in income. The top panel of Figure 2 shows the path of labor income and UI for a family that receives UI benefits for exactly one month. Labor income starts to decline two months before UI benefits are received and continues to decline through the month in which UI benefits are received. Because labor income drops before UI benefits arrive, the two-month period with the largest decline in income is from three months before UI receipt to one month before UI receipt. Throughout Section 3, this is the two-month window that we study. Because this UI recipient claimed only one month of benefits, she likely found a job during that month and labor income recovers over the subsequent two months.

We construct an aggregate series of income during unemployment and it shows a sharp decline at onset followed by modest declines through the second month in which UI checks are received in the bottom panel of Figure 2. Prior to onset, all future UI recipients are included in the sample. Once UI benefits begin, each point is estimated as

$$\Delta y_t = \frac{1}{n} \sum_{i \in \text{UI duration} > t} y_{i,t} - y_{i,t-1} \quad (1)$$

$$\bar{y}_t = \Delta y_t + \bar{y}_{t-1} \quad (2)$$

where i is a family, t is months since UI receipt began, y is income and n is the number of observations with duration $> t$. This restriction means that in months $t = \{-5, -4, -3, -2, -1\}$ prior to UI receipt, every future UI recipient is included in the sample. In month $t = 0$, everyone who gets UI through month 1 is included in the sample. In month $t = 1$, everyone who gets UI through month 2 is included in the sample, and so on.

From four months prior to onset to two months after onset, UI recipients' monthly labor income drops by \$1,950. Average monthly UI benefits are \$1,300 and an apparent replacement rate of 66% seems unusually large, given that average UI pre-tax replacement rates are around 45% in the US. Differences in the tax treatment of payroll and UI benefits can explain some of the gap. If a paycheck already has a 7.65% payroll deduction and 15% income tax withheld, a \$1,950 post-tax paycheck corresponds to a \$2,400 pre-tax paycheck. Because about 86% of payroll dollars are distributed by direct deposit, the observed drop in payroll is consistent with an average pre-tax replacement rate of 47%. Because of paper checks and the pre-tax to post-tax distinction, the drop in direct deposit family income from three months before UI to the first month in which UI is received is only about \$600 per month in the data.²⁴

Spending drops immediately *before* the start of a UI spell. Figure 3 shows event studies of spending for people who receive UI for different numbers of months. For recipients of all durations, spending falls in the month before UI receipt begins, which coincides with the start of unemployment, as discussed above. The vertical dashed lines in Figure 3 indicate the last month in which UI was received for the bolded data series. For short-duration UI recipients, spending jumps up at the end of a UI spell, although to less than its pre-unemployment level. This spending pattern is consistent with families drawing down savings at the start of an unemployment spell and then building up a buffer stock after the return to work. We explore the recovery in spending further in Section 4.

What exactly is captured by the drop in spending from three months before UI receipt to one month before UI receipt? With i indexing families, t indexing time, and $Post_{it}$ as a

²⁴Adding in payroll paid via paper check increases the estimated income drop to about \$800.

dummy for one month before UI receipt, we estimate $\hat{\beta}$ using the equation

$$c_{it} = \alpha + \beta Post_{it} + \varepsilon_{it} \quad (3)$$

and report the results in Table 4. Conceptually, this drop in spending reflects three distinct economic channels: (1) the direct loss in income from $t-3$ to $t-1$, (2) the news gained from $t-3$ to $t-1$ about the path of future income, and (3) the drop in work-related expenses, if the worker has stopped working between $t-3$ and $t-1$. For equation 3 to capture the causal impact of the three channels, we need to assume that $E(\varepsilon_{it}|Post_{it}) = 0$, which means that the timing of UI receipt is not correlated with something else that might affect spending directly. Because the start dates of UI spells are highly idiosyncratic, this orthogonality restriction seems plausible.

We construct an aggregate series of spending during unemployment and it shows a sharp decline at onset followed by modest declines in subsequent months. The top panel of Figure 4 plots spending separately for each duration group from Figure 3. Each series terminates before the last month of UI receipt, which is when spending recovers for short-duration UI recipients. The bottom panel of Figure 4 plots a composite series of spending while unemployed on the basis of this changing sample using the same methodology as in equation 1. Equation 2 is modified to $\bar{c}_t = \frac{1}{c_{base}} (\Delta c_t + \bar{c}_{t-1})$, where c_{base} is mean spending before UI onset. The small vertical bars around each point indicate the 95% confidence interval for Δc_t . In this composite spending series, spending drops by about 6% at the onset of unemployment (from $t-3$ to $t-1$) and then falls by less than 1% per month in subsequent months.

The drop in spending at onset is substantial relative to the drop in income. To facilitate comparisons of magnitudes, we summarize the drops in income and spending with regressions in Table 4. Labor income paid by direct deposit drops by \$400 from $t = -3$ to $t = -1$. Spending on goods and services consumed immediately – which is only 58% of non-saving outflows – falls by \$160, or about 16% of the drop in income. In Appendix B.1, we document that the shift in spending appears to reflect a true drop in family-wide spending rather than

a shift in spending to alternative payment channels and that our results for this sample are likely to have external validity for other UI recipients.

3.1.2 Decomposition – What Kinds of Spending Drop At Onset? How Is Consumption Smoothing Financed?

Families are able to protect their most important commitments and cut spending most on expenses which might have been related to work. Table 5 shows the drop in spending at onset for several selected categories. Student loans, cash withdrawals, food away from home, and auto expenses all drop sharply.²⁵ If the family owned a car with average gas mileage, the drop in auto expenditures corresponds to driving about 200 fewer miles per month. Notably, mortgage payments are stable at the onset of unemployment.²⁶

To the extent that families smooth their consumption, they do so mostly by drawing down liquid assets. Table 4 indicates that families increase inbound transfers from savings, money market accounts, investment accounts and checking accounts and cut outbound transfers to the same types of accounts.²⁷ Although we are only able to categorize electronic transfers, we believe that families also use paper checks to implement these types of transfers. Table 5 shows that paper check inflows rise during unemployment, even though paper checks from labor income almost surely fell. We find little evidence of actual smoothing on credit cards – the monthly increase in balances across all cards is equal to about 10% of the drop in spending and spending on Chase credit cards falls at onset. See Appendix B.1 for additional credit outcomes.²⁸

²⁵The drop in the fraction of families making student loan payments could reflect debtors becoming delinquent or obtaining deferments on the basis of their unemployment. We believe that it likely reflects delinquency because it takes substantial time to apply for deferment (and related options such as Income-Based Repayment) and debtors are advised to keep making payments until they obtain a deferment.

²⁶Our findings here differ from Gelman et al. (2015), who find that some federal workers delayed mortgage payments during the government shutdown of 2012. However, that shutdown was expected to end in a matter of weeks, meaning that mortgage payment delay carried little financial risk. In contrast, unemployment is of uncertain duration, and so mortgage payment delay carries more serious risks.

²⁷We do not know whether the source account was owned by the same person who owns the checking account, or if this is a transfer from a family member or friend in response to unemployment.

²⁸Herkenhoff et al. (2015) document that families in MSAs with high housing prices instrumented using land unavailability have more access to credit and longer nonemployment durations. This seems to conflict with our findings that average credit utilization is stable during an unemployment spell. One possible reconciliation is that there is some other feature of these MSAs such as higher wages or different skill mix which can explain the differences in nonemployment durations. Another is that increased access to credit affects search behavior even though little of that credit is used in practice.

3.2 Interpreting Onset: Temporary Income Loss, Permanent Income Loss, or Work-Related Expenses?

3.2.1 Spending Drops Most In States Where Income Drops Most

States that pay higher UI benefits show smaller drops in spending at onset, consistent with an important role for temporary income losses. The top panel of Figure 5 plots the change in spending at onset against the change in income at onset for the sixteen largest states in the data, which have at least 3,000 UI recipients. There are many complex rules which affect UI benefit levels and we summarize them by measuring the drop in the sum of labor income plus UI benefits at onset. Our estimated income drop measure accords with outside measures of UI benefit levels; Louisiana, Florida and Arizona are among the five states in the US with the lowest maximum UI benefit levels and New Jersey and Washington are among the three states in the US with the highest maximum benefit levels. States with large income drops also have large spending drops. The slope of the best fit line is 0.23. If we predict out of sample what would the spending drop be in a state which had no income drop at all, we estimate a drop in spending of \$75. In other words, of the \$160 drop in spending at onset, this exercise implies that the majority of the drop in spending is attributable to a temporary income drop rather than lost permanent income or a drop in work-related expenses.

3.2.2 Family Income Recovers Quickly

The bottom panel of Figure 5 shows that family labor income recovers to about 90% of its pre-spell level within 24 months and continues to trend upwards, suggesting that unemployment for this sample may not reflect a large shock to permanent income.²⁹ This finding may be surprising to readers familiar with Jacobson et al. (1993b), where mass layoffs of high-tenure workers cause long-term earnings losses of 30%.³⁰ Intuitively, high-tenure workers who separate in a mass layoff are the most likely of any worker to be adversely affected by a separation. Our paper, in contrast, focuses on the typical UI recipient, who may not have been part of a mass layoff and may not have had high tenure at their firm. We have

²⁹To be precise, income recovers in 24 months to 90% of the value of a control group. In the raw data, incomes for both UI recipients and the control group are trending up.

³⁰Similar results are present in Couch and Placzek (2010), Wachter et al. (2009), Davis and von Wachter (2011), and Jarosch (2015).

compared the path of earnings around UI receipt in the data to a sample in the SIPP, which also shows a similarly rapid recovery in family earnings.³¹ Consistent with the view that the high-tenure mass-layoff selection criteria induce larger earnings losses, we find using the SIPP that earnings losses are larger for high tenure workers and involuntary separations than for all UI spells.

Other government transfers provide additional insurance and, together with the recovery in labor income, family-level insurance is nearly complete. Average monthly government transfers – which include Social Security for the elderly, Disability Insurance, and tax refunds – rise from \$195 per month prior to UI receipt to \$345 per month two years after UI receipt. This increase is concentrated in payments to workers age 59 or older from the Social Security Administration, so we believe that this is driven by people retiring. By month 24, labor income plus government benefits are equal to 95% of their pre-onset level and are trending upwards.

3.2.3 Work-Related Expenses Explain 22-31% of Total Drop At Onset

A person without a job may use her time and money differently, even without any change in family income. A series of papers by Aguiar and Hurst (2005, 2013) has argued that the drop in expenditure at retirement reflects a shift to home production, rather than a failure of consumption smoothing. Non-employment may enable someone to avoid work-related expenses (e.g. fuel to drive to work) and offers an increase in leisure time, with possible substitution to home production (e.g. cooking at home instead of eating out) and increased time spent shopping for low prices (Aguiar and Hurst (2007)). To assess the empirical relevance of these arguments for unemployment, we first categorize expenditures

³¹Appendix Figure 4 compares the monthly path of earnings in the JPMCI data and in the 2004 SIPP. We use the 2004 SIPP rather than the 2008 SIPP because long follow-up horizons in the 2008 SIPP are available only for people who separated at the start of the Great Recession and therefore faced unusually bad job opportunities. This is consistent with findings in Jacobson et al. (1993a) that income for UI recipients recovers after six years to its level immediately prior to separation. Another strand of the literature focuses on displaced workers in surveys such as the PSID and the Current Population Survey (CPS), and does find evidence of persistent earnings losses (Stephens (2001), Farber (2015)). Understanding why the SIPP and administrative records deliver different results from the PSID and CPS is a valuable area for future work. See Appendix B.3 for additional discussion.

by whether they decline at retirement and then examine the drop in spending for these retirement-sensitive categories at the onset of unemployment.

We use changes in spending at retirement to identify which expenditure categories are sensitive to labor force status. We identify retirement transitions using people ages 62 to 70 who started receiving Social Security, and had liquid assets above \$100,000, suggesting that they should be relatively able to smooth their consumption at retirement. The top panel of Figure 6 plots the change in expenditure for 16 merchant categories at retirement and unemployment. The darkness of a each bar is proportional to dollar spending on the category. Some of the merchant categories which drop the most during unemployment are Auto, Food Away From Home, Flights/Hotels, and Department Stores. This aligns well with Aguiar and Hurst (2013)’s findings that Food Away From Home, Transportation, and Clothing decline in the cross-section with age in the CEX. We estimate that work-related expenditures account for 41% of our spending measure.³²

The spending drop at the onset of unemployment is concentrated in work-related expenses, consistent with the predictions of Aguiar and Hurst (2013). The bottom panel of Figure 6 plots the three components of our headline spending measure – work-related expenses on debit or credit cards, other spending on debit and credit cards, and cash withdrawals and bills. While other categories fall by about 5%, work-related expenses fall by 9%.

We estimate that the excess drop in work-related expenses can account for 22-31% of the total drop in spending at onset, by comparing the actual drop in work-related expenses to two counterfactuals with no change in labor force status.³³ The causal impact of interest is how spending would have changed if someone switched from working to not working and began receiving a monthly government income payment of equal value. One way to calculate this is to take the actual drop in work-related spending at onset and subtract a

³²Work-related *card* expenditures are 29% of total spending on nondurable goods and services. If we assume that cash withdrawals are allocated proportionally to the same categories as card expenditures, then work-related expenditures are 41% of total spending. For comparison, Aguiar and Hurst (2013) estimate that work-related expenses are 31% of nondurable expenditures.

³³Baker and Yannelis (2015) estimate the role of work-related expenses using federal government furloughs. They find an estimate larger than ours, but with a confidence interval which contains our point estimate.

counterfactual for how much work expenditures would have changed given a \$500 change in income and no change in work status. One counterfactual comes from using the drop in non-work-related expenses at onset, which implies a \$35 fixed cost of working. Another counterfactual comes from multiplying the marginal propensity to consume out of work-related expenses at benefit exhaustion (8 cents for each dollar of lost income) by the drop in income at onset, which implies a \$50 fixed cost of working.

4 Spending Remains Depressed After Re-employment

In this section, we study the path of spending for workers who find jobs prior to exhausting UI benefits. As already shown in Figure 3, spending recovers slowly upon re-employment. This slow recovery is consistent with a model where agents who have depleted their buffer stock during unemployment rebuild it after they find a job. First, in Section 4.1 we show that this slow recovery after re-employment led to an upward bias in prior estimates of the spending drop during unemployment. We also discuss how our findings might be used by economists studying optimal UI formulas and models of the business cycle. Second, in Section 4.2, we show that the slow recovery in spending is concentrated among families who had little assets at onset. This evidence is consistent with a central role for liquidity in explaining spending fluctuations during unemployment and in particular with models which predict that agents have a target ratio of wealth to income.

4.1 Prior Literature Overstated Spending Drop During Unemployment

A key challenge for prior studies of unemployment was the absence of reliable high-frequency expenditure data. Chodorow-Reich and Karabarbounis (2015) (henceforth CRK) use annual spending data in the CEX to estimate the spending drop during unemployment.³⁴ Kolsrud et al. (2015) (henceforth KLNS) use annual income and asset register data from Sweden. Without higher-frequency data on spending, analysts typically assumed that monthly spending took two values, c^e when employed and c^u when unemployed. For example, if someone

³⁴Although the CEX has quarterly spending data, it only has employment information on an annual basis.

was unemployed for 1 month and spent 3% less annually, the CRK-KLNS method would estimate a drop in spending of 36% during unemployment. Formally, with \bar{c}^e as the pre-unemployment sample mean, and $c_{i,D}$ as the annual spending of someone unemployed for D months, this methodology would estimate the average drop during unemployment as

$$\widehat{c^u/c^e} = \frac{1}{n} \sum_{i \in u} \sum_{D \in \{1 \dots 12\}} \frac{c_{i,D}/\bar{c}^e}{D/12} \quad (4)$$

Families engage in substantial smoothing *within* the year of an unemployment spell and methodologies which neglect this overstate the drop in spending during unemployment. The top panel of Figure 7 plots the average monthly spending of a family with a completed UI duration of three months. Average spending during unemployment was 6% lower than the pre-onset mean while receiving UI, and 2.5% lower than the pre-onset mean in the subsequent 9 months. The light blue arrows indicate the estimated spending drop using equation 4; an analyst using this equation would have estimated a drop in spending of 13% during unemployment. The bottom panel repeats the exercise separately for families of different UI durations and shows that the bias is substantial at short durations. Table 6 reports the drop in spending at onset for various categories as well as the drop estimated from implementing equation 4. The drop at onset is 6% for all nondurables and 6% for food. Applying equation 4 in the data, we estimate drops of 20% and 10% respectively.

Suitably adjusted, our estimates are broadly in line with prior work using survey data. CRK estimate that spending on nondurables drops by 13% in the CEX and spending on food in the PSID drops by 8%.³⁵ Replicating their methodology in the data yields a 16% drop in nondurables and a 10% drop in food expenditures. Browning and Crossley (2001) study a survey which asked UI recipients after six months how much their monthly expenditure had fallen since the time of their job separation. The mean drop in spending was 14%, which is a bit larger than our estimate of a 10% drop from onset to six months later.³⁶

³⁵Table 2 in their paper reports a drop in nondurables spending in the CEX of 23% and of food spending in the PSID of 14% for a family transitioning from all its adult members being employed to all its members being unemployed. Separated workers are responsible on average for 57% of family earnings (Table 1), so we adjust the CRK estimate to 13% for nondurables and 8% for food respectively. Finally, because CRK are interested in the average spending of an unemployed family relative to an employed family, we weight each family's estimated spending drop using equation 4 by its duration of UI receipt.

³⁶However, this method of estimating spending drops may be biased upward due to telescoping, where

Our estimates are also in line with estimates by KLNS for the short-term unemployed in Sweden. They estimate that spending drops by 4% at onset in a point-in-time expenditure survey and 11% using annual income and asset register data.³⁷ Our comparable statistics are a 6% drop at onset and a 19% drop using equation 4. Time aggregation may help explain why KLNS estimated a much larger drop in spending for the short-term unemployed in the annual data than in the point-in-time survey data.³⁸ These arguments do not apply, however, to people who are unemployed for a full calendar year. If time aggregation is a problem for their estimates in the way we hypothesize, this further bolsters their economic argument for local welfare gains from an increasing benefit profile.

There are two distinct research literatures which are interested in the drop in spending during unemployment – economists evaluating optimal UI using the Baily (1978)-Chetty (2006) formula and economists building models of the business cycle. Substituting our estimates of the spending drop during unemployment for CRK’s or KLNS’s and subtracting the fixed cost of work shrinks the estimated gap in marginal utilities between the employed and unemployed states, lowering the apparent benefits of UI. However, a dynamic model which incorporated decreased spending after re-employment would offset this to some extent. We leave a formal reevaluation of the Baily-Chetty formula to future work. Separately, for business cycle modelers who are interested in how unemployment affects output through product demand, the relevant statistic is probably the annual spending decrease associated with an unemployment spell, rather than the drop in spending during unemployment (CRK, Kaplan and Menzio (2015)).

respondents accidentally include expenditures prior to the sample reference period, as discussed in Browning et al. (2014).

³⁷Table 4 of KLNS reports a spending drop of 19% for individuals with short UI durations. We adjust this to the family level using a ratio of family spending drop to individual spending drop reported in Section 5.2 ($15.8/26.3=0.6$).

³⁸Time aggregation may also help explain why KLNS do not observe smoothing out of liquid assets during short unemployment spells. If agents decumulate and rebuild their buffer all within one year then this will look like no smoothing out of assets when using an annual point-in-time asset estimate.

4.2 Low-Asset Families Have a Slower Spending Recovery Upon Re-employment

A broad class of consumption models predict that agents will have a “target ratio” of wealth to permanent income, but this prediction is not a feature of the permanent income model or the hand-to-mouth model. By target ratio, we mean that when wealth is below this level, agents will consume less until their wealth returns to this level. Examples of models with this property include Carroll (1997), Laibson et al. (2015), Gourinchas and Parker (2002), and Kaplan and Violante (2014). In a permanent income model, in contrast, agents consume the annuity value of their wealth each period and so their consumption is insensitive to small wealth fluctuations. In a hand-to-mouth model, by definition, consumption is insensitive to wealth.

For an income shock of a fixed size, families with little initial assets need to draw down a larger *share* of their assets in order to achieve the same amount of consumption smoothing. Then, having drawn down assets to weather the income shock, models with a target ratio, and sufficient curvature of utility around that target ratio, predict that spending will remain depressed longer for families with little initial assets. We test this prediction by studying high-, medium- and low-asset families that receive UI for exactly three months.³⁹ The top panel of Figure 8 shows that the path of labor income plus UI benefits is very similar for three groups. Integrating over the path of income for families with a three month UI spell indicates a total income loss equal to 0.58 months of pre-onset income for the low-asset group, 0.61 months for the medium-asset group, and 0.66 months for the high-asset group.

The low-asset group uses up a larger share of its assets and recovers spending more slowly, which is consistent with target ratio behavior. On the basis of the gap between the drop in spending and the drop in income, we estimate that high-asset families use up 0.46 months of assets (which is 15% of the median total liquid assets within this group), while low-asset families use up 0.27 months of assets (which is 97% of the median total liquid

³⁹We stratify families using JPMCI’s internal estimate of a family’s total liquid assets – across all financial institutions. This estimate is based on a wide variety of data sources which update at different frequencies and is suitable for examining heterogeneity in long-run asset holdings, but not for understanding month-to-month changes in total liquid assets.

assets within this group). (Additional statistics are reported in Appendix Table 3.) The bottom panel of Figure 7 shows spending recovers quickly for the high-asset group and more slowly for the low-asset group. Quantitatively, after re-employment, high-asset families cut spending enough to rebuild 0.11 months of lost income, while low-asset families cut spending enough to rebuild 0.43 months of lost income. Understanding the source of heterogeneity in asset holdings would be useful to interpret our findings in this section further. Low-asset groups could have lower optimal target ratios because of different time preferences or different income risk profiles or they could simply have experienced a series of negative income shocks.

5 UI Benefit Exhaustion

UI benefit exhaustion provides an informative test of theories of consumption behavior because exhaustion causes no change to opportunities for home production and no change to labor market productivity.⁴⁰ The change in income at benefit exhaustion is large, with \$1,350 of lost benefits, and is predictable. With a monthly job-finding rate of 25%, the probability of exhaustion is 75% one month before, 56% two months before, and so on.⁴¹ What *should* happen to spending at exhaustion? A liquidity-constrained consumer with no assets at the onset of unemployment may cut spending gradually, but will have no excess drop in the month in which she exhausts benefits (Jappelli and Pistaferri (2010), Section 3.2). We formalize this prediction in the model.

In practice, we find that spending drops sharply by \$259 in the month benefits are exhausted. Spending drops when benefits are exhausted, so it drops sooner in Florida, which offers at most 16 weeks of benefits, than it does in most states, where benefits last for 26 weeks. Spending drops across a wide variety of categories, including food at home, retail

⁴⁰Formally, UI recipients are required to search for jobs and so UI recipients might have more time for home production after benefit exhaustion. However, our understanding is that these search requirements are rarely enforced.

⁴¹To study the experience of a typical UI recipient, our analysis studies people who exhausted benefits in February 2014 or later. These people were eligible for at most 26 weeks of benefits. Some states had lower potential benefit durations: Kansas (20 weeks), Michigan (20 weeks), Florida (16 weeks) and Georgia (18 weeks).

purchases, entertainment and medical copays. To the extent that families are able to smooth this income shock, they do so by drawing down their liquid assets. Such a discontinuous drop is quite surprising and we explore possible explanations in the model section.

5.1 Income Drops Sharply at Exhaustion

The exhaustion of UI benefits causes a substantial negative loss in monthly family income, as shown in the top panel of Figure 9.⁴² Lost UI benefits were about \$1,350 per month, or 37% of median income prior to onset. Labor income rises by about \$400 and other income rises by \$50 per month, so the drop in monthly family income is about \$900. Labor income rises at exhaustion for three reasons: (1) some UI recipients would have found jobs even if benefits continued, (2) other family members may increase their labor supply (Cullen and Gruber (2000), Stephens (2002), Rothstein and Valetta (2014), Blundell et al. (2015)), and (3) search effort and job-finding rates are higher at benefit exhaustion (Katz and Meyer (1990), Schmieder et al. (2012), Card et al. (2007), Krueger and Mueller (2010), DellaVigna et al. (2014)).

5.2 Spending Drops Sharply At Exhaustion

Spending drops by \$22 per month in the months leading up to exhaustion and by \$260 (12%) in the month after benefits are exhausted, as shown in the bottom panel of Figure 9.⁴³ The estimating equation for exhaustion is the same as the equation for onset: $(c_{it} = \alpha + \beta Post_{it} + \varepsilon_{it})$. The orthogonality restriction for this regression is $E(\varepsilon_{it}|Post_{it}) = 0$, which means that the timing of exhaustion is not correlated with something that might affect spending directly. Because the start dates of UI spells are highly idiosyncratic, exhaustion

⁴²We define exhautees as families who received UI benefits equal to the maximum number of allowed weeks in each state, with a window of two weeks to allow for administrative noise. Some UI recipients (perhaps 20%) with limited earnings histories are eligible for less than the maximum duration of benefits and we are unable to identify these exhautees. To adjust for differences in benefit duration across states, we organize our plots in this section around the month in which the last UI check was received for benefit exhautees. Appendix Figure 6 shows an event study of UI benefits for exhautees for the six largest states in the JPMCI sample – the shorter duration of UI benefits for Florida and Michigan is clearly evident.

⁴³Table 4 reports the percent change at exhaustion relative to the pre-onset mean, which is 10%. Here, we report the drop as a percent of the spending level prior to exhaustion, which is 12%.

dates are also idiosyncratic and so this orthogonality restriction seems plausible. Note this restriction does not rule out extra job search at exhaustion or that exhaustion causes a family to make a new plan for their spending; this is part of the causal impact of exhaustion. In the rest of our analysis, to deal with time aggregation, we define the drop at exhaustion as the change in spending over a two-month window so that we can study all exhaustees.⁴⁴

The best evidence that the drop in spending at benefit exhaustion is caused directly by benefit exhaustion comes from differences across states. Appendix Figure 6 shows the path of spending over time for UI exhaustees for the six largest states in the data. Florida and Michigan offer maximum durations of UI benefits less than 26 weeks. Spending declines at the same time benefits are exhausted in these states, which is well before the time when spending declines in states that offer the traditional 26 weeks of benefits.

5.3 Decomposition – What Kinds of Spending Drop? How Is Consumption Smoothing Financed?

The drop in spending at benefit exhaustion appears to reflect a change in a family’s actual consumption bundle from the prior month, rather than simply a delay in purchases of durable goods or a decrease in payments on outstanding debts. The top half of Table 7 decomposes the drop in outflows into nine different categories. In a reversal of the patterns we documented at onset, non-work-related expenses on cards fall *more* than work-related expenses. The categories which drop most are food at home, retail purchases and the presence of any medical copay, as shown in Table 7. Aguiar and Hurst (2005) compare the diets of employed and unemployed people, controlling for a wide variety of observables, and report a similarly-sized gap in spending on food at home between the employed and unemployed (9-15%) to the drop we see at exhaustion. They estimate that unemployment causes a five percentage point increase in any hot dog consumption and a nine percentage point decrease

⁴⁴One important technical wrinkle for estimating the spending drop at benefit exhaustion comes from time aggregation – we have monthly income and spending data, but benefits are paid on a weekly or biweekly basis. In our plots in Figure 9, we limited the sample to exhaustees who received their last UI check on the 25th of the month or later. These families have a sharp drop in UI income from one month to the next and also a sharp drop in spending. However, the monthly structure of the data means that UI benefits appear to phase out over two months for most families. Appendix Figure 7 shows that the magnitude of the two-month spending drop for all UI exhaustees is very similar to the magnitude of the one-month drop for exhaustees who get their last check at the end of the month.

in any fresh fruit consumption, suggesting that there is a substantial change in diet quality at exhaustion. In addition, the share of families with any entertainment expenditures, which was stable at onset, drops by about 10% at exhaustion.

At exhaustion, families appear to prioritize their most important financial commitments, which show relatively small drops in spending. Table 7 shows that the drop in spending is smallest for utility payments, auto loans and mortgage payments. Delinquency measured in credit bureau records and credit scores are all relatively stable (Appendix Table 2). There is little evidence to suggest that benefit exhaustion does immediate damage to a family’s long-term financial health.⁴⁵ The data are consistent with prior work on consumption commitments by Chetty and Szeidl (2007), where families cut spending on some flexible expenditure categories sharply to protect their long-run commitments.

To the extent agents smooth their spending at exhaustion, they do so by drawing down liquid assets. Dissaving inflows spike, as do paper checks, as shown in Table 7. Agents also draw down their checking account balance, as shown in Table 4. Again, we find only a modest increase in credit card borrowing; spending on Chase credit cards does not increase, and balances rise because families make smaller payments on their outstanding credit card debt.

6 Performance of Benchmark Consumption Models

In this section, we compare the actual path of spending during unemployment to benchmark models of consumption. First, in Section 6.1 we show that unemployment is a good way to test alternative consumption models, since unemployment is a large shock to income, implying that hand-to-mouth behavior cannot be consistent with near-rationality. Then, we describe the setup of our model in Section 6.2. To capture “buffer stock” consumers in the tradition of Deaton (1991), and Aiyagari (1994), we do not allow agents to borrow

⁴⁵The decline in the presence of medical copayments, however, could imply that families are delaying important health expenditures.

at all in our baseline parameterization. As an alternative scenario, to capture “permanent income” consumers in the spirit of Friedman (1957), Modigliani and Brumberg (1954), and Hall (1978), we allow agents to borrow against their future income at interest rate R .⁴⁶ The drop in spending from onset through exhaustion in the data matches the buffer stock model, assuming that agents start their unemployment spell with liquid assets equal to one month of income. Next, in Section 6.3, we show that the buffer stock model does a better job of fitting the data than either a permanent income model, or a hand-to-mouth model. Finally, in Section 6.4, we explore two major shortcomings of the buffer stock model relative to the data – it predicts substantially more asset holdings at onset and it predicts a much smoother path of spending around benefit exhaustion.

6.1 Why Unemployment is a Good Test of Alternative Consumption Models

Unemployment is a powerful setting for testing alternative consumption models, since it causes a large shock to income, implying that myopic behavior is not approximately rational using a welfare metric. A large literature uses the spending response to temporary income shocks such as tax rebates to test between models with and without liquidity constraints. Most papers in this literature consistently find a higher marginal propensity to consume (MPC) than would be predicted for a permanent income consumer without liquidity constraints. Many authors interpret these high MPCs as evidence in favor of buffer stock models. However, an alternative interpretation is that agents’ choices are consistent with near-rationality (Cochrane (1989)).⁴⁷ Proponents of this view argue that the welfare costs of failing to smooth income shocks of the magnitude observed in the literature are quite small, and that such small deviations from optimality are not sufficiently compelling evidence to reject the permanent income model.

⁴⁶In the “permanent income” models cited above, because agents can borrow against their future income, spending is insensitive to temporary income fluctuations. Not all models which allow agents to borrow against their future income have a low sensitivity of spending to current income (see Carroll (1997) for a counterexample), but our model does have this feature.

⁴⁷Papers which use near-rationality to explain consumption fluctuations include Kueng (2015), Reis (2006) and Caballero (1995).

Fuchs-Schuendeln and Hassan (2015) develop a framework to evaluate the near-rationality claim. For a given temporary income change, they calculate the welfare cost of behaving like a hand-to-mouth consumer and failing to adjust spending in order to perfectly smooth the shock. Specifically, consider an agent with regular monthly income y who receives a one-time tax rebate of x . They calculate a measure of equivalent variation as the additional monthly income v that a consumer would require to be indifferent between consuming all of the tax rebate x in one month plus v in every month over the year, and smoothing the tax rebate over one year. In other words, they find the v which solves

$$\underbrace{u(y + x + v)}_{MPC=1} + 11 \cdot u(y + v) = 12 \cdot \underbrace{u\left(y + \frac{x}{12}\right)}_{\text{perm income}}, \quad (5)$$

for CRRA utility with $\gamma = 2$, and they define $EV = \frac{v}{y}$. They perform this calculation for the income changes examined in 18 recent empirical papers in this literature. Their findings are shown in the blue bars in Figure 10. They find that acting like a hand-to-mouth consumer who fails to smooth spending has a welfare loss smaller than losing 1% of monthly consumption over a year in most cases, and no more than 5% in any case. We perform the same calculation for the income loss associated with an unemployment spell which lasts at least six months, and show this as the yellow bar in Figure 10. The welfare cost of failing to smooth the income loss associated with a UI spell terminating in exhaustion, and instead acting like a hand-to-mouth consumer, is equivalent to 20% of annual consumption.⁴⁸

The large income change associated with unemployment also enables us to test theories of excess sensitivity motivated by transaction costs. Kaplan and Violante (2014) build a model with a transaction cost of accessing an illiquid asset which offers higher returns than liquid asset holding. A key prediction of their model is that the excess sensitivity of spending to tax rebates is falling in rebate size: agents immediately consume 15% of a \$500 rebate, but only immediately consume 3% of a \$5,000 rebate, as shown in Figure 8 of their paper. The average UI spell entails an average loss of \$8,500 of income. Because the size of the income

⁴⁸Formally, we calculate this as the scalar v in monthly consumption which solves $\sum_j w_j \sum_{t=1}^{15} \beta^t u(c_{t,j}^{PIH}) = \sum_j w_j \sum_{t=1}^{15} \beta^t u(c_{t,j}^{H2M} + v)$ where j indexes different employment histories after benefit exhaustion and w_j is the probability of each employment history.

loss is uncertain, the motive to liquidate at UI onset is even stronger than when the rebate size is known with certainty. As a result, the logic of the model suggests that Kaplan and Violante (2014) predict a withdrawal from the illiquid asset at the start of an unemployment spell, followed by relatively stable consumption during unemployment. However, we have not explicitly modeled the dynamics of when the agent would choose to pay the liquidation cost and this is a fruitful area for further research.

6.2 Model Setup

We calibrate a finite-horizon buffer stock model of consumption and savings. Agents have Constant Relative Risk Aversion (CRRA) utility, and choose their level of consumption each month, c_t , to maximize their expected discounted flow of lifetime utility. Agents earn a monthly return of R on their beginning of month assets a_t . Income z_t is risky because of unemployment; this risk is partially insured by unemployment benefits, which expire after six months. Employment follows a Markov process Π where agents transition between employment and unemployment. The agent's problem in month t can be written as

$$\begin{aligned} \max_{\{c_t\}} \quad & \mathbb{E} \sum_{n=0}^{T-t} \beta^n u(c_{t+n}) \\ \text{subject to } & c_t + a_{t+1} = Ra_t + z_t \\ & c_t \geq 0 \\ & a_{t+1} \geq -b_t \\ & Ra_T + z_T - c_T \geq 0 \end{aligned}$$

where β is the monthly discount factor, $u(c) = \frac{c^{1-\gamma}}{1-\gamma}$, z_t evolves according to transition matrix Π , T is the number of months in the agent's life, and b_t is the borrowing limit. The last inequality is a budget balance condition at the end of life.

To capture two different benchmark models of consumption, we consider two different

asset constraints. First, to capture buffer stock consumers, we consider a case where agents cannot borrow ($b_t = 0 \forall t$). Second, to capture permanent income consumers we allow agents to borrow against their future income at interest rate R . A “natural borrowing constraint” (Aiyagari (1994)) arises because the agent must pay all her debts before death and have positive consumption in every period. Therefore, in any period the natural borrowing constraint is the present discounted value of the minimum possible future income flows, which are bounded below by the income value for an agent who has exhausted UI benefits.⁴⁹

Given an environment $\{R, z, \Pi, b\}$ and preferences $\{\beta, \gamma\}$, there is an optimal consumption path $c_t^*(a, z)$ which satisfies

$$u'(c_t) = \max\{\beta R \mathbb{E}_t[u'(c_{t+1})], u'(Ra_t + z_t + b_t)\}$$

We calibrate the model using the JPMCI data and the standard preference parameters summarized in Table 8.

- Income – We normalize income to 1.0 in the employed state. To match the data, we set income to 0.84 while receiving UI benefits and 0.53 after UI benefit exhaustion. Income does not fall to zero after exhaustion because our income concept includes labor income from all family members, non-labor income, and government transfers.
- Transition Rates – The transition rate from unemployment to employment is 25%, which matches the UI exit rate in the data. We do not observe job-finding after benefit exhaustion; we assume that it is 25% in all months *except* the month benefits are exhausted, when we set the job-finding rate to 30% to match evidence from Card et al. (2007). In a robustness check, we consider an alternative specification where the job-finding rate is permanently lower after exhaustion. We choose a separation rate to UI of 3.25% in order to match the 11.5% of families with an unemployed member during 2013 and 2014 (Bureau of Labor Statistics, 2014).

⁴⁹Formally, we set $b_t = \sum_{s=0}^{T-t-1} \frac{z_{min}}{R} \left(\frac{1}{R}\right)^s$ where z_{min} equals the income for an agent who has exhausted UI benefits.

- Preferences and Environment – For the preference parameters β and γ we choose standard values of 0.996 (translating to an annual discount rate of 5%) , and 2.0. We choose a monthly real interest rate of 0.25%, which translates to an annual interest rate of 3%. We consider a time horizon of 240 months, corresponding to a middle-aged worker with 20 years left in her career.

Given these parameter values, we solve the consumer’s problem numerically using the method of endogenous gridpoints suggested in Carroll (2006). This method returns optimal consumption $c_t^*(a_t, e_t)$ as a function of the agent’s beginning of month assets and their employment status $e_t = \{E, U1, U2, U3, U4, U5, U6, U7, U8\}$, where $U8$ represents benefit exhaustion.⁵⁰

6.2.1 An Over-identification Test Using Liquid Asset Holdings

We choose the asset holdings at onset which best match the spending drop for an agent who cannot borrow and this comes very close to matching the actual liquid asset holdings in the data. An agent who becomes unemployed after $t = 0$ with assets a_0 and stays unemployed through benefit exhaustion sees a consumption drop of $\Delta c^{model}(a_0) \equiv c_{post-exhaust,t}^*(a_t^*(a_0), U8)/c_0^*(a_0, E)$. We choose $a_0^{best-fit}$ such that

$$\Delta c^{model}(a_0^{best-fit}) = c_{post-exhaust}^{data}/c_{pre-onset}^{data}$$

Because our model has exactly one free parameter ($a_0^{best-fit}$) and we match one sample moment ($c_{post-exhaust}^{data}/c_{pre-onset}^{data}$), the model is exactly identified and we estimate assets at onset equal to 0.84 months of income. We do not observe total liquid asset holdings in the Chase data, so we estimate them using an adjustment factor from the SCF. Specifically, we

⁵⁰The combination of asset level and employment status determines beginning-of-period cash on hand $m_t = Ra_t + z_t(e)$, which is formally how the model is solved. In Section 3.1, we documented that the decline in family income occurs one month *before* UI receipt begins because of a time lag between job separation and the beginning of UI receipt. To match this feature of the data in the model, we assume that UI benefits actually last 7 months rather than six months.

estimate

$$a_0^{data} = \frac{(\text{Total liquid assets})_{SCF}}{(\text{Checking account balance})_{SCF}} \cdot \frac{(\text{Checking account balance})_{Chase}}{(\text{Pre unemployment monthly income})_{Chase}} = 0.71,$$

This is very close to the 0.84 months which fits the spending drop from onset through exhaustion.⁵¹ There are many reasons that liquid assets held by agents at onset might not reflect the total amount of assets they might have available to help them smooth a large shock such as unemployment. For example, agents might receive transfers from their parents, or they might be able to sell consumer durables. Our results suggest that these channels are not quantitatively important channels for consumption smoothing.

6.3 Buffer Stock Fits the Data Better Than Permanent Income or Hand-to-Mouth Consumers

We compare the path of consumption predicted by our buffer stock model to the path of spending observed in the data. To enable this comparison, we need to assume that the nondurables spending in the data is the same as consumption in the model.⁵² In our robustness checks, we examine total spending as well. The buffer stock model fits some aspects of the spending path during unemployment, as shown in the top panel of Figure 11. By construction, the buffer stock model matches the level of spending at onset and exhaustion. Not by construction, the buffer stock model matches the drop at the onset of unemployment. In the model, families cut spending additionally each month that they stay unemployed. This matches the data qualitatively – families in the data are cutting spending from months two through 5 – but not quantitatively, since the model predicts larger spending cuts while receiving UI and no excess drop in spending at benefit exhaustion. We focus on

⁵¹This is slightly smaller than the one month's income worth of liquid asset holdings prior to unemployment estimated by Chetty (2008) in survey data.

⁵²To ensure comparability between the model and the data, we make two adjustments to our prior data analysis. First, we analyze the subset of UI spells where potential benefit duration was 26 weeks at the start and end of UI receipt. Second, we adjust the spending series to reflect the spending of agents who remain unemployed after benefit exhaustion. In the data, we observe average spending in the month after benefit exhaustion for the unemployed and the re-employed together. We assume that spending is constant for the 30% of agents that are re-employed in the month of benefit exhaustion (as it is for agents who are re-employed after 3, 4, or 5 months of unemployment) and estimate the drop in spending for the unemployed alone as 1.43 times the drop for the pooled sample.

this failing of the model in Section 6.4.

Next, we show that the buffer stock model outperforms the permanent income benchmark and the hand-to-mouth benchmark using the Cochrane (1989)-Fuchs-Schuendeln and Hassan (2015) welfare metric. As discussed above, we implement a permanent income benchmark by allowing the agent to borrow out of her future income. We also consider a hand-to-mouth agent who sets consumption equal to current income each period.⁵³ We calculate v , the increment to monthly spending needed to make the agent indifferent between her choices in the data and her predicted choices under the different benchmark models. This is given by a modified equation (5) where w_j reflects probabilities of different employment histories:

$$\sum_j w_j \sum_{t=1}^{15} \beta^t u(c_{t,j}^{model}) = \sum_j w_j \sum_{t=1}^{15} u(c_{t,j}^{data} - v) \quad (6)$$

We aggregate over all possible job-finding histories in the eight months after exhaustion.⁵⁴ We find that the path of spending we observe in the data represents a 7% gain relative to the hand-to-mouth path, and a 13% loss relative to the smooth permanent income path, as shown in the bottom panel of Figure 11. In contrast, the welfare loss of the deviations from the buffer-stock path shown in Figure 10 is about 1%. We interpret this as evidence in favor of the buffer-stock model with little assets at onset, relative to either of the alternatives considered here.

Two key lessons from the model are that we can fit the drop in spending from onset through exhaustion assuming families hold little liquid assets at onset, but that we cannot fit the monthly drop at exhaustion. These conclusions continue to hold under a number of alternative assumptions which we discuss in Appendix B.4 and show graphically in Appendix Figure 8.

Nevertheless, our conclusions are highly sensitive to assumptions about families' assets

⁵³This corresponds to a special case of the rule-of-thumb consumer in Campbell and Mankiw (1989), where $c_t = \alpha z_t$, with $\alpha = 1$. This model is commonly used in the public economics literature when studying unemployment. Examples include Mortensen (1977), Shimer and Werning (2007), Rothstein (2011), and Krueger and Mueller (2014).

⁵⁴For c_t^{data} we assume that the agent behaves optimally between exhaustion and re-employment according to the buffer-stock model given the assets they have left at this point, and then once re-employed, they adjust their spending such that they match the assets of buffer-stock agents by month 15.

prior to unemployment. In the bottom right panel of Appendix Figure 8 we show the model predictions assuming agents either have no assets, or have assets equal to one year’s worth of income. Agents with assets equal to one year’s worth of income smooth spending considerably throughout the spell, whereas agents with no assets cut their spending substantially more as the spell progresses.

6.4 Failings of the Buffer Stock Model

While the buffer-stock model does a reasonable job of matching the overall path of spending, it has two major failings relative to the data. First, it predicts substantially more asset holdings at onset. Second, it predicts that spending does not drop discontinuously at benefit exhaustion.

6.4.1 Failure 1: Agents Hold Too Little Liquid Assets at Onset

A key prediction of buffer stock models is that agents should accumulate precautionary savings to self-insure against income risk. In our model with only temporary income risk, we calculate that agents should hold liquid assets equal to about 2.4 months of income, which is three times the asset holdings which fit the spending drop from onset to exhaustion. Models with realistic income processes – including permanent income risk and retirement – predict much higher asset holdings. Gourinchas and Parker (2002) estimate that an agent’s target buffer stock is about 12 months of assets early in life and rises to over 60 months as retirement approaches. Laibson et al. (2015) estimate a model where they match illiquid wealth holdings equal to 31 months of income.

Why might agents be holding so little liquid assets, even when this means that spending appears to be so sensitive to income? There are two broad classes of reasons why this might be the case. First, monthly spending on nondurable goods from bank accounts may not accurately capture fluctuations in consumption. Purely from a measurement perspective,

this could arise if consumption rises through in-kind transfers or purchases made with cash not deposited in the bank account. Even if bank accounts accurately capture the goods a family purchases each month, even nondurables have a shelf life such that consumption flows are more stable than expenditures. Second, even if consumption does fluctuate from month to month, there are some preferences which can rationalize this behavior. With a low coefficient of risk aversion, a family could be very willing to substitute consumption across periods. A model with quasi-hyperbolic preferences such as Laibson (1997) predicts low liquid asset holdings from highly impatient consumers.

A related puzzle which merits further work is why agents do not seem to use the borrowing channels which are available to them. For example, we have documented almost no change in credit card borrowing during unemployment. The *monthly* interest rate on credit cards is about 1% in the UI recipient sample. And to the extent that agents can default on credit card debt if their income remains low as in Herkenhoff (2015), the argument for borrowing on credit cards while unemployed is even stronger.

6.4.2 Failure 2: Agents Cut Spending Too Slowly During UI Receipt and Too Much at Exhaustion

We have not been able to find a parameterization of our model in which agents have rational expectations which matches the very slow average monthly decline during UI receipt (0.6% per month) and the 12% drop in spending at benefit exhaustion. Two specific scenarios shown in the top panels of Figure 12 help clarify why this pattern is difficult to model. First, we consider a scenario where agents have no assets two months prior to their unemployment spell. These agents cut spending rapidly at the start of an unemployment spell to the level of UI benefits. As the unemployment spell wears on, they cut spending further, *below* the level of UI benefits, in order to build a buffer which will help offset the income drop at benefit exhaustion. Second, we consider a scenario where agents have a 10% *monthly* discount rate (e.g. DellaVigna and Paserman (2005)). These agents draw down their assets at the start of an unemployment spell such that by month two, consumption is equal to the level of UI benefits. As exhaustion approaches, even these agents build a small buffer.

Benefit exhaustion does not appear to be associated with a permanent change in re-employment wages (von Wachter et al. (2015)) and this enables us to rule out certain theories about the drop at benefit exhaustion. First, if agents discretely received negative news about their productivity at benefit exhaustion, then we would expect re-employment wages to be permanently lower. Second, if agents were present-biased then they would face a liquidity shortfall in every month after benefit exhaustion and again we would expect permanently lower re-employment wages. If von Wachter et al. (2015)'s findings also hold in the US then our results are hard to reconcile with productivity updating or present-bias.

One simple deviation from rationality which can explain the drop at exhaustion is over-optimism about job-finding at the end of a UI spell combined with pessimism (or lack of effort) earlier on in the a spell. Spinnewijn (2015) finds that on average unemployment spells last more than three times longer than workers expect at onset. Why might workers be over-optimistic and then cut their spending at exhaustion? One possibility is that if they searched little while receiving UI, they might have an inflated view of how easily they can get a job. When they raise their search effort at exhaustion and do not find a job, this leads them to update their beliefs about how quickly they can get a job. Another possibility is that they were expecting to be recalled to their previous job but this did not pan out (Katz and Meyer (1990)).

We find that we can match both the drop from onset to the last month of benefits and the drop at exhaustion if we assume that agents believe their job-finding rate is only 10% while receiving UI, but jumps dramatically to 70% in the last month of benefits. (Recall that in fact the job-finding rate is about 25% in most months of UI receipt and 30% in the month UI benefits are exhausted.) The path of consumption predicted by such a model is plotted by the yellow line in the bottom right panel of Figure 12. In this scenario, exhaustion without finding a job is much more unexpected than it is with accurate beliefs about job-finding probabilities. One month before exhaustion, families believe there is only a 30% chance of being unemployed at exhaustion. Two months before, the probability is 23%. Since the potential income drop associated with exhaustion is (erroneously) assumed to be

a low-probability event, families rationally choose not to cut spending much in anticipation of this event.

Another possibility is that agents have correct beliefs about their job-finding probabilities, but are inattentive in their monthly consumption decisions. We showed in Section 6.3 that an agent whose optimal spending path followed a buffer stock model would incur little welfare loss from making the choices observed in the data. One example of a specific friction comes in a model developed by Reis (2006). In his model, agents rationally respond to the costs of processing information about their finances by infrequently updating their budgets, remaining inattentive between updating dates. Inattention among some agents during UI receipt, followed by attention from all agents at benefit exhaustion, might explain the patterns we see in the data. Estimating a model with inattention using these spending patterns is an interesting area for future research.

7 Conclusion

In this paper, using spending records from the JPMorgan Chase Institute, we built a dataset to study how unemployment affects spending. To summarize our results, we find that families do insufficient self-insurance, in the sense that spending is quite responsive to income. We document that unemployment causes a large but short-lived drop in income, generating a need for liquidity. Spending on nondurables falls by 6% at the onset of unemployment and work-related expenses explain about one-quarter of the drop in spending. People receiving UI keep their spending low after re-employment, perhaps in order to rebuild their financial buffer. For people who exhaust UI benefits, spending drops by an additional 12%.

We compare the path of spending in the data to three benchmark consumption models: buffer stock, permanent income and hand-to-mouth. Prior work on excess sensitivity of spending to income had been criticized on the grounds that the observed behavior was consistent with near-rationality; because unemployment is such a large shock to income, this criticism is less relevant for our work. The predictions of the buffer stock model are

much closer to the data than the alternatives. However, there are two important failings of the buffer stock model: families in the data have less assets at onset than predicted by the model and spending drops much more in the data at exhaustion than predicted by the model.

We see at least three fruitful avenues for future work. First, we find that families act during unemployment as if they have little liquid assets and little access to credit. But we see in the data that these families have room to borrow on their credit cards and from surveys that these cash-poor families have substantial illiquid assets in housing and retirement accounts (Angeletos et al. (2001), Kaplan and Violante (2014)). Why do families not use these mechanisms to help smooth spending? And why do families not hold more of a liquid buffer stock against risks like unemployment and health shocks? Second, we documented a sharp drop in spending at the exhaustion of UI benefits which is hard to fit into a model with forward-looking agents who have rational expectations about job finding. Future work should try to understand which theories of unemployment and/or consumption best explain this drop. Finally, we have focused in this paper entirely on a partial equilibrium model of unemployment and spending. Understanding the general equilibrium effects of spending by UI recipients is important for both models of optimal UI and models of the business cycle (Kekre (2015)).

References

- Aguiar, M. and Hurst, E. (2005). Consumption versus Expenditure. *Journal of Political Economy*, 113(5):919–948.
- Aguiar, M. and Hurst, E. (2007). Life-Cycle Prices and Production. *American Economic Review*, 97(5):1533–1559.
- Aguiar, M. and Hurst, E. (2013). Deconstructing Life Cycle Expenditure. *Journal of Political Economy*, 121(3):437 – 492.
- Aiyagari, S. R. (1994). Uninsured Idiosyncratic Risk and Aggregate Saving. *The Quarterly Journal of Economics*, 109(3):659–84.
- Akerlof, G. A. and Yellen, J. L. (1985). A Near-rational Model of the Business Cycle, with Wage and Price Intertia. *The Quarterly Journal of Economics*, 100(5):823–38.
- Angeletos, G.-M., Laibson, D., Repetto, A., Tobacman, J., and Weinberg, S. (2001). The Hyperbolic Consumption Model: Calibration, Simulation, and Empirical Evaluation. *Journal of Economic Perspectives*, 15(3):47–68.
- Baily, M. N. (1978). Some aspects of optimal unemployment insurance. *Journal of Public Economics*, 10(3):379–402.
- Baker, S. R. (2014). Debt and the Consumption Response to Household Income Shocks. *working paper*.
- Baker, S. R. and Yannelis, C. (2015). Income Changes and Consumption: Evidence from the 2013 Federal Government Shutdown. SSRN Scholarly Paper ID 2575461, Social Science Research Network, Rochester, NY.
- Baugh, B., Ben-David, I., and Park, H. (2013). Disentangling Financial Constraints, Precautionary Savings, and Myopia: Household Behavior Surrounding Federal Tax Returns. *working paper*.
- Blundell, R., Pistaferri, L., and Preston, I. (2008). Consumption Inequality and Partial Insurance. *American Economic Review*, 98(5):1887–1921.
- Blundell, R., Pistaferri, L., and Saporta-Eksten, I. (2015). Consumption inequality and family labor supply. *American Economic Review*.
- Browning, M. and Crossley, T. (2001). Unemployment insurance benefit levels and consumption changes. *Journal of Public Economics*, 80(1):1–23.
- Browning, M., Crossley, T. F., and Winter, J. (2014). The Measurement of Household Consumption Expenditures. *Annual Review of Economics*, 6(1):475–501.
- Caballero, R. J. (1995). Near-Rationality, Heterogeneity, and Aggregate Consumption. *Journal of Money, Credit and Banking*, 27(1):29–48.
- Campbell, J. Y. and Mankiw, N. G. (1989). Consumption, Income and Interest Rates: Reinterpreting the Time Series Evidence. NBER Chapters, National Bureau of Economic Research, Inc.

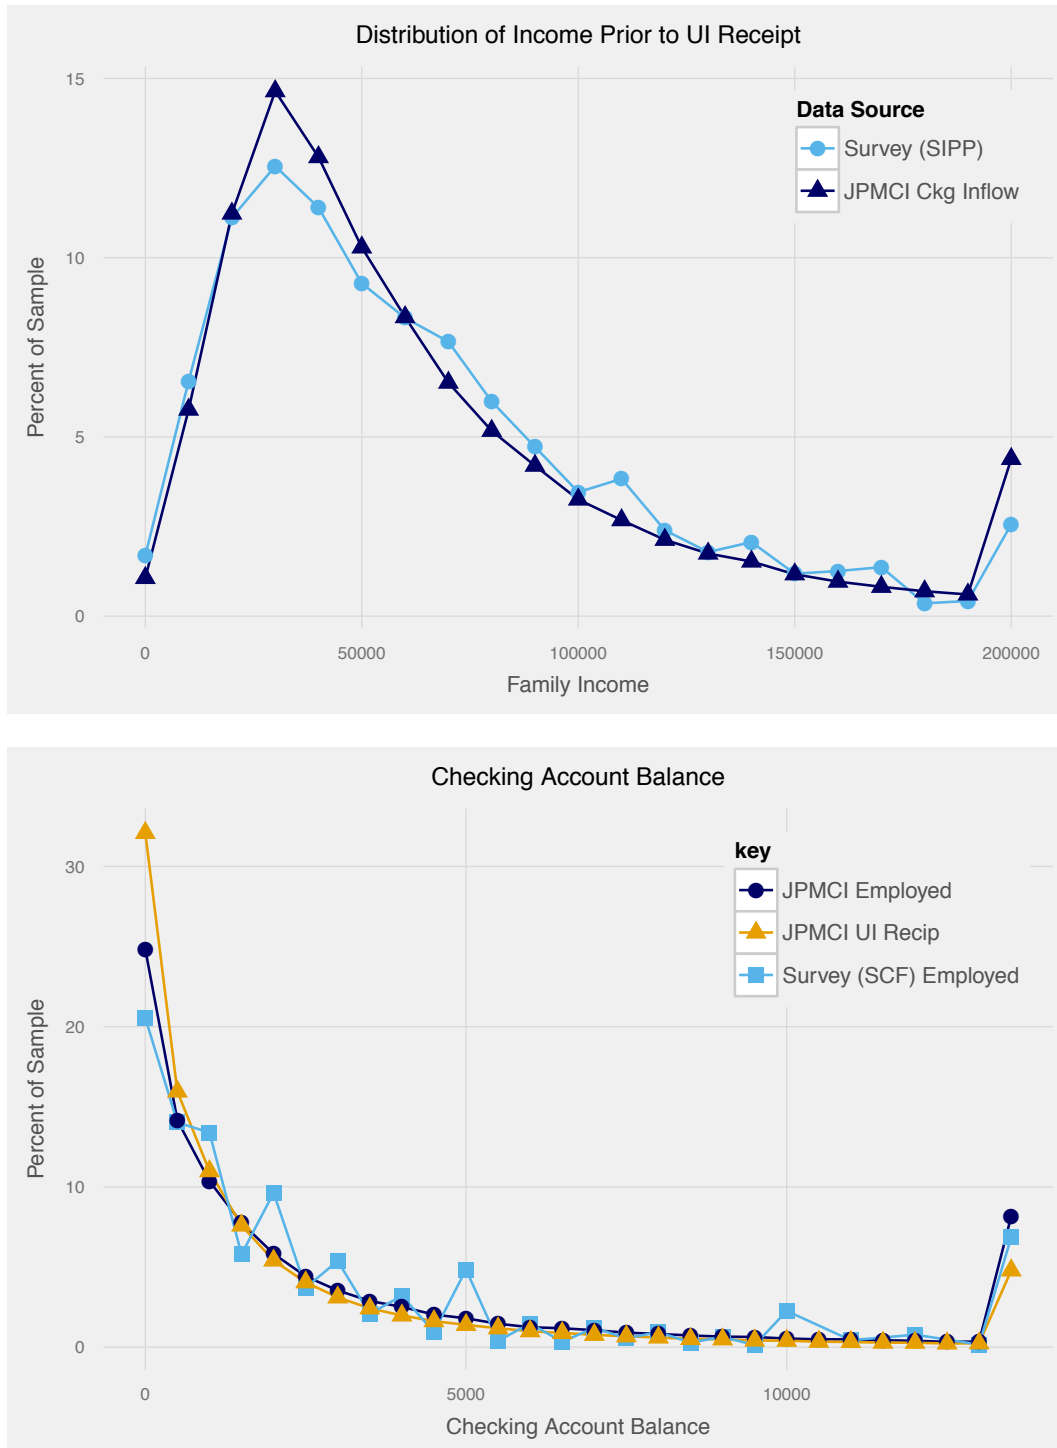
- Card, D., Chetty, R., and Weber, A. (2007). The Spike at Benefit Exhaustion: Leaving the Unemployment System or Starting a New Job? *American Economic Review*, 97(2):113–118.
- Carroll, C. D. (1997). Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis. *The Quarterly Journal of Economics*, 112(1):1–55.
- Carroll, C. D. (2006). The method of endogenous gridpoints for solving dynamic stochastic optimization problems. *Economics Letters*, 91(3):312–320.
- Chetty, R. (2006). A general formula for the optimal level of social insurance. *Journal of Public Economics*, 90(10-11):1879–1901.
- Chetty, R. (2008). Moral Hazard versus Liquidity and Optimal Unemployment Insurance. *Journal of Political Economy*, 116(2):173–234.
- Chetty, R. and Szeidl, A. (2007). Consumption Commitments and Risk Preferences. *The Quarterly Journal of Economics*, 122(2):831–877.
- Chodorow-Reich, G. and Karabarbounis, L. (2015). The Cyclicalities of the Opportunity Cost of Employment. *Journal of Political Economy*, forthcoming.
- Cochrane, J. H. (1989). The Sensitivity of Tests of the Intertemporal Allocation of Consumption to Near-Rational Alternatives. *American Economic Review*, 79(3):319–37.
- Cochrane, J. H. (1991). A Simple Test of Consumption Insurance. *Journal of Political Economy*, 99(5):957–76.
- Cohen, M. and Rysman, M. (2013). Payment Choice with Consumer Panel Data. *Federal Reserve Bank of Boston Working Paper*.
- Couch, K. and Placzek, D. (2010). Earnings Losses of Displaced Workers Revisited. *American Economic Review*, 100(1):572–589.
- Cullen, J. B. and Gruber, J. (2000). Does Unemployment Insurance Crowd Out Spousal Labor Supply? *Journal of Labor Economics*, 18(3):546–72.
- Davis, S. J. and von Wachter, T. (2011). Recessions and the Costs of Job Loss. *Brookings Papers on Economic Activity*, 43(2, Fall):1–72.
- Deaton, A. (1991). Saving and Liquidity Constraints. *Econometrica*, 59(5):1221–48.
- DellaVigna, S., Lindner, A., Reizer, B., and Schmieder, J. (2014). Reference-Dependent Job Search: Evidence from Hungary. *working paper*.
- DellaVigna, S. and Paserman, M. D. (2005). Job Search and Impatience. *Journal of Labor Economics*, 23(3):527–588.
- Farber, H. (2015). Job Loss in the Great Recession and its Aftermath: U.S. Evidence from the Displaced Workers Survey. *working paper*.
- Friedman, M. (1957). *A Theory of the Consumption Function*. National Bureau of Economic Research, Inc.

- Fuchs-Schuendeln, N. and Hassan, T. A. (2015). Natural Experiments in Macroeconomics. NBER Working Paper 21228, National Bureau of Economic Research, Inc.
- Gelman, M., Kariv, S., Shapiro, M., Silverman, D., and Tadelis, S. (2015). How Individuals Smooth Spending: Evidence from the 2013 Government Shutdown Using Account Data. *working paper*.
- Gourinchas, P.-O. and Parker, J. A. (2002). Consumption Over the Life Cycle. *Econometrica*, 70(1):47–89.
- Gruber, J. (1997). The Consumption Smoothing Benefits of Unemployment Insurance. *The American Economic Review*, 87(1):192–205.
- Hall, R. E. (1978). Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence. *Journal of Political Economy*, 86(6):971–87.
- Hannagan, A. and Morduch, J. (2015). Income Gains and Month-to-Month Income Volatility: Household evidence from the US Financial Diaries. *working paper*.
- Hendren, N. (2015). Knowledge of Future Job Loss and Implications for Unemployment Insurance. *working paper*.
- Herkenhoff, K. (2015). The Impact of Consumer Credit Access on Unemployment. *working paper*.
- Herkenhoff, K., Phillips, G., and Cohen-Cole, E. (2015). How Credit Constraints Impact Job Finding Rates, Sorting & Aggregate Output. *working paper*.
- Hsieh, C.-T. (2003). Do Consumers React to Anticipated Income Changes? Evidence from the Alaska Permanent Fund. *American Economic Review*, 93(1):397–405.
- Hurd, M. and Rohwedder, S. (2010). Effects of the Financial Crisis and Great Recession on American Households. *working paper*.
- Jacobson, L., LaLonde, R., and Sullivan, D. (1993a). *The Costs of Worker Dislocation*. Upjohn Institute for Employment Research., Kalamazoo, Michigan.
- Jacobson, L., LaLonde, R., and Sullivan, D. (1993b). Earnings Losses of Displaced Workers. *American Economic Review*, 83(4):685–709.
- Jappelli, T. and Pistaferri, L. (2010). The Consumption Response to Income Changes. *Annual Review of Economics*, 2(1):479–506.
- Jarosch, G. (2015). Searching for Job Security and the Consequences of Job Loss. *working paper*.
- Johnson, D. S., Parker, J. A., and Souleles, N. S. (2006). Household Expenditure and the Income Tax Rebates of 2001. *American Economic Review*, 96(5):1589–1610.
- Kaplan, G. and Menzio, G. (2015). Shopping Externalities and Self-Fulfilling Unemployment Fluctuations. *Journal of Political Economy*, forthcoming.
- Kaplan, G. and Violante, G. L. (2014). A Model of the Consumption Response to Fiscal Stimulus Payments. *Econometrica*, 82(4):1199–1239.

- Katz, L. F. and Meyer, B. D. (1990). Unemployment Insurance, Recall Expectations, and Unemployment Outcomes. *The Quarterly Journal of Economics*, 105(4):973–1002.
- Kekre, R. (2015). Unemployment Insurance in Macroeconomic Stabilization. *working paper*.
- Kolsrud, J., Landais, C., Nilsson, P., and Spinnewijn, J. (2015). The Optimal Timing of Unemployment Benefits: Theory and Evidence from Sweden. *working paper*.
- Kostøl, A. R. and Mogstad, M. (2015). Earnings, Disposable Income, and Consumption of Allowed and Rejected Disability Insurance Applicants. *American Economic Review*, 105(5):137–41.
- Krueger, A. B. and Mueller, A. (2010). Job search and unemployment insurance: New evidence from time use data. *Journal of Public Economics*, 94(3-4):298–307.
- Krueger, A. B. and Mueller, A. (2014). A Contribution to the Empirics of Reservation Wages. *working paper*.
- Kuchler, T. (2014). Sticking to Your Plan: Hyperbolic Discounting and Credit Card Debt Paydown. *working paper*.
- Kueng, L. (2015). Explaining Consumption Excess Sensitivity with Near-Rationality: Evidence from Large Predetermined Payments. *working paper*.
- Laibson, D. (1997). Golden Eggs and Hyperbolic Discounting. *The Quarterly Journal of Economics*, 112(2):443–77.
- Laibson, D., Repetto, A., and Tobacman, J. (2015). Estimating Discount Functions with Consumption Choices over the Lifecycle. *working paper*.
- Mankiw, N. G. (1985). Small Menu Costs and Large Business Cycles: A Macroeconomic Model. *The Quarterly Journal of Economics*, 100(2):529–38.
- Modigliani, F. and Brumberg, R. (1954). Utility analysis and the consumption function: an interpretation of cross-section data. In Kurihara, K., editor, *Post- Keynesian Economics*, pages 388–436. Rutgers University Press, New Brunswick, NJ.
- Mortensen, D. T. (1977). Unemployment insurance and job search decisions. *Industrial and Labor Relations Review*, 30(4):505–517.
- of Labor Statistics, B. (2014). Employment Characteristics of Families. Technical report, Bureau of Labor Statistics.
- Parker, J. A., Souleles, N. S., Johnson, D. S., and McClelland, R. (2013). Consumer Spending and the Economic Stimulus Payments of 2008. *American Economic Review*, 103(6):2530–53.
- Passero, B., Blair, C., Cage, R., Garner, T., and McCully, C. (2011). The Devil’s in the Details —A Concordance for the CE, CPI, and PCE.
- Reis, R. (2006). Inattentive consumers. *Journal of Monetary Economics*, 53(8):1761–1800.
- Rothstein, J. (2011). Unemployment Insurance and Job Search in the Great Recession. *Brookings Papers on Economic Activity*, 43(2 (Fall)):143–213.

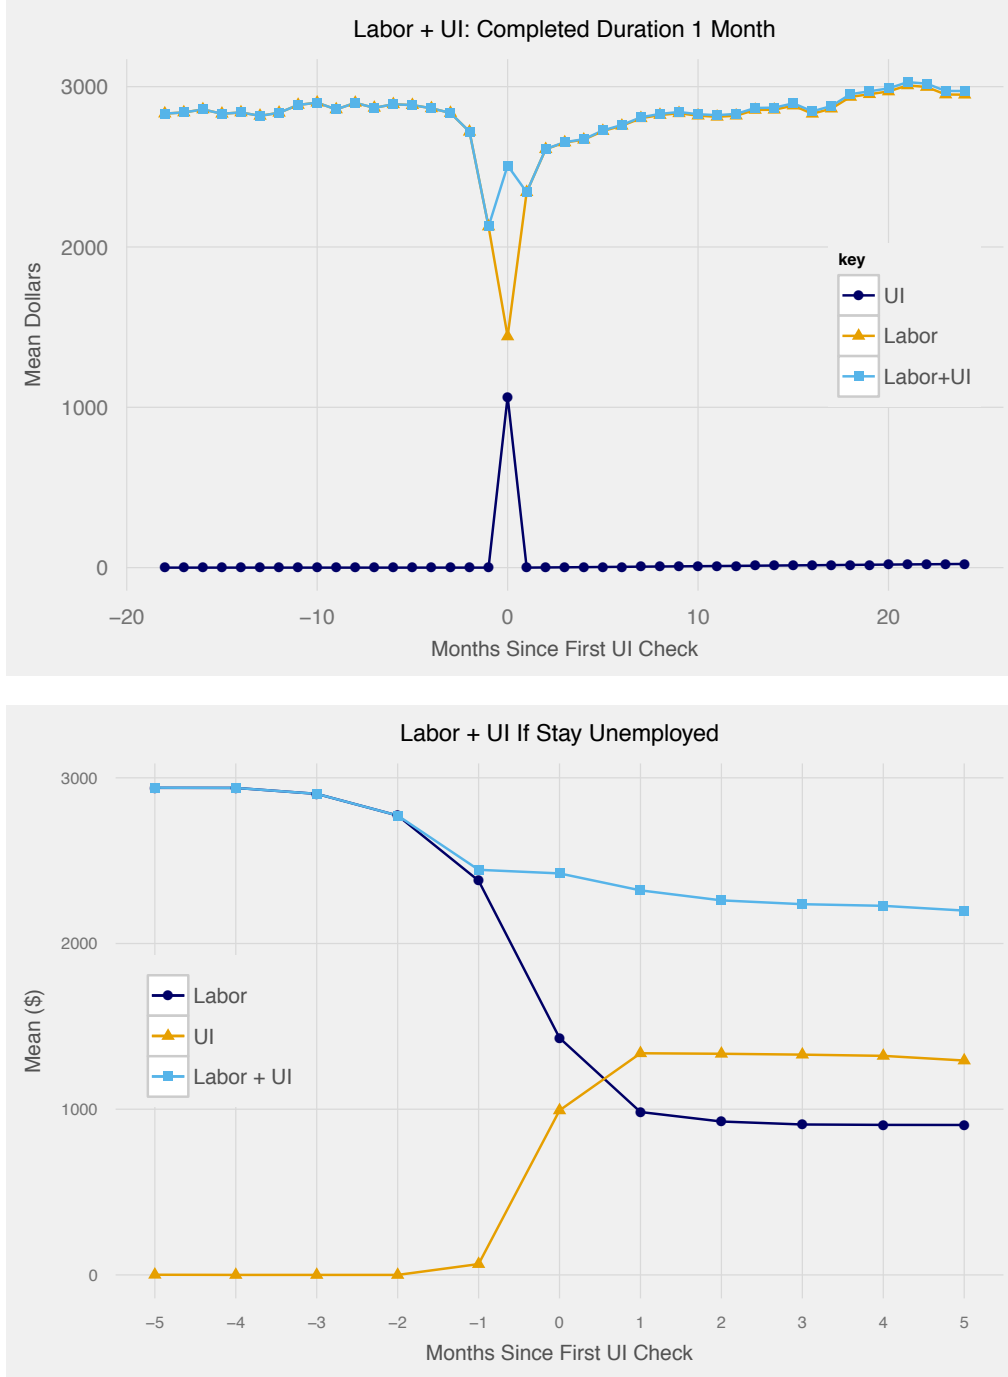
- Rothstein, J. and Valetta, R. (2014). Scraping By: Income and Program Participation After the Loss of Extended Unemployment Benefits. *working paper*.
- Saporta-Eksten, I. (2014). Job Loss, Consumption and Unemployment Insurance. *working paper*.
- Saunders, L. and McLaughlin, J. (2013). Survey of Unemployment Prepaid Cards. Technical report, National Consumer Law Center.
- Schmieder, J. F., Wachter, T. v., and Bender, S. (2012). The Effects of Extended Unemployment Insurance Over The Business Cycle: Evidence From Regression Discontinuity Estimates Over 20 Years. *The Quarterly Journal of Economics*, page qjs010.
- Shapiro, M. D. and Slemrod, J. (2009). Did the 2008 Tax Rebates Stimulate Spending? *American Economic Review*, 99(2):374–79.
- Shimer, R. and Werning, I. (2007). Reservation Wages and Unemployment Insurance. *The Quarterly Journal of Economics*, 122(3):1145–1185.
- Shimer, R. and Werning, I. (2008). Liquidity and Insurance for the Unemployed. *The American Economic Review*, 98(5):1922–1942.
- Shy, O. (2013). How Many Cards Do You Use? *Federal Reserve Bank of Boston Working Paper*.
- Souleles, N. S. (1999). The Response of Household Consumption to Income Tax Refunds. *American Economic Review*, 89(4):947–958.
- Spinnewijn, J. (2015). Unemployed But Optimistic: Optimal Insurance Design With Biased Beliefs. *Journal of the European Economic Association*, 13(1):130–167.
- Stephens, M. (2001). The Long-Run Consumption Effects of Earnings Shocks. *Review of Economics and Statistics*, 83(1):28–36.
- Stephens, M. (2002). Worker Displacement and the Added Worker Effect. *Journal of Labor Economics*, 20(3):504–537.
- System, F. R. (2013). The 2013 Federal Reserve Payments Study. Technical report.
- von Wachter, T., Schmieder, J., and Bender, S. (2015). The Effect of Unemployment Insurance and Nonemployment Duration on Wages. *American Economic Review*, forthcoming.
- Wachter, T. v., Song, J., and Manchester, J. (2009). Long-Term Earnings Losses due to Mass Layoffs During the 1982 Recession: An Analysis Using U.S. Administrative Data from 1974 to 2004. *working paper*.
- Welander, T. (2014). Trends in Consumer Payments and Retail Banking: Report 1 of 4. Technical report, GC Insights Marketing Research Services.

FIGURE 1 – REPRESENTATIVENESS: INCOME AND ASSET DISTRIBUTION



Notes: The top panel plots the distribution of pre-tax family income in the year prior to UI receipt in the 2004 Survey of Income and Program Participation and in the JPMCI data. The bottom panel plots the distribution of checking account balances for employed families in the 2013 Survey of Consumer Finances, employed families in the JPMCI data, and families three months before UI receipt in the JPMCI data. See Sections 2.2 and 2.4 for details.

FIGURE 2 – EVENT STUDY: INCOME AT UI ONSET



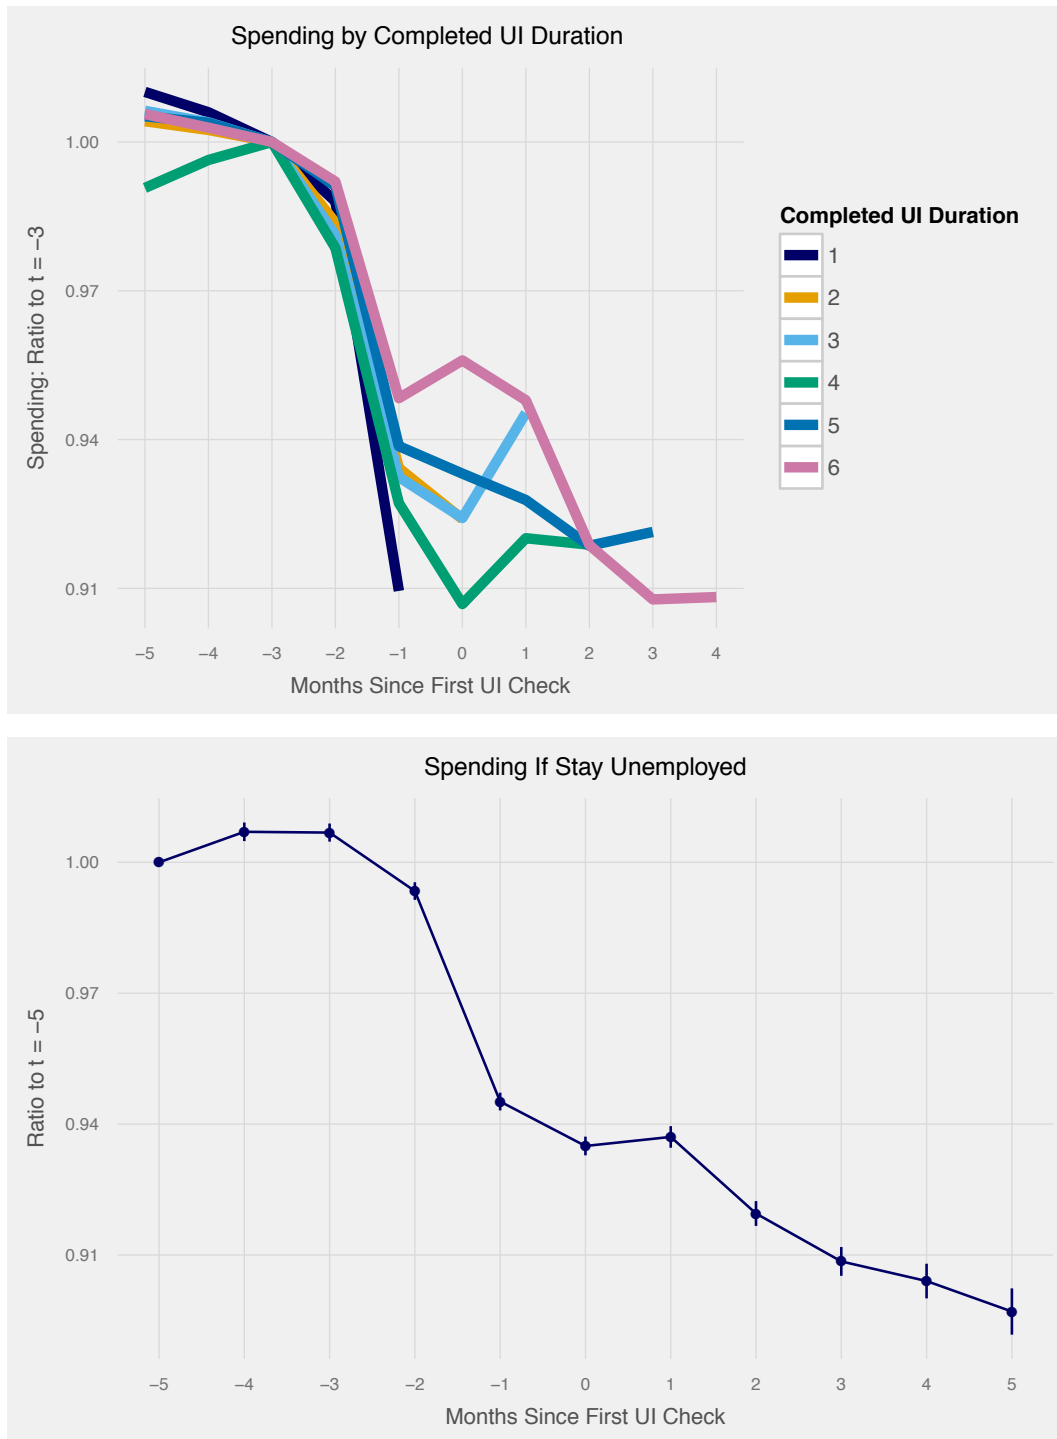
Notes: The top panel shows the path of labor income for a family that receives UI benefits in exactly one month. Direct deposit labor income declines in the three months leading up to UI receipt. The bottom panel plots average labor and UI income for the sample of agents who stay unemployed. In months $t = \{-5, -4, -3, -2, -1, 0\}$, this includes everyone who receives UI at date 0. In month $t = 1$, this includes only families who continue to receive UI and excludes families who received their last UI check in month 0. In month $t = 2$, this excludes families who received their last UI check in month 0 or month 1, and so on. Mean labor income is positive during UI receipt because sometimes other family members continue to receive labor income. See Section 3.1 for details. These estimates are relative to a control group described in Section 2.5.

FIGURE 3 – EVENT STUDY: SPENDING AT UI ONSET



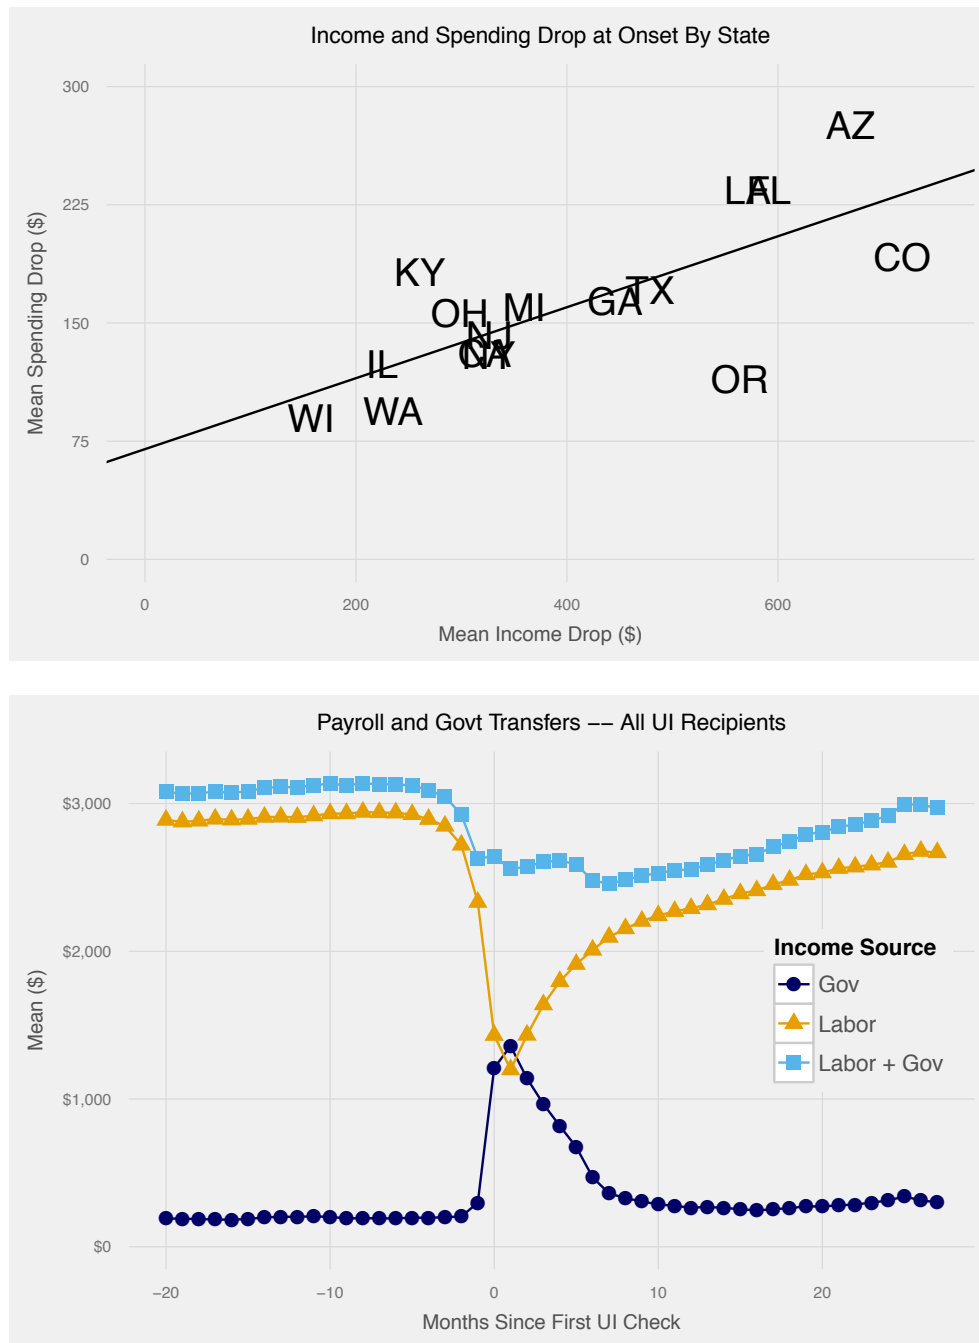
Notes: The top-left panel shows the path of spending for a family that received UI benefits in exactly one month. The gray dashed vertical line indicates the last month in which UI benefits were received. The subsequent panels plot the path of spending for families that received UI for 2, 3, 4, and 5 months. The last panel plots spending for families that received UI for 6 months and exhausted benefits. These estimates are relative to a control group described in Section 2.5.

FIGURE 4 – SPENDING IF STAY UNEMPLOYED



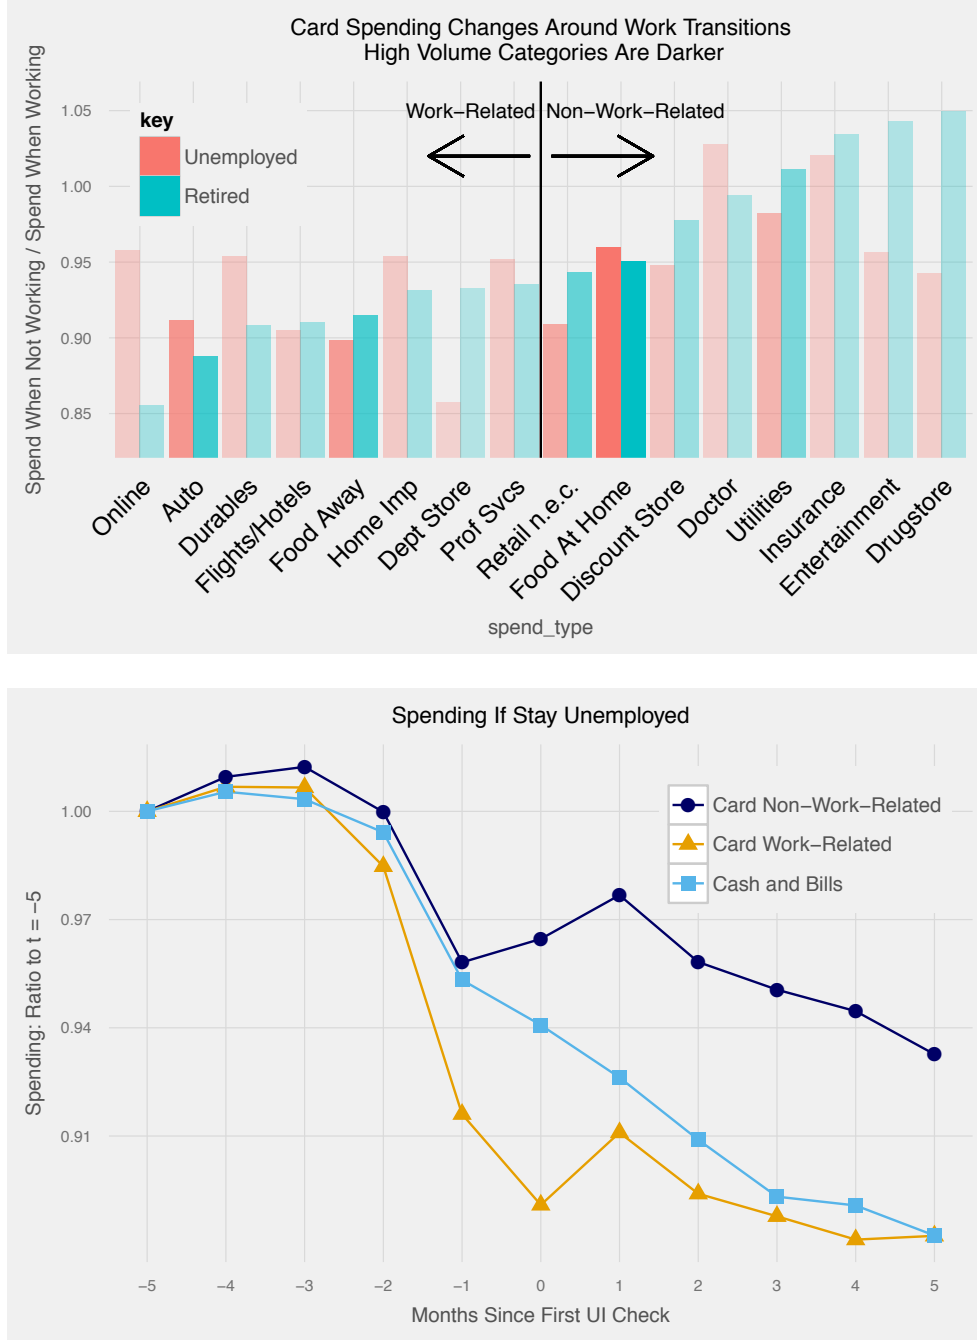
Note: The top panel plots the same spending series as Figure 3, zoomed in from five months prior to UI onset until one before UI benefits are terminated. The bottom panel shows the composite path of spending for families who remain unemployed using the data in the top panel using the same methodology as in Figure 2. The vertical bars are 95% confidence intervals for the change from the prior month. Section 3.1 describes the methodology for building this plot. These estimates are relative to a control group described in Section 2.5.

FIGURE 5 – INTERPRETING ONSET: TEMPORARY INCOME LOSS,
PERMANENT INCOME LOSS



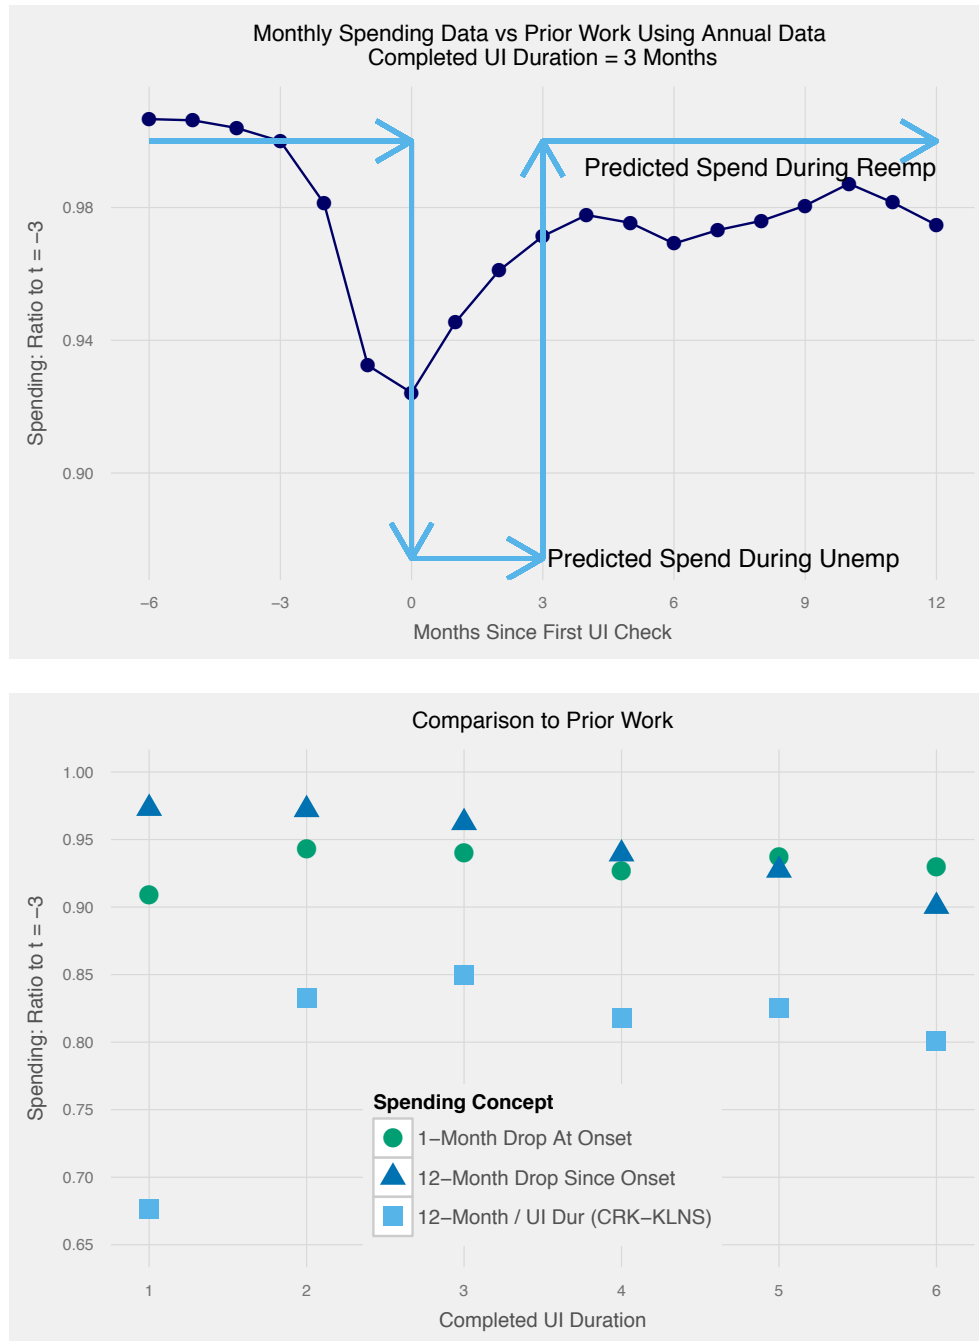
Note: The top panel plots the change in income and the change in spending at onset for the sixteen largest states in the JPMCI sample. States where families have a bigger drop in income at onset also have a bigger drop in spending at onset. The bottom panel plots the change in labor income and government transfers (UI, SSA, DI and tax refunds) for all UI recipients, relative to the first month in which they received a UI check. Transfers fall and labor income rises each month as people find employment. These estimates are relative to a control group described in Section 2.5.

FIGURE 6 – INTERPRETING ONSET: WORK-RELATED EXPENSES



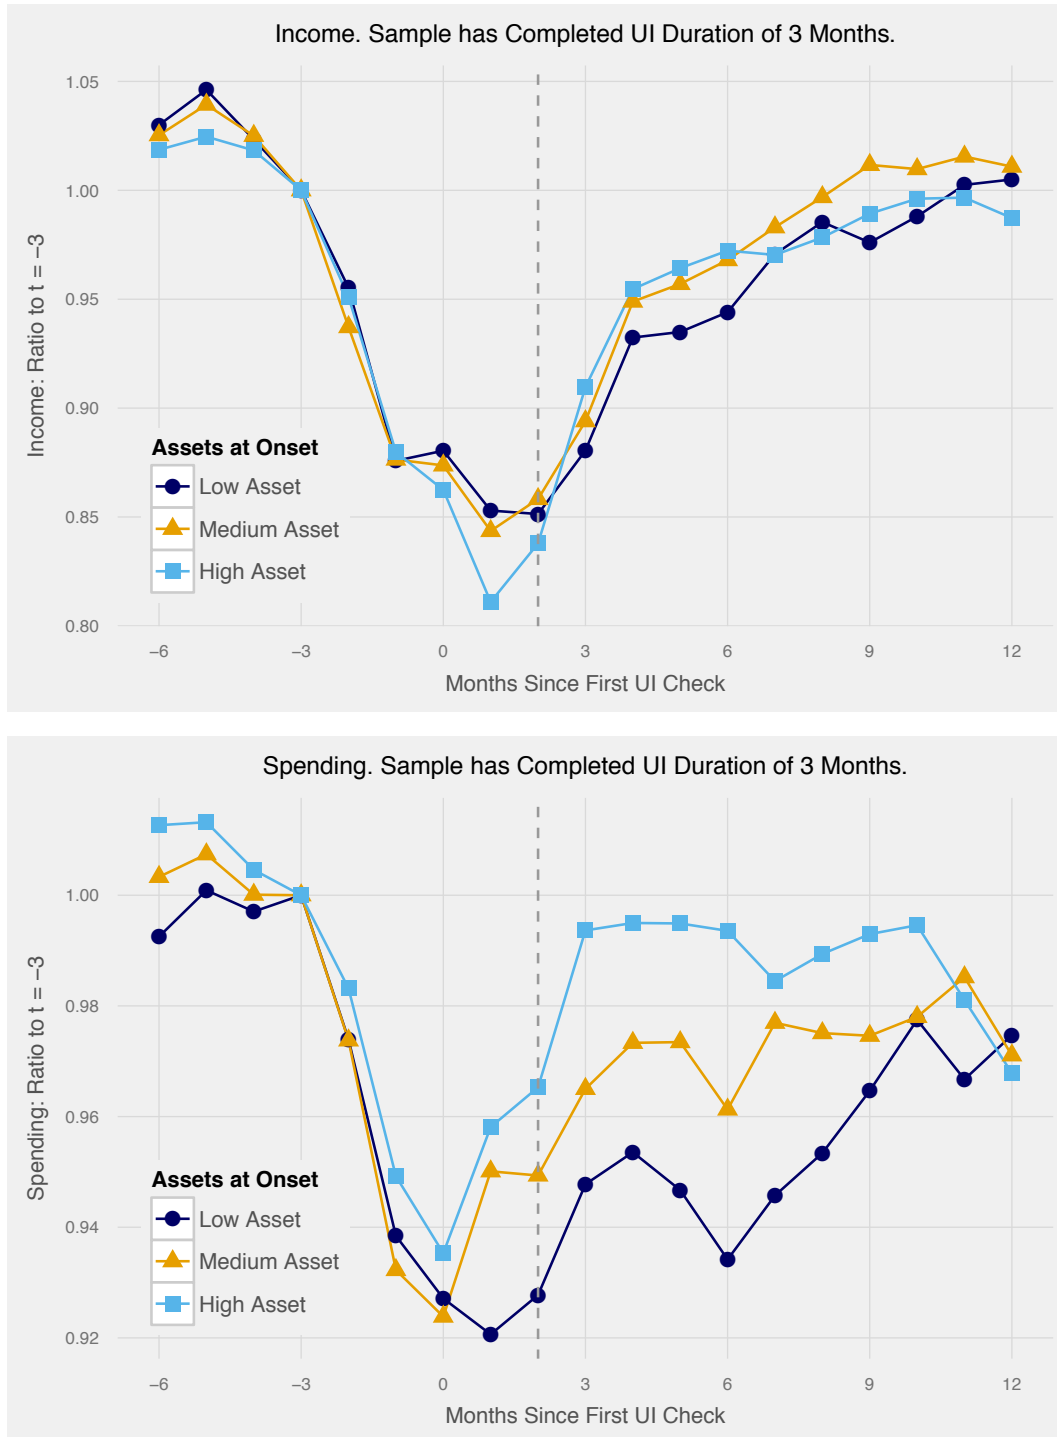
Note: The top panel compares the change in spending at retirement to the change in spending at the onset of unemployment for debit and credit card expenditures in 16 different merchant groups. Darker bars indicate larger spending categories. We classify expenditure groups with drops greater than 6% at retirement (to the left of the vertical line) as “work related.” The bottom panel re-constructs the composite spending series while unemployed from Figure 4 separately for card work-related expenditures (29% of pre-onset spending), card non-work-related expenditures (33%) and cash withdrawals and bills (38%). In Section 3.2.3, we estimate that 22-31% of the drop in spending at onset is attributable to the excess drop in work-related expenditures. These estimates are relative to a control group described in Section 2.5.

FIGURE 7 – SPENDING DROP DURING UNEMPLOYMENT:
COMPARISON TO PRIOR WORK



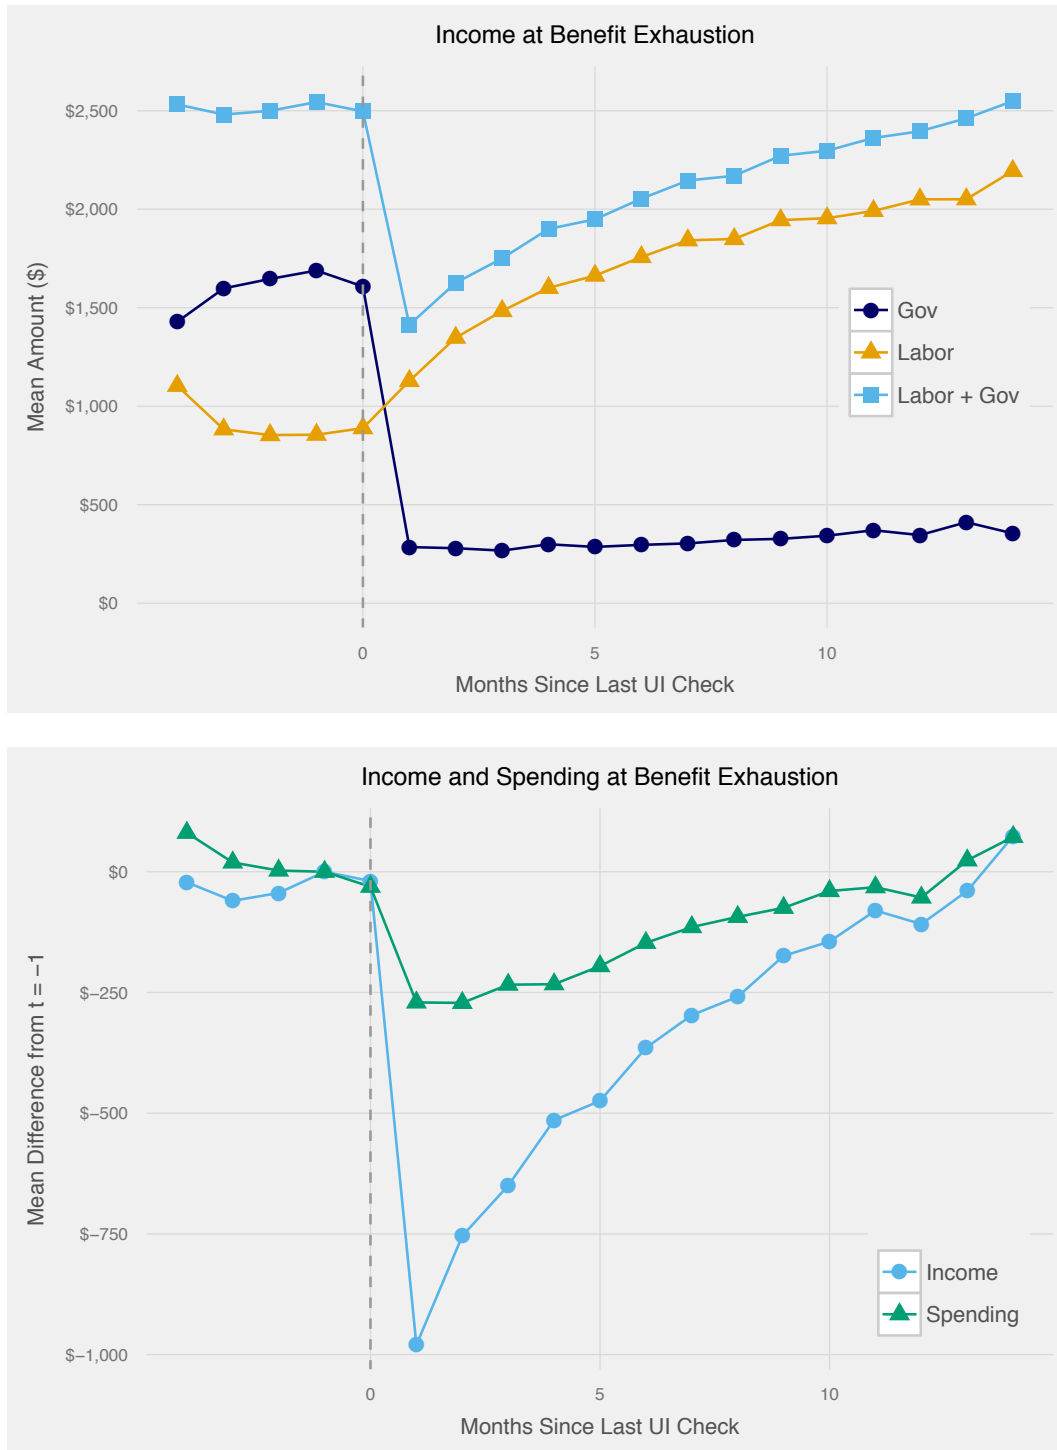
Note: The top panel plots the path of spending for families that received UI for exactly three months with navy blue circles. Chodorow-Reich and Karabarbounis (2015) (CRK) and Kolsrud et al. (2015) (KLNS) analyze annual spending data and assume spending drops *only* during unemployment. The light blue arrows depict their calculation methodology applied to the data. This overstates the true drop in spending because families engage in smoothing from month to month. The bottom panel shows: (1) the annual drop in spending in the 12 months following onset in orange triangles, (2) the calculated drop in spending during unemployment using the CRK-KLNS methodology in light blue squares and (3) the monthly drop in spending at onset in blue circles.

FIGURE 8 – EVENT STUDY FOR 3-MONTH COMPLETED UI SPELLS:
HETEROGENEITY BY ASSET HOLDINGS



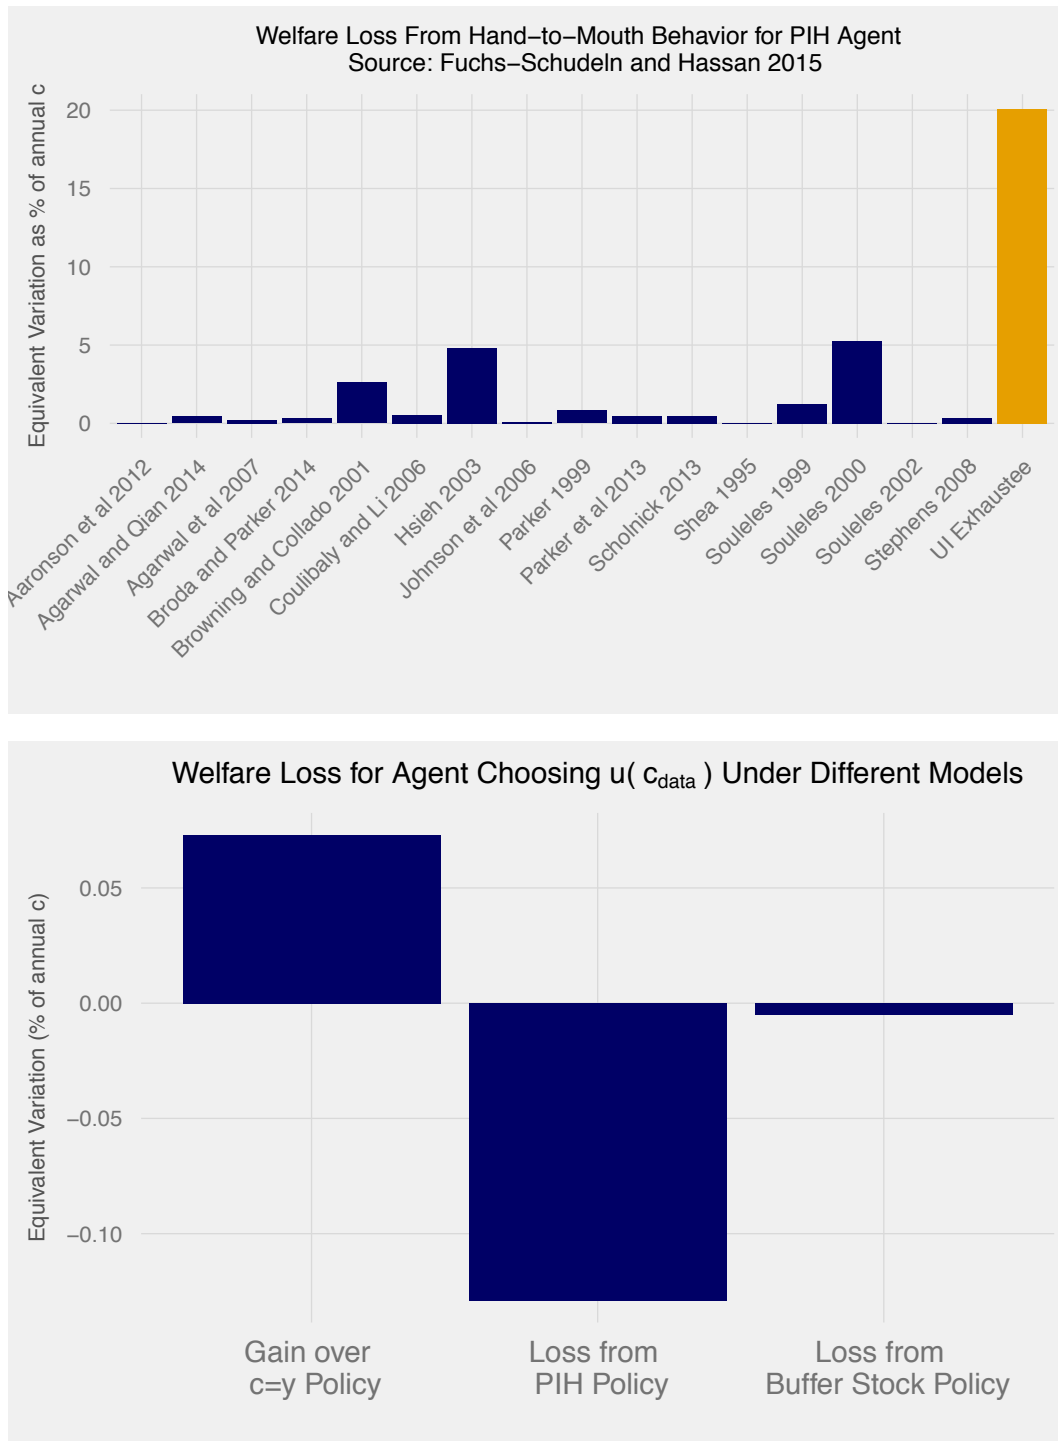
Note: The top panel plots the path of income for families with completed UI spells of three months, stratified by asset terciles. The vertical dashed gray line indicates the last month in which UI benefits were received. Families in these three groups with completed UI spells of three months have relatively similar paths of income. The bottom panel plots the path of spending by asset group. Families with little assets at onset have a much slower recovery in spending. These estimates are relative to a control group described in Section 2.5.

FIGURE 9 – UI BENEFIT EXHAUSTION



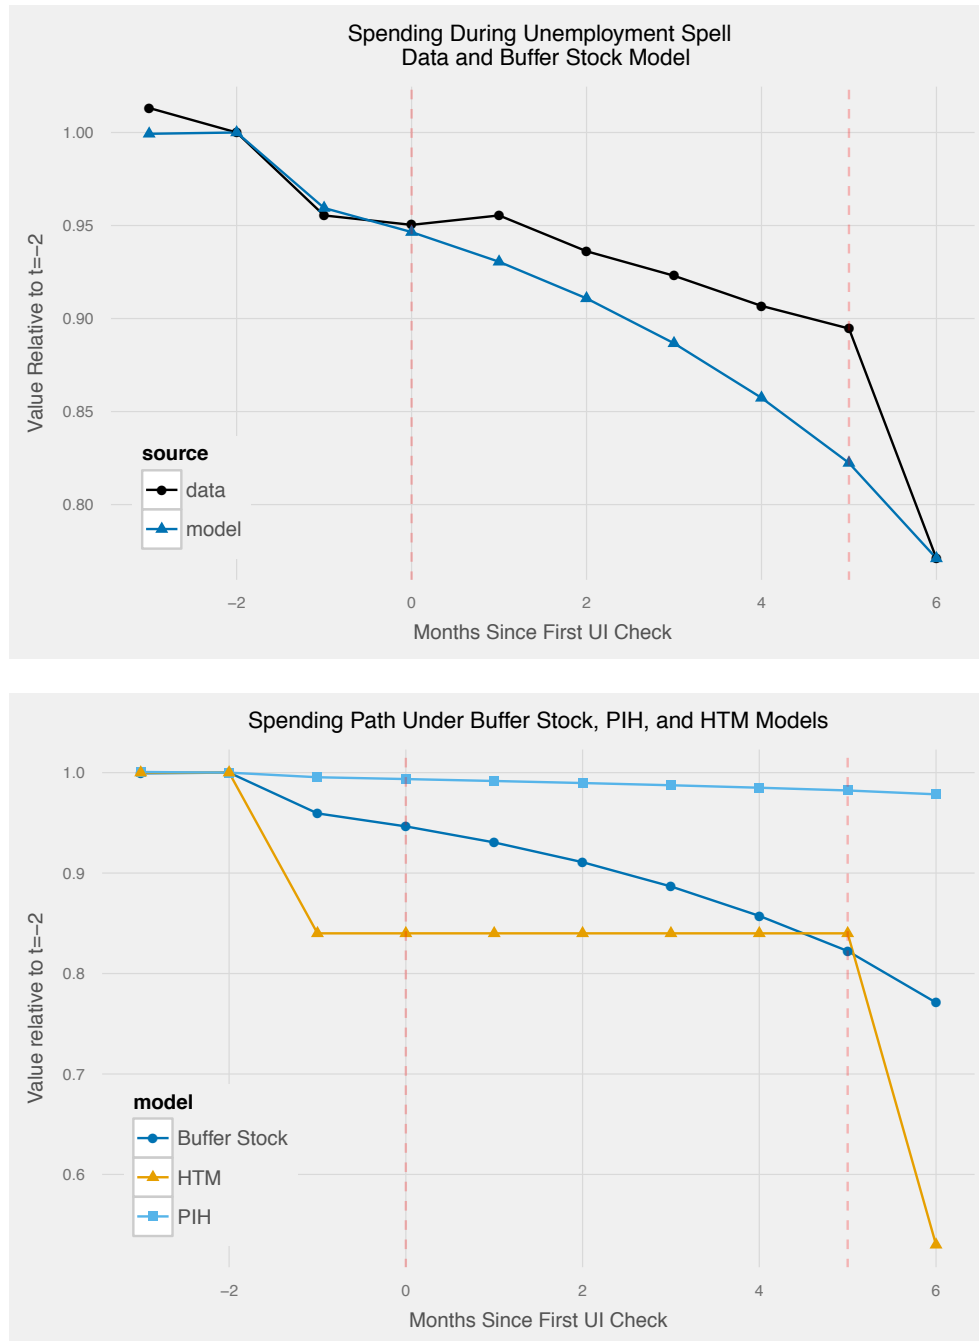
Notes: The top panel plots UI benefits and labor income relative to benefit exhaustion. The bottom panel plots the change in income (labor income plus government transfers) and spending around benefit exhaustion. See Section 5 for details. These estimates are relative to a control group described in Section 2.5.

FIGURE 10 – WELFARE LOSSES BY MODEL



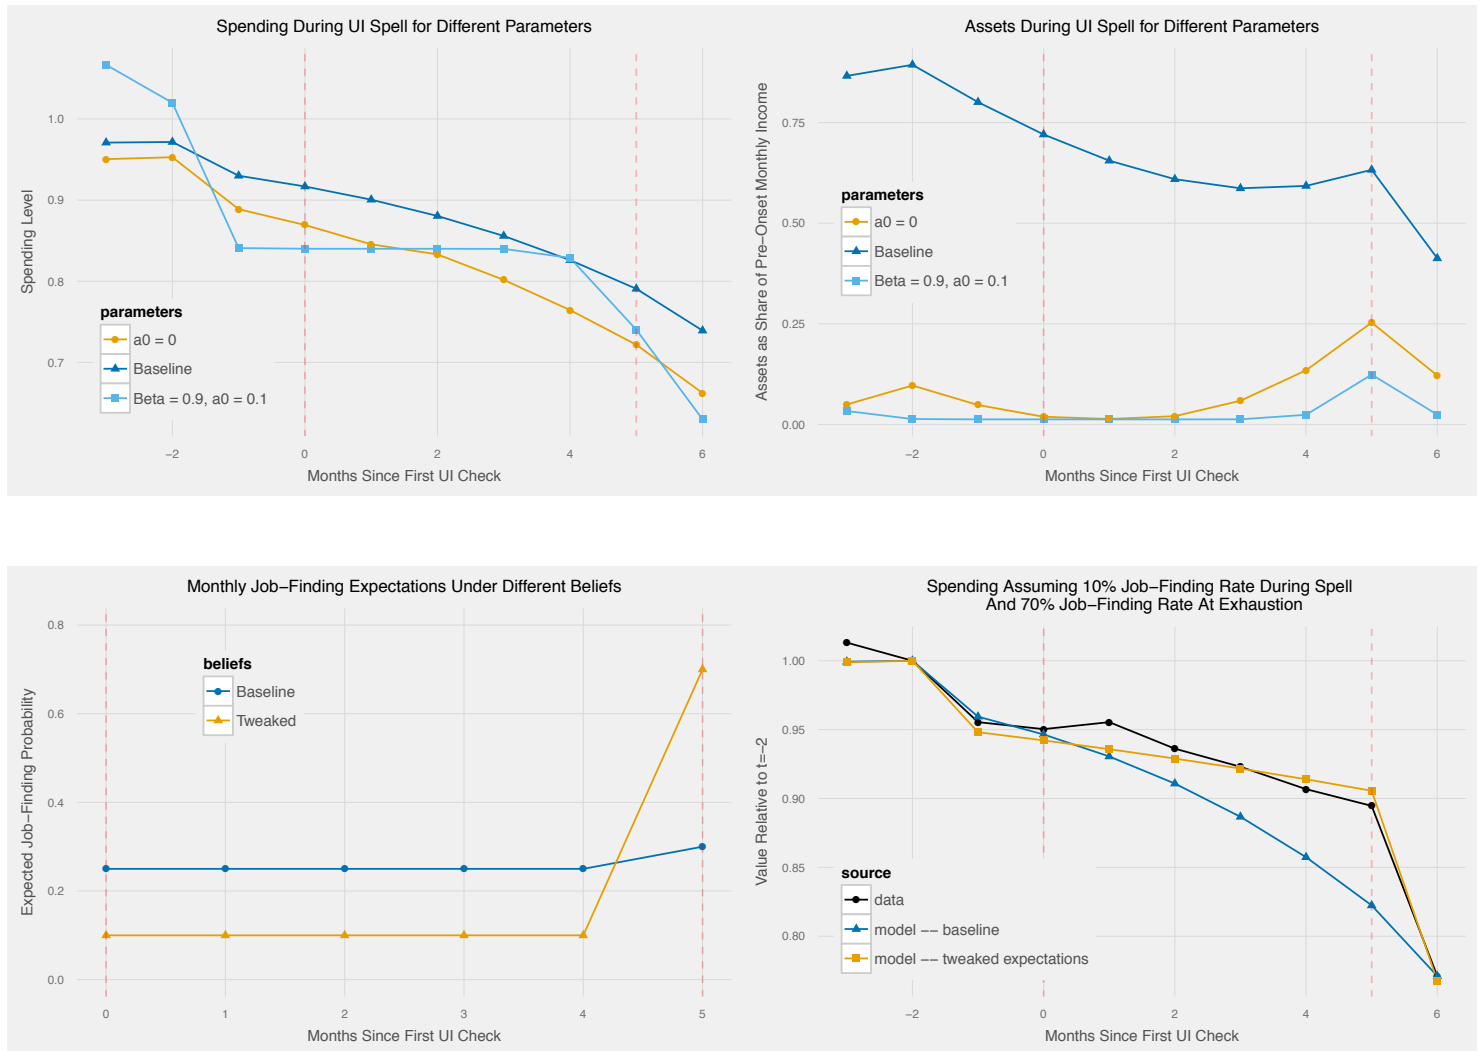
Note: The blue bars in the top panel show the welfare loss of failing to smooth consumption out of a temporary income change in 18 studies. Fuchs-Schuendeln and Hassan (2015) calculate this as z that solves $12u(y + \frac{x}{12}) = 11u(y + z) + u(y + x + z)$. The yellow bar is our calculation of a comparable statistic for the income change associated with an unemployment spell of at least six months. The bottom panel shows the welfare gain or loss associated with consumption paths predicted by the hand-to-mouth, permanent-income-hypothesis, and buffer stock models relative to the spending path observed in the data. See Section 6.1 for details.

FIGURE 11 – SPENDING IF STAY UNEMPLOYED – MODELS VS. DATA



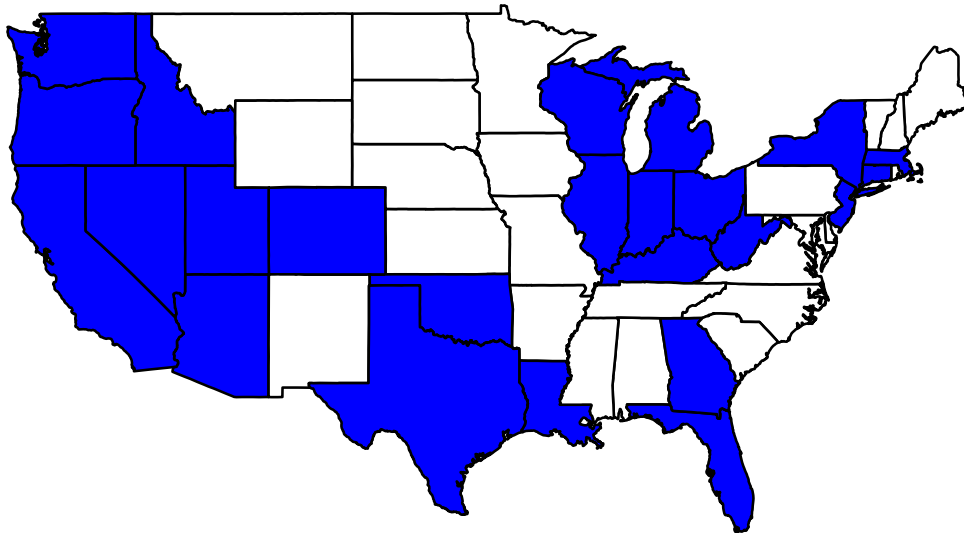
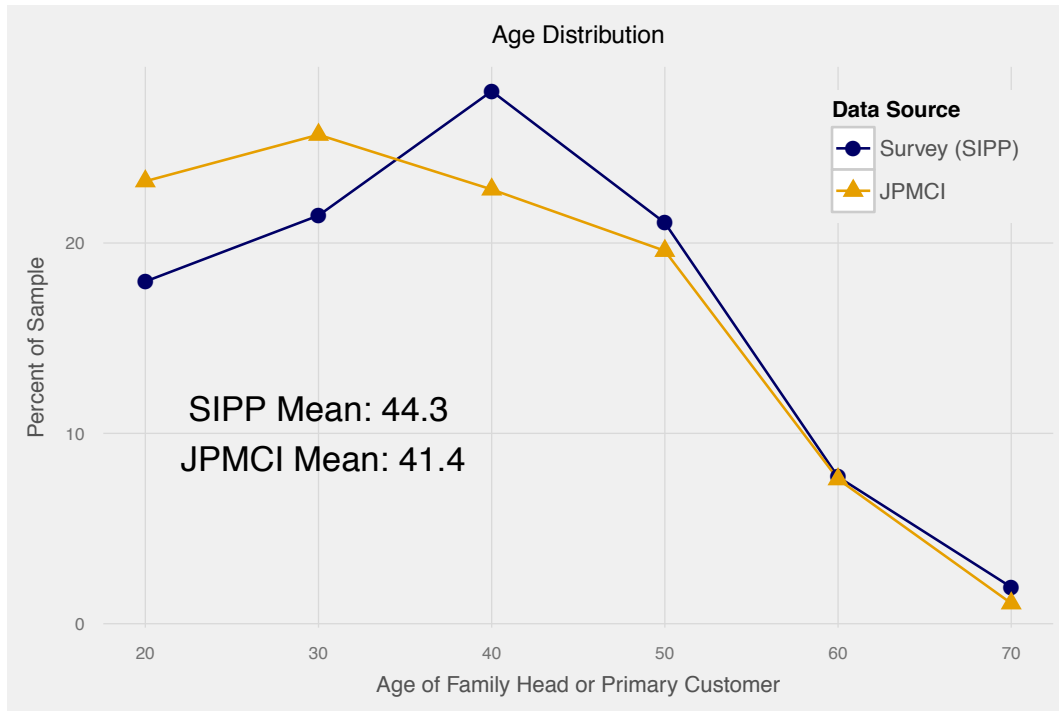
Note: The top panel plots the path of spending predicted by the buffer-stock model against the path of spending observed in the data for families that exhaust UI benefits. See Section 6.3 for details. The bottom panel plots the path of spending predicted by the buffer-stock, permanent income hypothesis (PIH), and hand-to-mouth (HTM) models described in the text.

FIGURE 12 – MATCHING SPENDING DROP AT EXHAUSTION



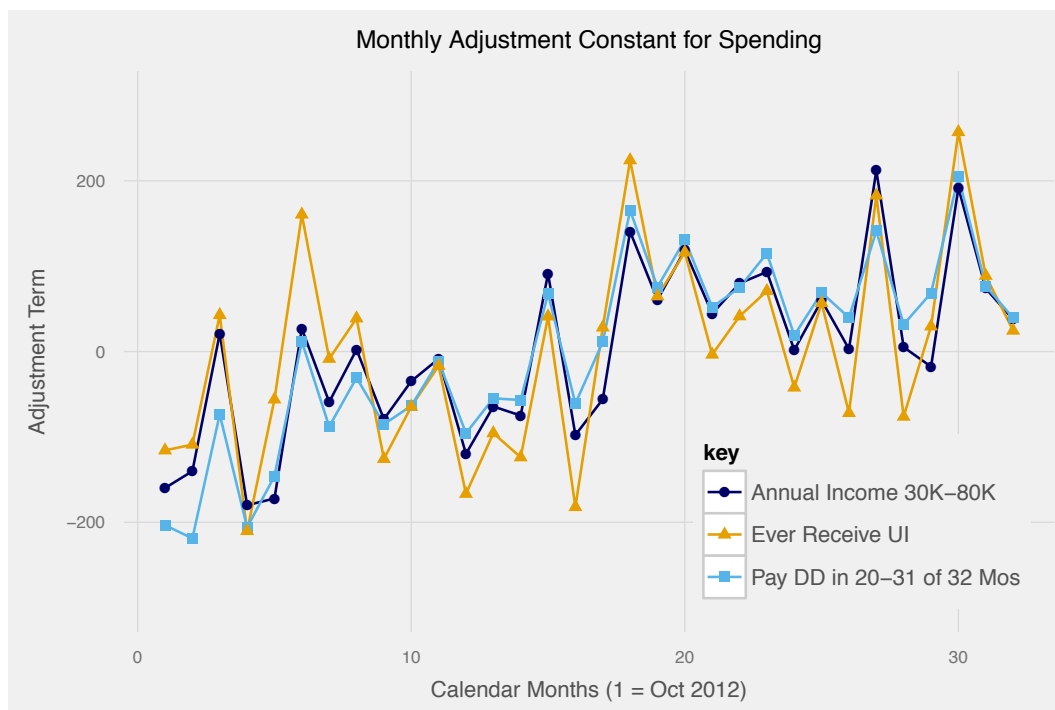
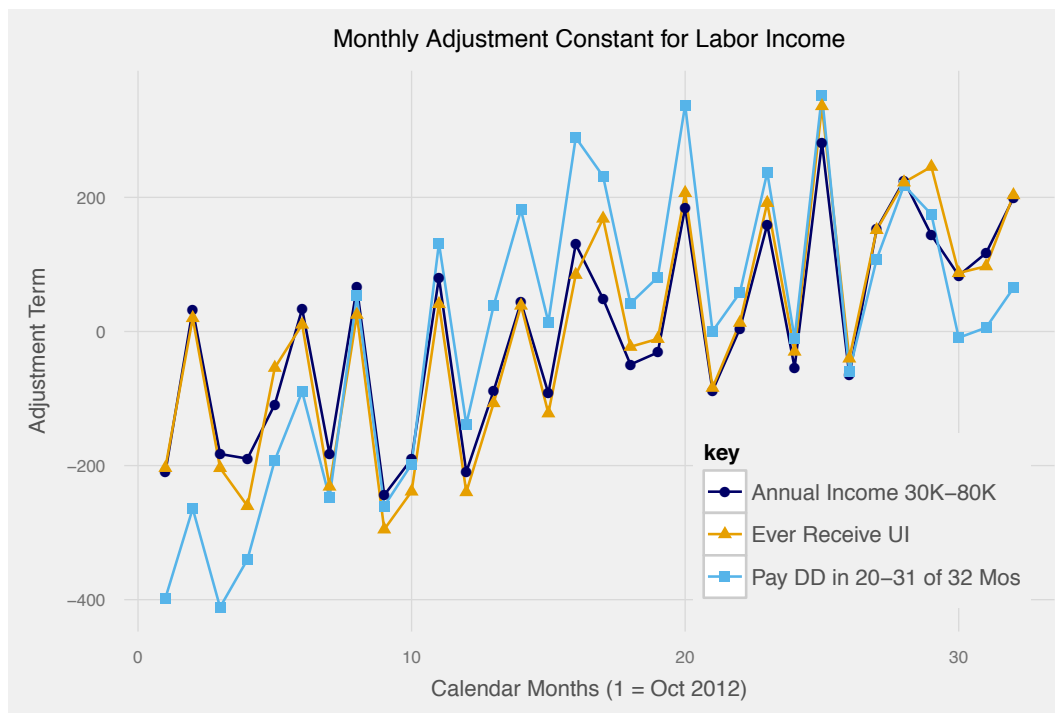
Note: The top left panel plots the path of spending predicted by the buffer-stock model under different parameter assumptions: (1) the baseline set of assumptions, (2) an alternative where initial assets are set to zero, and (3) an alternative where initial assets are set to zero and the monthly discount factor is 0.9. The top right panel plots the path of assets under the same three sets of assumptions. The bottom left panel plots monthly job-finding expectations under the baseline assumptions (which match the data), and tweaked assumptions where agents believe their monthly job-finding probability is 10% in the first five months of unemployment, and jumps to 70% in the final month of benefits. The bottom right panel shows the path of spending in the data, the model under baseline job-finding beliefs, and the model under the tweaked job-finding beliefs plotted in the previous panel.

APPENDIX FIGURE 1 – REPRESENTATIVENESS: AGE AND GEOGRAPHY



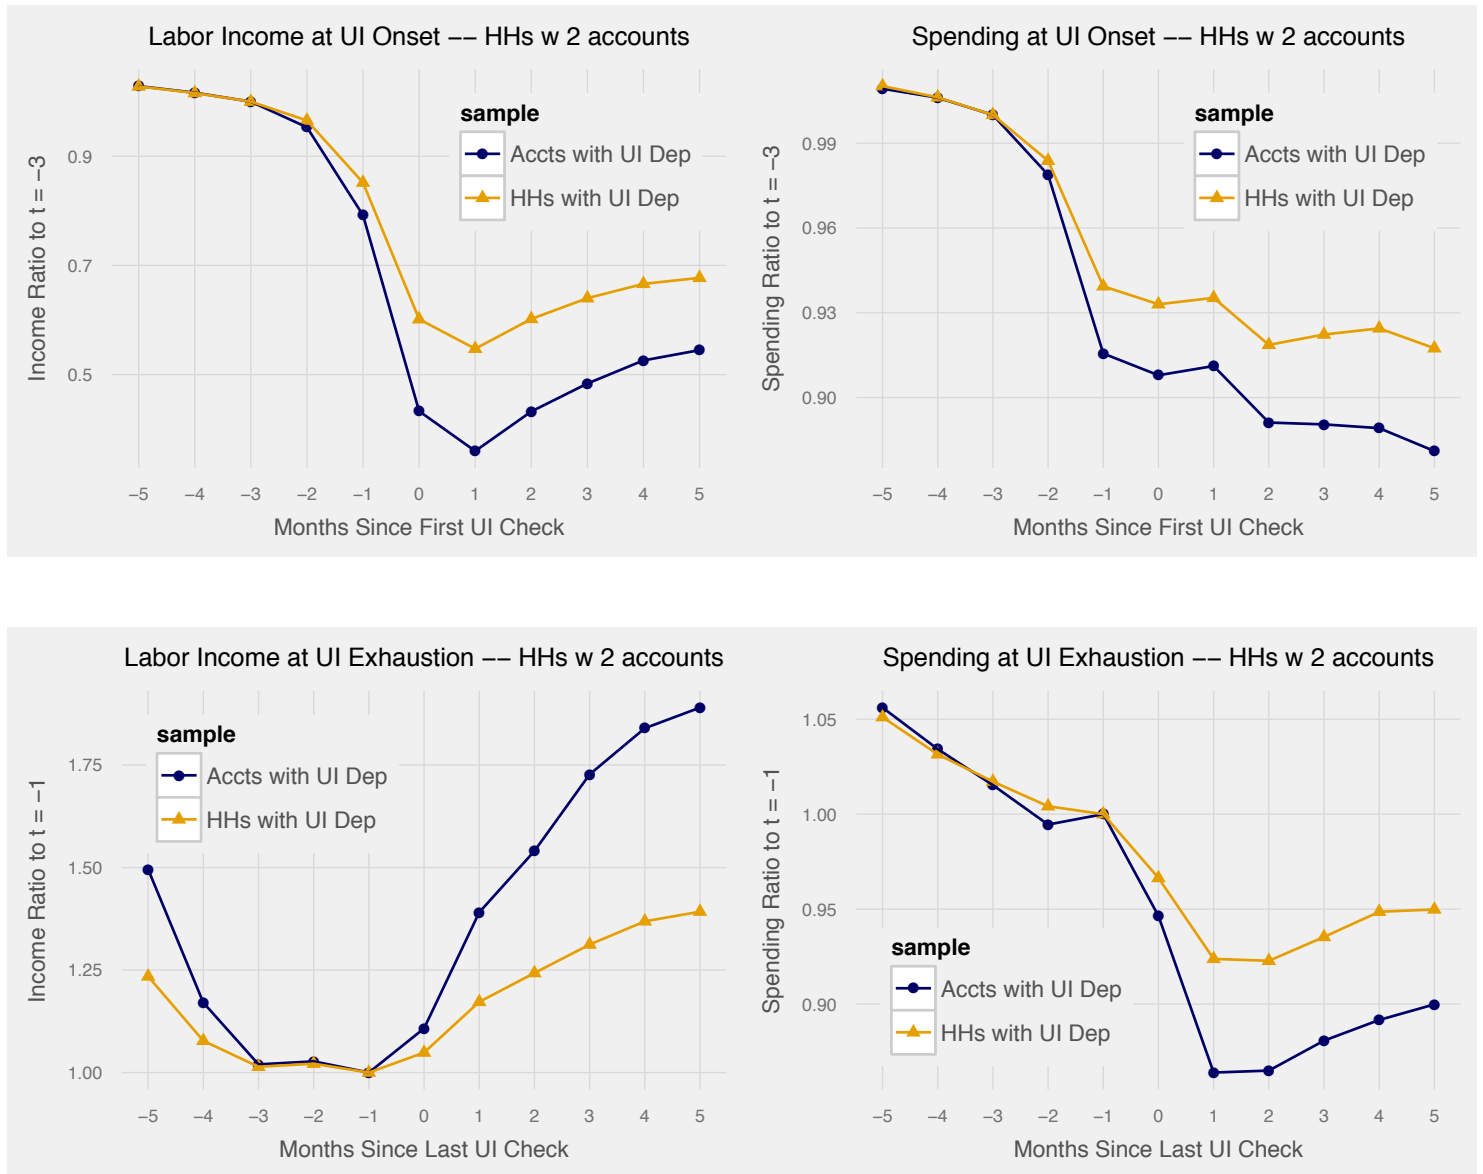
Notes: The top panel plots the age of family head for UI recipients in the Survey of Income and Program Participation and in the bank data. The bottom panel shows the states in which the bank has a physical footprint based on ATM locations publicly posted on Chase.com.

APPENDIX FIGURE 2 – CALENDAR MONTH ADJUSTMENT FACTORS



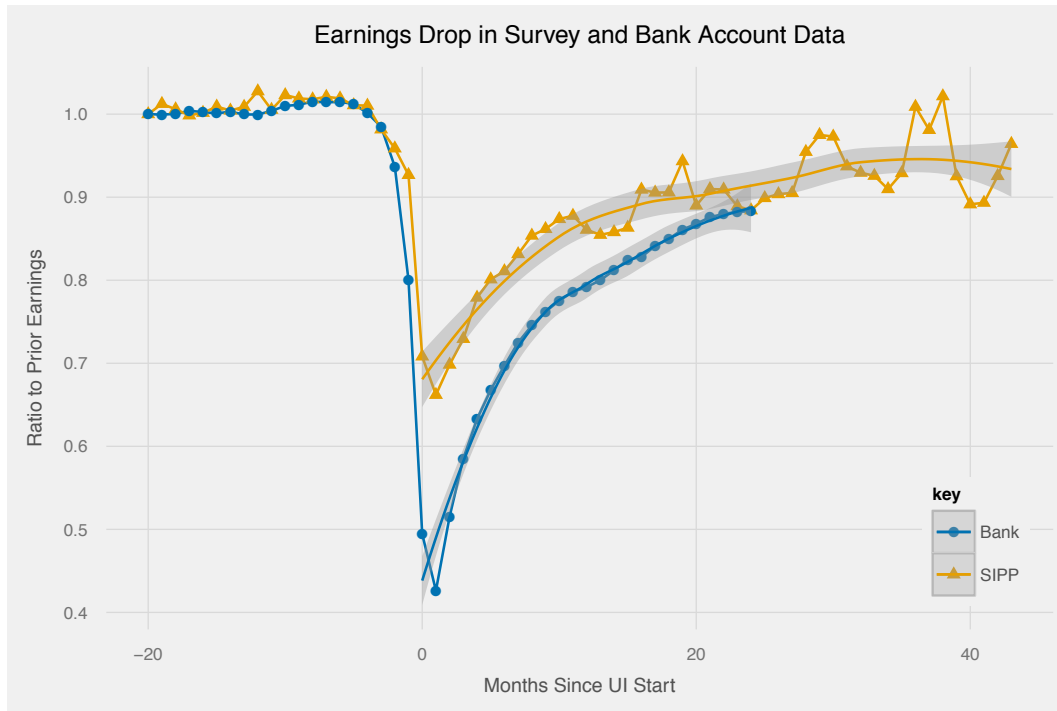
Note: To eliminate seasonality, inflation, and upward secular trends related to increased use of electronic payment methods, all results for income and spending are presented relative to a comparison group. This figure shows the monthly dollar adjustments associated with three different comparison groups: families which (1) received UI in at least one month, (2) received direct deposit payroll in 21-31 of 32 months and (3) had third-party annual income estimates between \$30,000 and \$80,000. See Section 2.5 for details.

APPENDIX FIGURE 3 – MEASURING FAMILY-WIDE SPENDING WITH UNLINKED ACCOUNTS



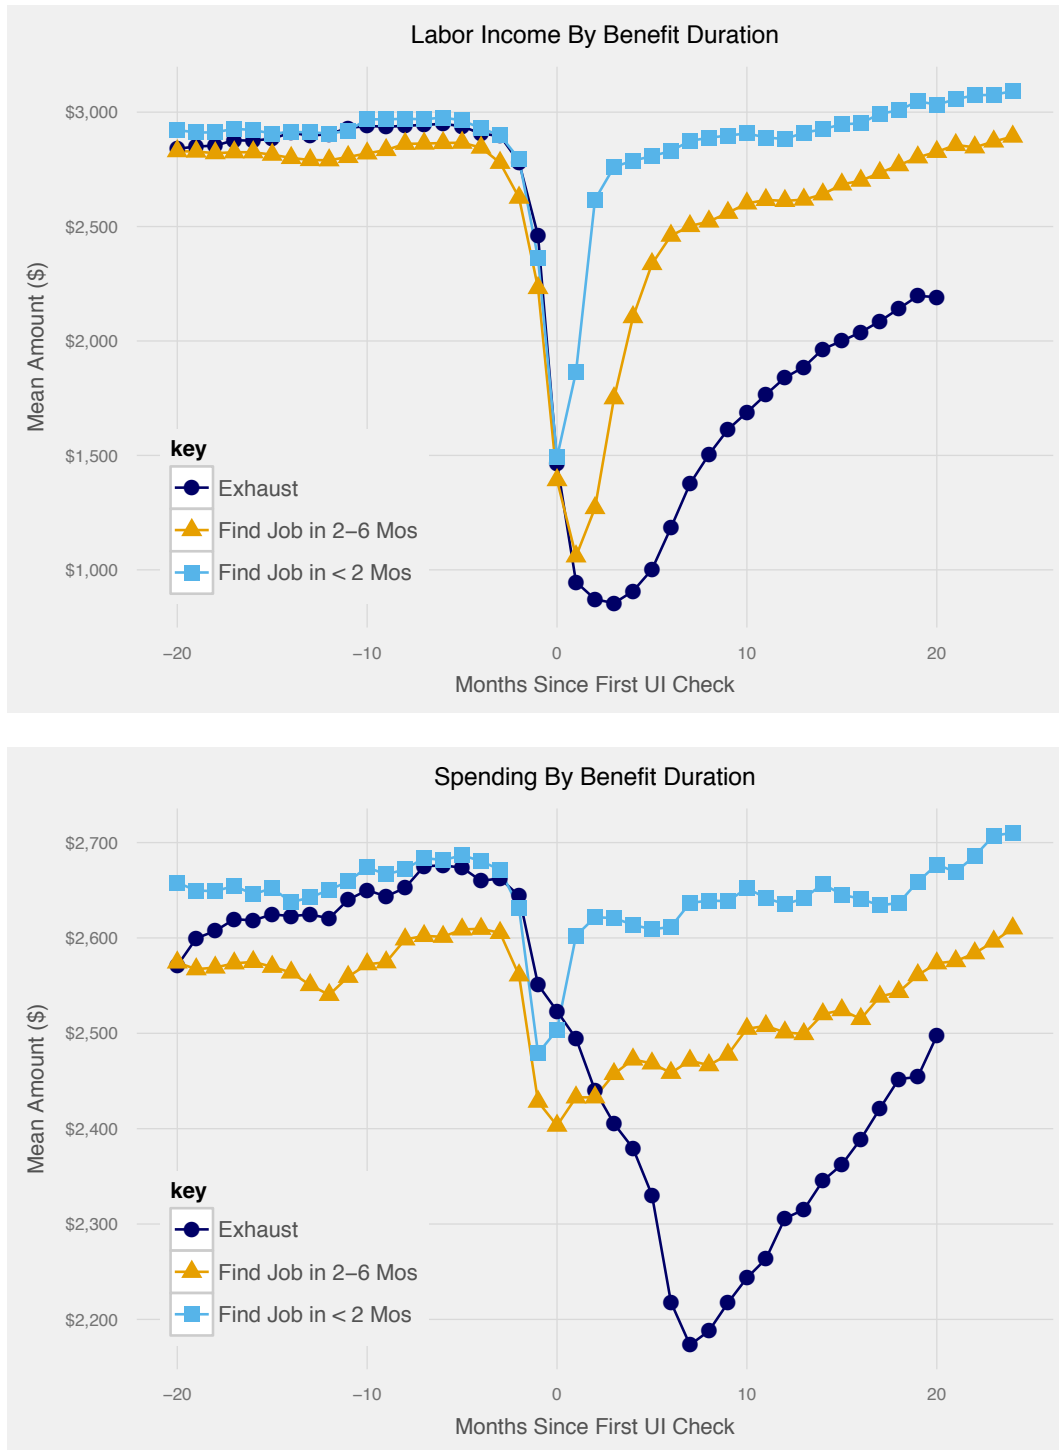
Note: About one-quarter of families have checking accounts at multiple banks. To understand how checking accounts outside the bank might bias our results, we study income and spending out of accounts which have not been linked together administratively, but have the same last name and address, suggesting that they belong to the same family.

APPENDIX FIGURE 4 – REPRESENTATIVENESS: INCOME RECOVERY AFTER ONSET



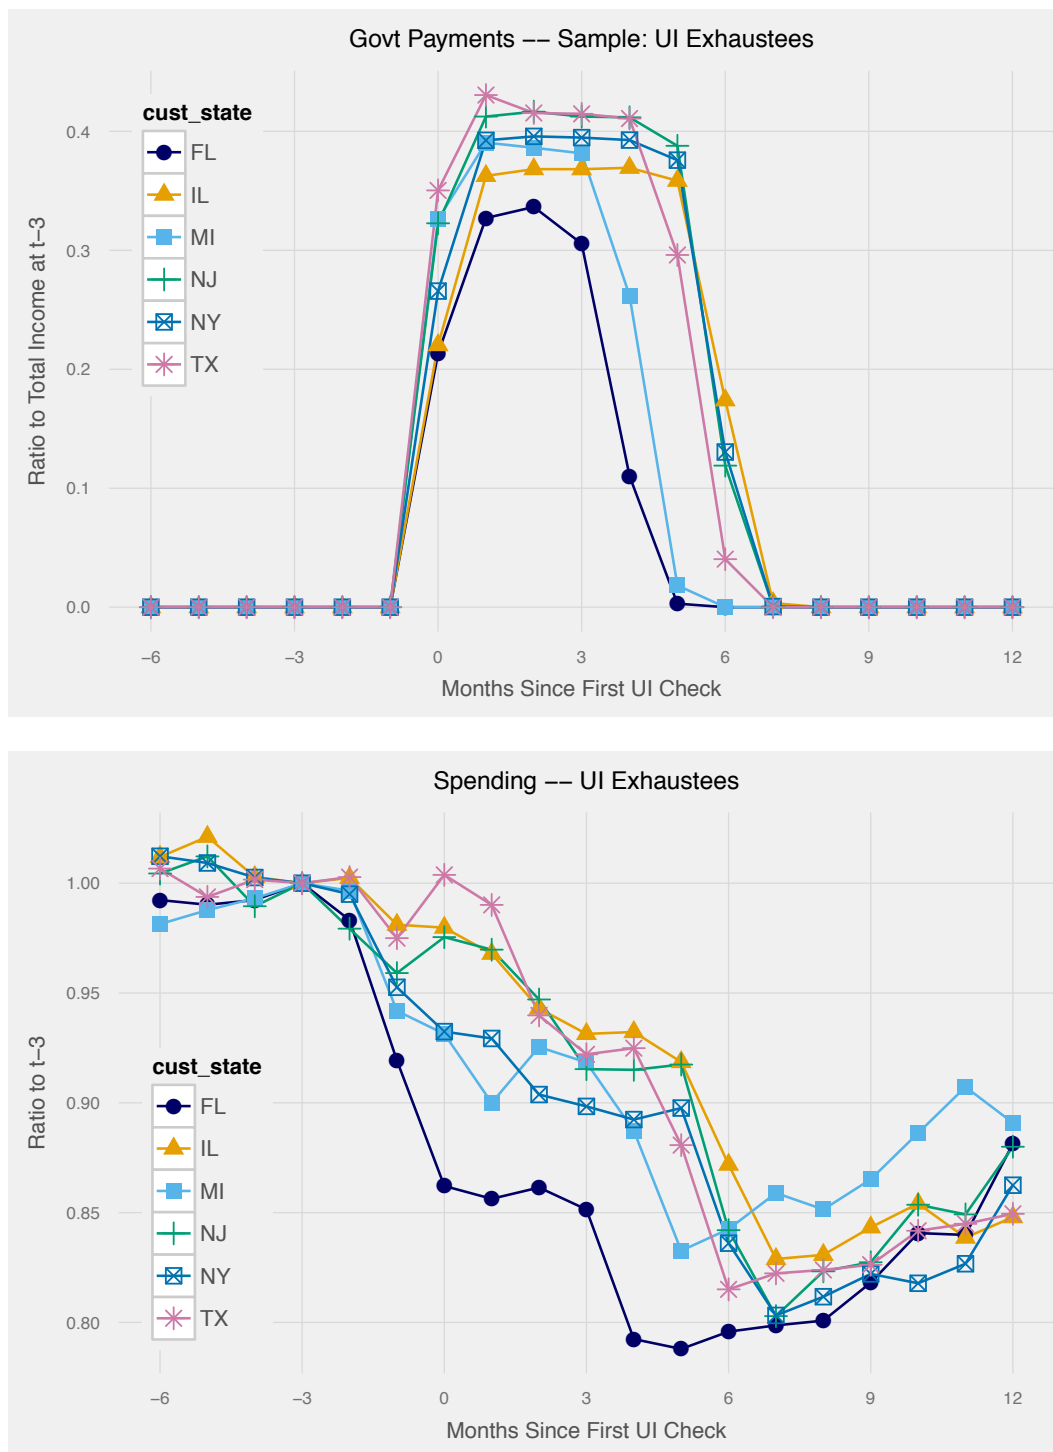
Note: This figure compares mean monthly family labor income around a UI spell in the 2004 SIPP and in the bank data. The SIPP shows a smaller drop in income than the bank data. This may be attributable to “seam bias”, where respondents who were re-employed report having positive earnings in all four months about which they are surveyed, even though in fact they were earning less in prior months. We use the 2004 SIPP rather than the 2008 SIPP because long follow-up horizons in the 2008 SIPP are available only for people who separated at the start of the Great Recession and therefore faced unusually bad job opportunities.

APPENDIX FIGURE 5 – INCOME AND SPENDING BY UI DURATION



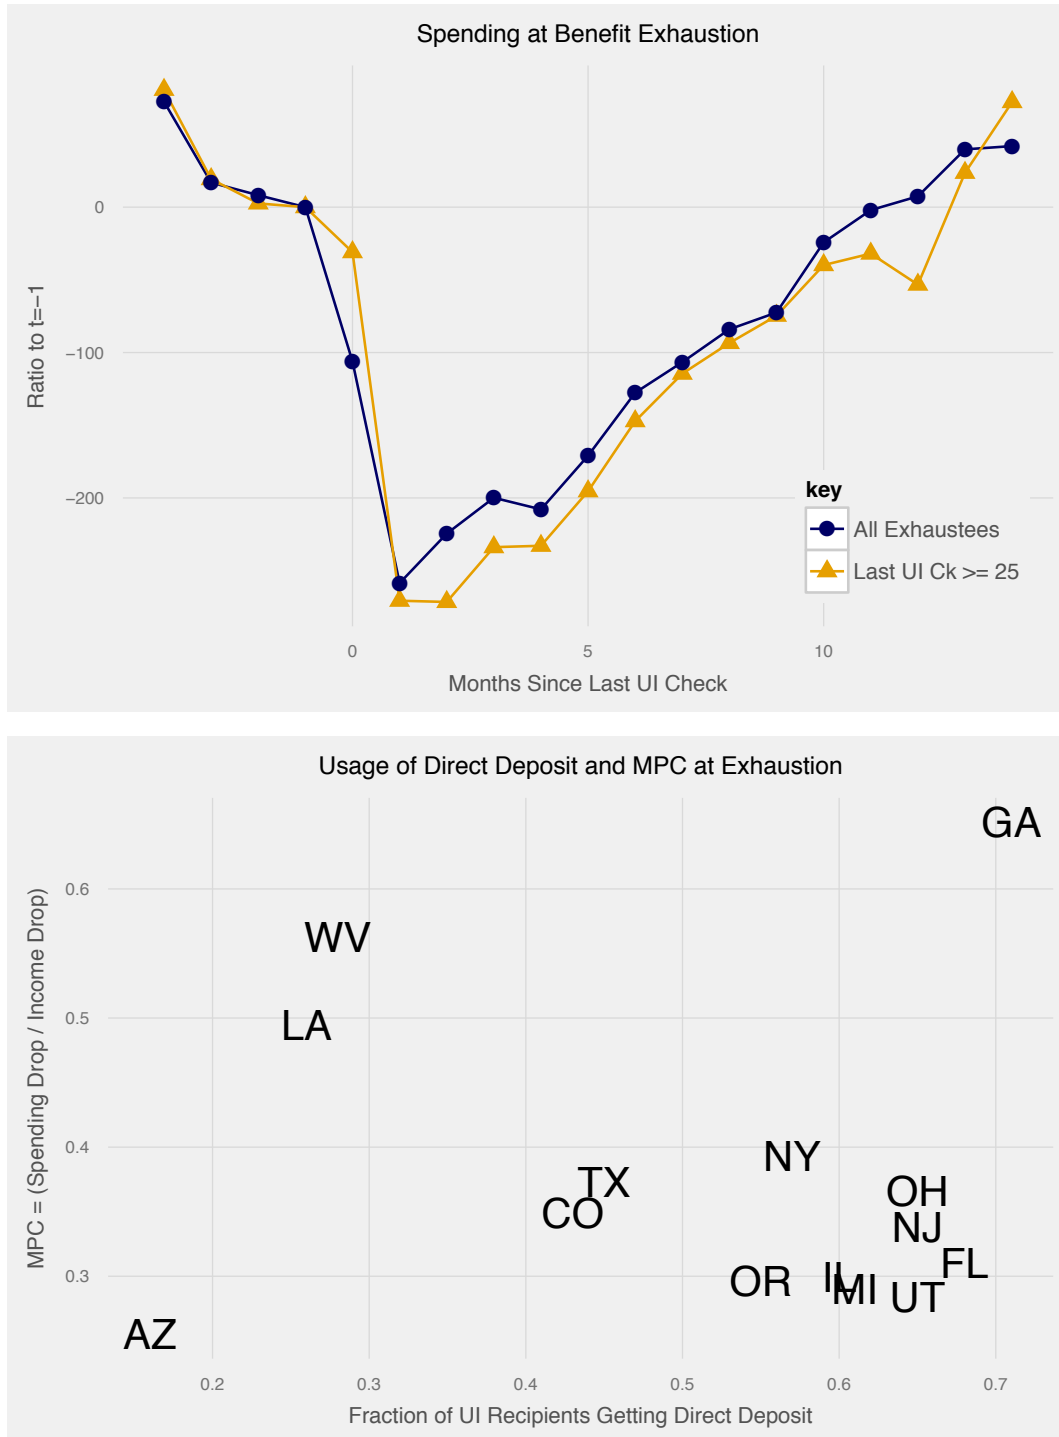
Note: The top panel replicates the bottom panel of Figure 5, which estimated the labor income drop and recovery for all UI recipients, but stratifies families by completed UI duration. The bottom panel examines the path of spending for the same three groups.

APPENDIX FIGURE 6 – EVENT STUDY FOR SIX LARGEST STATES



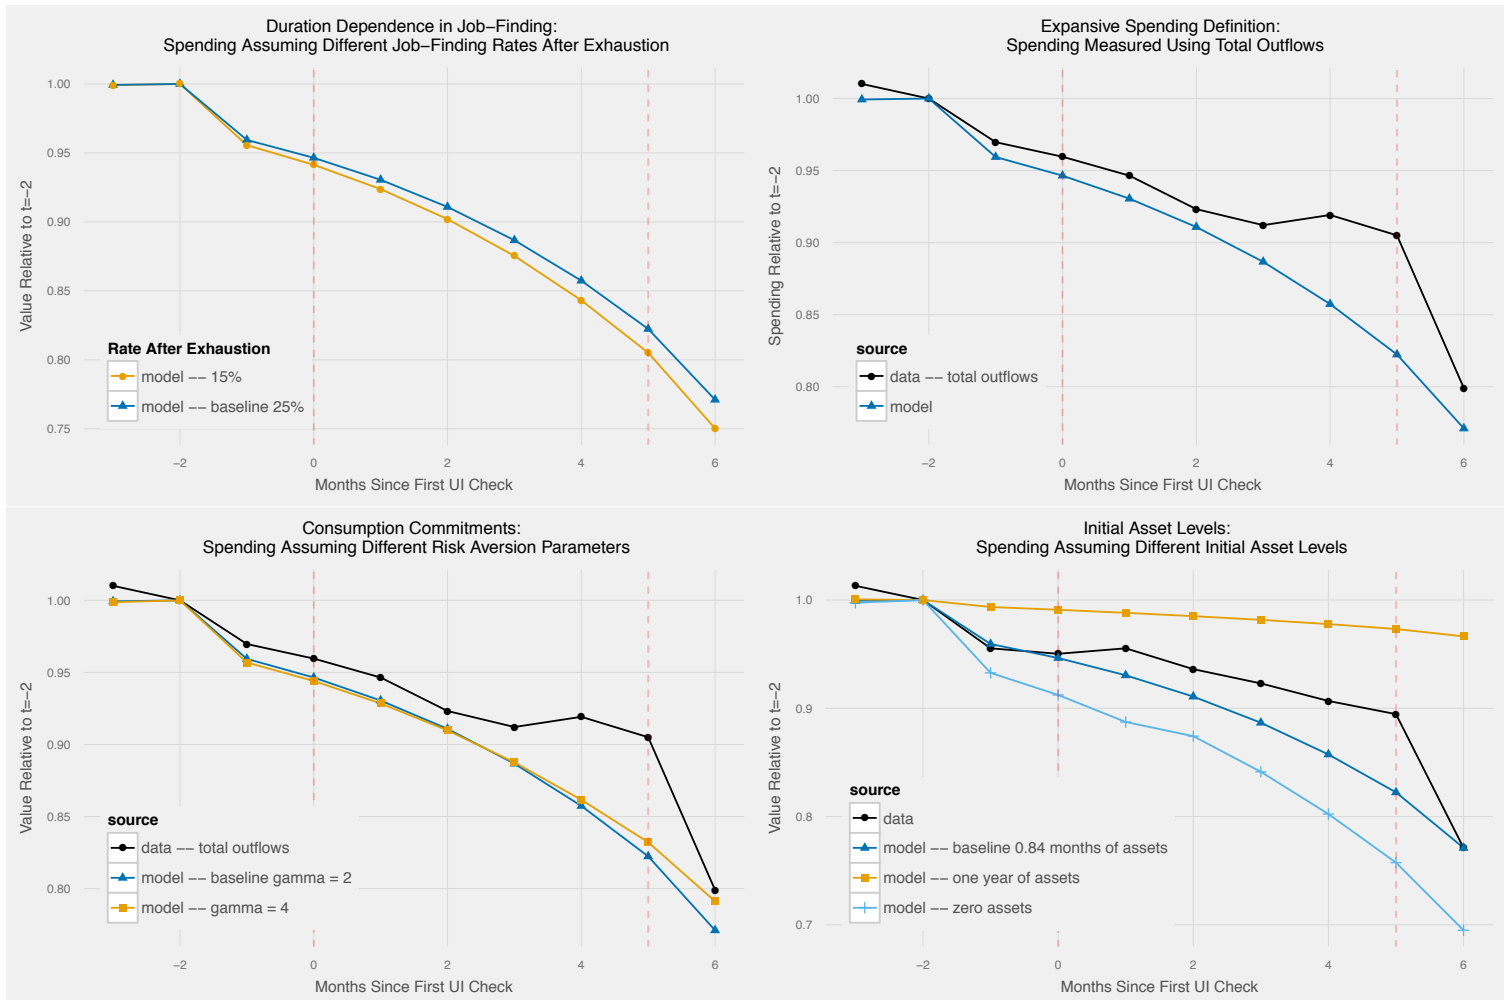
Note: The top panel plots the path of UI benefits for exhaustees in the six largest states in the data. Maximum benefit durations were shorter in Florida (16 weeks) and Michigan (20 weeks) than the 26 weeks of benefits available in most states. The bottom panel plots the path of spending for exhaustees in each of these six states.

APPENDIX FIGURE 7 – BENEFIT EXHAUSTION: ROBUSTNESS CHECKS



Note: The top panel compares the path of spending for UI exhaustees who received their last UI check on the 25th of the month or later to the path of spending for all UI exhaustees. The latter group appears to have benefits phase out over two months due to monthly time aggregation. The two-month magnitude of the spending drop is very similar for between the two groups. The bottom panel plots the estimated state-level marginal propensity to consume around benefit exhaustion against the share of UI recipients in each state who receive direct deposit (Saunders and McLaughlin (2013)).

APPENDIX FIGURE 8 – ROBUSTNESS CHECKS OF THE BUFFER STOCK MODEL



Note: The top left panel plots the path of spending predicted by the buffer-stock model under baseline job-finding beliefs post-exhaustion (25%) and assuming that the job-finding rate permanently drops to 15% post-exhaustion. The top right panel plots the predicted income path by the buffer-stock model against the path of spending in the data measured as total outflows net of transfers to savings accounts. The bottom left panel is the same as the top right, but adds a line showing the predicted path of spending from the buffer-stock model assuming γ , the coefficient of relative risk aversion, is equal to 4. The bottom right panel shows the path of spending in the data compared to the path predicted by the buffer-stock model assuming agents have different initial asset levels.

A Data Appendix

A.1 Sample Construction, Subsamples, and Winsorization

Unit of Observation. The core unit of observation is a set of bank accounts linked around a single primary account owner in the JPMCI. Many of these accounts have secondary owners who can also access the account. Some accounts have two people who jointly own the account. Sometimes, members of a family will not administratively link their accounts together; we exploit this feature of the data in Section B.1 to understand how missing accounts affect our analysis.

Classifying Primary Accounts with UI Spells.

Errors in transaction classification lead to measurement error of UI receipt, so we developed three criteria to establish whether a UI spell is plausible. First, a family must receive at least two UI payments. Second, the checks must have an amount and frequency which is reasonable given UI program rules – less than \$3,000 per month and fewer than 6 checks per month. Third, months with UI payments must be contiguous and observed duration must be less than or equal to program rules on potential benefit duration.⁵⁵ These restrictions serve to reduce measurement error due to erroneously classified non-UI transactions and provide a clear benefit exhaustion date, which is necessary for our analysis in Section 5. Of the roughly 1 million families with any potential UI receipt, 57% meet these criteria.

We classify families who bank primarily with Chase as those with five outflows from their checking accounts each month. To be conservative, we select families who have five outflows in each of the three months prior to their UI spell, five monthly outflows during their UI spell, and five monthly outflows in each of the three months following their UI spell, if their UI spell ends before the end of the panel. This restriction also reduces sample size – of the 586,000 families with a UI spell, about 376,000 meet this primary account criteria. In our robustness analysis in Section B.1, we repeat our analysis dropping this sample criteria and also using subsamples which offer even better coverage of family finances, but come at the expense of studying a more highly-selected sample.

Finally, we study UI spells which start in January 2013 or later, so that we have at least three months of pre-UI data on each family. This screen brings us to our baseline analysis sample of about 235,000 families.

Subsamples. We use three subsamples of these data in our analysis:

- **While Unemployed** – In some places, we analyze the spending of families where a member is unemployed. Because spending jumps up in the month *prior* to the last UI check, if a family received benefits for T months, we define this sample using months 0 to $T - 1$. Figures 2, 4 and 6 as well as Table 4 use this sample.
- **Exhausted Benefits** – We analyze exhaustees who had a potential benefit duration of 26 weeks or less, by focusing on UI recipients whose last UI check was paid in February 2014 or later. We measure duration in weeks as the date from the first UI check to the

⁵⁵In practice, this means that we also require the UI spell to be fifteen months or less. We are only able to compute the duration of UI spells which begin in November 2012 or later. Extended benefits which were legislated in response to the Great Recession expired in December 2013 and the last payments for these benefits were made in January 2014, which is why fifteen months is the maximum in the JPMCI data.

last UI check. Exhaustees are those who received benefits for a number of weeks equal to the current potential benefit duration in their state (usually 26 weeks, but lower in Florida, Michigan and Georgia), plus or minus two weeks for administrative error. We use this sample in Table 4, Table 7, Figure 9, Appendix Figure 6 and Appendix Figure 7.

- **26-Week Potential Benefit Duration** – In the model, we are specifically interested in the forward-looking behavior of a family eligible for 26 weeks of benefits. Here, we analyze UI recipients whose last UI check was paid in February 2014 or later and did not live in Florida, Michigan or Georgia.

Winsorization In general, we winsorize all variables at the 95th percentile of the set of observations with positive values. The one exception is in Table 5, 7 and Appendix Table 1 to preserve an additive decomposition across inflow and outflow categories, we instead *drop* families with inflows or outflows greater than the 95th percentile.

A.2 Categorizing Income and Spending

Group	JPMCI Category (Selected Examples)	% of Flows
Inflows		
Labor	Payroll, Direct Deposit ⁵⁶	61%
Government Income	Tax Refunds, Social Security (Old Age and Disability), Child Support, Unemployment Insurance, Veterans' Benefits, Supplemental Security Income	4%
Other Income	Cash, Investment Income, Interest, Refunds	4%
Unclassified	Paper Checks	21%
Dissaving	Transfers from Checking, Savings, Money Market, and Investment Accounts	10%
Outflows		
Debit Card		33%
Cash Withdrawal		14%
Bill Payments	Telecom Bill by ACH, Electric Bill by ACH, or Payment Method Used Primarily for Bills	7%
Installment Debt	Mortgage, Home Equity, Auto Loan, Student Loan	10%
Credit Card Debt		7%
Paper Checks		13%
Unclassified	PayPal, Misc ACH, tax payments	10%
Saving	Transfers to Money Market, Savings, and Investment Accounts	
Notes: % of flows measured for UI recipients three months prior to UI spell. This sample is defined in Section 2.1.		

B Robustness Checks

B.1 Empirics – Onset of Unemployment

The decline in spending at onset appears to reflect a true drop in family-wide spending rather than a shift in spending to alternative payment channels. First, as discussed in Section 2.1, 27% of families have checking accounts at multiple banks. One way to estimate if unemployment affects spending at outside checking accounts is to examine *unlinked* checking accounts within Chase for customers who share a last name and mailing address.⁵⁷ This could occur if, for example, two Chase customers formed a family unit without linking their accounts administratively. We find that spending in these unlinked accounts falls by \$51 at the onset of unemployment. Because the spending drop is computed using a larger denominator, we now find a 6% drop at onset across all accounts in this subsample, rather than an 8% drop in only the linked accounts. Second, families could shift spending from the debit card linked to their checking account to a credit card which did not need to be paid immediately. Outstanding balances on all credit cards in the credit bureau records rise by \$60 over a *two-month* period, so families either increase card spending or reduce card payments by \$30 each month. It is unclear whether this reflects increased spending on credit cards or reduced payments on outstanding credit card debt. We estimate that the change in spending through alternative payment channels is \$35 ($27\% \times \$51 + 72\% \times \30).⁵⁸

Our results for the sample of families with direct deposit of UI and at least five outflows per month appear to have external validity for other UI recipients. One concern is that families who adopt direct deposit of UI will be more financially sophisticated and better at smoothing than the typical family. We analyze the drop in spending at onset for the five states in the data with the highest adoption rate of direct deposit of UI: Georgia, Ohio, New Jersey, Florida, and Utah. According to Saunders and McLaughlin (2013), at least 65% of UI claimants receive their benefits using direct deposit. In these states, the drop in spending at onset is 8%, which is close to our overall estimate of 6%.

B.2 Empirics – Benefit Exhaustion

We implement the same robustness checks for internal and external validity at exhaustion as we did at onset. The empirical strategies are described in detail in Section B.1 and here we review only the results. Spending out of unlinked accounts rises slightly at benefit exhaustion, by \$48 per month. Because the spending drop is computed using a larger denominator, we now find an 8% drop at exhaustion across all accounts in this subsample, rather than an 14% drop in only the linked accounts. Remember that this modification applies to only the one-quarter of families with accounts at multiple banks, so the impact on our modification on our overall results is limited. Borrowing on Chase credit cards rises by about \$30 per month (Table 2), which appears to be driven by decreased payments rather than substituting consumption to credit cards. Credit bureau records show no additional borrowing on non-Chase credit cards.

⁵⁷About 10% of families that receive UI have not linked all of their accounts together. At no point during this analysis did we see personally identifiable information. Rather, the dataset included a numeric identifier which grouped together unlinked accounts which had the same last name and street address.

⁵⁸The Survey of Consumer Payment Choice estimated that 72% of people have at least one credit card.

B.3 Empirics – Income Recovery Rates in Other Datasets

One area where the literature has not reached consensus is in understanding the path of earnings *prior* to a separation. Jacobson et al. (1993b) and Jacobson et al. (1993a) find that mass layoff separators as well as UI recipients show declining wages in the years *prior* to separation. JLS argue that this reflects declining worker productivity as well as declining firm labor demand (e.g. overtime). The JPMCI data as well as our plots from the SIPP show roughly constant earnings prior to separation. Wachter et al. (2009) show sharply *rising* earnings in the years prior to separation. Understanding the reasons for these disparate trends is an important area for future work.

Our analysis uses the 2004-2007 SIPP panel, because the economic climate during this survey better reflects the labor market in 2013 and 2014 than the 2008-2012 SIPP panel. Earnings losses are deeper for UI recipients in the 2008 SIPP panel, where UI recipients searched for work in the midst of a severe recession.

One additional challenge for this exercise is estimating the earnings counterfactual in the absence of the UI separation. The analysis above has focused on whether earnings return to their pre-separation level. Some researchers have used workers who never separate as a control group. This choice seems problematic because UI recipients have lower labor income prior to separation and education than the typical employee and earnings rise faster over the lifecycle for employees with more education. Finding a suitable control group that matches UI recipients on observables seems necessary for accurately calculating a counterfactual.

B.4 Model

Below, we describe some alternative parameterizations of the model which do not change our substantive results. The results are shown graphically in Appendix Figure 8.

- Duration dependence in job-finding – We change the model by assuming that the job-finding rate falls permanently after exhaustion from 25% to 15%. With this change, we find that agents reduce their consumption slightly more during UI receipt to prepare for the possibility of longer unemployment.
- More expansive definition of spending – We consider an alternative expenditure series where we categorize all non-saving outflows as consumption. The path of spending using this definition is slightly smoother, with slightly smaller discrete drops at onset and exhaustion and a larger monthly drop as the spell progresses. In addition, in the two months prior to exhaustion there is a slight uptick in spending.
- Higher risk aversion to reflect consumption commitments – Chetty and Szeidl (2007) find that individuals with large consumption commitments effectively have larger risk aversion while unemployed. To allow for this case, we consider a case with coefficient of risk aversion $\gamma = 4$. As the figure shows, the risk aversion parameter has very little impact on the predicted consumption path.

Table 1 -- Representativeness: Income in JPMCI Data Compared to External Benchmarks

Dataset	Sample	Share <	Median	Mean	Poverty	Mean	Person	Other	Others'
		Age 21	Monthly	Monthly	Rate	Earnings	Earn	Earn > 0	Earn Mean
		(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
SIPP	Employed	0.06	6029	7405	0.07	6866	3739	0.60	3126
SIPP	All Unemployed	0.22	4374	5596	0.16	5064	2042	0.56	3023
SIPP	Get UI	0.02	5106	6290	0.08	5750	3273	0.54	2477
JPMCI	Get UI		4510	5445		3652			
JPMCI	Exhaust UI		4521	5414		3590			

Notes: All income statistics are monthly, for the 12-month period prior to the onset of unemployment.

SIPP The first three rows are from the Survey of Income and Program Participation panel (SIPP) and are inflated to 2014 \$ using CPI-U. This survey covered years 2004-2007. "All unemployed" are people with a reported job separation followed by unemployment in the subsequent month. "Get UI" are people who report positive UI income.

JPMCI data are for Oct 2012-May 2015. We define income as all inflows which are not explicitly categorized as dissaving and we rescale these inflows into pre-tax dollars. Earnings includes only labor income paid by direct deposit. About 86% of payroll dollars in the US are paid by direct deposit.

Table 2 -- Representativeness: Spending in JPMCI Data Compared to External Benchmarks

Category	JPMCI			External Benchmarks			
	Unadjusted Mean (\$)	Adj Factor	Adjusted Mean (\$)	CEX (\$)	Ratio	BEA (\$)	Ratio
<u>Headline</u> ^a Nondurable Goods and Services	--	--	1797	1912	0.94	4130	0.44
<u>Specific Nondurables</u> ^b							
Food At Home	254	0.59	432	331	1.31	580	0.74
Food Away From Home	148	0.59	252	219	1.15	471	0.53
Fuel	126	0.59	214	218	0.98	277	0.77
Utilities	300	1.00	300	312	0.96	--	--
<u>Debt Payments</u> ^c				SCF (\$)	Ratio		
Mortgage	212	0.2	1060	1368	0.77		
Auto Loan	73	0.22	332	465	0.71		
Credit Card	595	0.66	902	1613	0.56		
Student Loan	30	0.15	200	304	0.66		

Notes: All spending estimates are monthly. For external benchmarks, we use published 2013 Consumer Expenditure Survey (CEX) statistics, Bureau of Economic Analysis Table 2.3.5 for 2013 divided by 125 million consumer units, and 2013 Survey of Consumer Finances (SCF) microdata for employed families. Estimates from JPMCI data use all families with at least five outflows per month in 2013.

a. Headline We exclude healthcare because checking account data miss lots of healthcare spending and utilities because the BEA does not report them separately.

b. Specific Nondurables To capture a family's total spending on each category, we adjust food and fuel spending estimates upward by the ratio of Chase card spend to cash + debit card + all credit card spend (0.59). BEA reports food services together with accommodations, so the BEA estimate overstates true spending on food away from home.

c. Debt Payments We are only able to identify debt payments made by direct deposit for a small fraction of households. We compare the average payment made by households making any payment in the JPMCI data to comparable estimates in the SCF.

Table 3 -- Representativeness: Assets in JPMCI Data Compared to External Benchmarks

Data Source	Sample	Asset Balance	p10	p50	p90	Mean
SCF	All Employed	All Liquid Assets	270	4900	54000	29952
SCF	All Employed	Checking Account	150	1500	10000	4920
JPMCI	All Employed	Checking Account	80	1460	10940	5766
JPMCI	Employed, Pre-UI Receipt	Checking Account	40	1020	6900	3465

Notes: This table compares liquid assets in the 2013 Survey of Consumer Finances (SCF) to families with primary accounts at JPMCI from October 2012 through May 2015. Liquid assets include checking and saving accounts, money market accounts, certificates of deposit, savings bonds, non-retirement mutual funds, stocks and bonds. When households have multiple checking accounts, the primary checking account is defined in the SCF as "the one you use the most." Employed is defined as \$15,000 of annual pre-tax labor income in the SCF and \$1,000 of monthly post-tax labor income in the bank.

Table 4 -- Summary of Changes at Onset, During UI Receipt, and Benefit Exhaustion

	Pre-Onset Mean	Two-Month Drop at Onset (t = -3 to t = -1) ^a	Monthly Drop During UI Receipt ^b	Two-Month Drop at Exhaustion ^c
	(1)	(2)	(3)	(4)
Checking Account: Income and Spending				
Income (% of Pre-Onset Mean)		-0.114 (0.001)	-0.021 (0.001)	-0.228 (0.004)
Income (\$)	3520	-401 (5)	-73 (2)	-802 (13)
Total Inflows (\$)	5822	78 (37)	-194 (13)	-281 (92)
Spending on Nondurables (% of Pre-Onset Mean)		-0.061 (0.001)	-0.008 (0.0004)	-0.098 (0.003)
Spending on Nondurables (\$)	2644	-161 (3)	-22 (1)	-259 (8)
Total Outflows (\$)	5739	-59 (33)	-99 (12)	-243 (72)
Checking Account: Asset Flows				
Net Dissaving from External Accts (\$)	210	38 (2)	14 (1)	97 (5)
Balance Pre - Balance Post (\$) (Outflows - Inflows)	-16	26 (4)	31 (1)	106 (9)
n Checking Account Outcomes		235,245	677,628	35,507
Chase Credit Cards^d				
Revolving Balance (\$)	2288	-3 (5)	8 (2)	56 (13)
New Charges (\$)	208	-10.1 (0.9)	0.4 (0.3)	-0.7 (2.3)
Credit Bureau Records				
All Credit Cards -- Balance (\$)	6883	36 (9)	31 (4)	66 (24)
n Credit Card Outcomes		91,779	269,223	14,496

Notes: Standard errors are shown in parentheses underneath regression coefficients.

a. Changes at the onset of unemployment. We define this as from three months before the first UI payment to one month before the first UI payment. Each observation is a family.

b. Monthly changes while receiving UI. Each observation is a family-month. Standard errors in this column are clustered at the family level.

c. Changes at the exhaustion of UI benefits. We define this as from one month before the last UI payment to one month after the last UI payment for benefit exhaustees. Sample is exhaustees eligible for 26 weeks of benefits or less. Each observation is a family.

d. Credit card balance variables capture stocks rather than flows. For example, a \$36 increase in credit card balance at onset corresponds to spending \$18 extra on the card each month.

Table 5 -- Income and Spending at Onset of Unemployment

	Pre: 3 Months Before First UI Check (1)	Post: 1 Month Before First UI Check (2)	% Change (2)/(1) (3)
A. Total Inflows	4371	4157	-4.9%
Labor Direct Deposit	2656	2118	-20.3%
Govt: IRS, SS, DI, SSI	192	314	63.5%
Paper Checks	928	1057	13.9%
Other Income	154	173	12.3%
Unclassified	8	9	12.5%
Dissaving	433	486	12.2%
B. Total Outflows	4360	4191	-3.9%
Card: Work-Related	689	626	-9.1%
Card: Non-Work-Related	780	743	-4.7%
Cash Withdrawal	589	543	-7.8%
General Bills	308	319	3.6%
Credit Card Bills	302	302	0.0%
Installment Debt	454	443	-2.4%
Paper Checks	545	534	-2.0%
Unclassified	443	435	-1.8%
Saving	252	246	-2.4%
C. Selected Categories Ranked By Size of Drop			
Any Student Loan Pay	0.127	0.108	-15.2%
Transport	180	160	-11.1%
Food Away From Home	185	164	-11.1%
Any Medical Copay	0.25	0.228	-8.7%
Any Flights/Hotels	0.149	0.137	-7.9%
Retail	356	336	-5.7%
Food At Home	297	287	-3.3%
Any Auto Loan Pay	0.176	0.171	-2.8%
Telecom	105	103	-2.4%
Utilities	162	161	-0.7%
Any Entertainment	0.439	0.442	0.9%
Any Mortgage Pay	0.149	0.151	1.9%
Any Credit Card Pay	0.542	0.553	2.1%

Notes: n=235,245. The top two panels presents a decomposition of the change in inflows and outflows at onset. To make this decomposition less sensitive to outliers, we drop observations with inflows above the 95th percentile in the pre or post period for panel A and outflows above the 95th percentile for panel B. In panel C, we winsorize each continuous outcome variable at the 95th percentile.

Table 6 -- Comparison of Spending Drop to Prior Literature

Spending Drop Compared to 3 Months Before UI Onset					
	Pre-Onset Mean (1)	Onset ^a (t=-1) (2)	While Receiving UI ^b (3)	Annual ^c (t=-1,0,...10) (4)	Literature Replication ^d (5)
(1) Total Nondurables (2 + 3 + 4)	2641	-5.9%	-7.7%	-7.3%	-20.4%
(2) Work-Related	671	-8.8%	-10.6%	-8.3%	-22.3%
(3) Non-Work-Related	773	-5.2%	-4.7%	-5.5%	-16.2%
(4) Cash and Bills	1197	-4.7%	-8.0%	-7.9%	-22.1%
(5) Food	489	-6.4%	-5.6%	-4.3%	-10.3%

Notes: This table computes the spending drop for various time horizons and various spending measures.

Time subscripts are relative to the first month of UI receipt and T is the last month of UI receipt. In each column, we compute Loss/Spend₃.

a. Loss = Spend₁ - Spend₃.

b. Loss = Mean(Spend₁, Spend₀ ... Spend_T) - Spend₃

c. Loss = Mean(Spend₁, Spend₀ ... Spend₁₀) - Spend₃.

d. Loss = (Mean(Spend₁, Spend₀ ... Spend₁₀) - Spend₃)/(T/12). This is the calculation done by Chodorow-Reich and Karabarbounis (2015) and Kolsrud et al. (2015).

Table 7 -- Income and Spending for Families Who Exhaust UI Benefits

	Pre Onset	Pre Exhaustion	Post Exhaustion	% Change (3)/(2)
	(1)	(2)	(3)	(4)
A. Total Inflows	4068	3649	3145	-13.8%
Labor Direct Deposit	2522	683	1086	59.0%
Govt: IRS, SS, DI, SSI	178	1662	296	-82.2%
Paper Checks	827	663	944	42.4%
Other Income	155	176	223	26.7%
Unclassified	7	12	15	25.0%
Dissaving	379	454	582	28.2%
B. Total Outflows	4059	3747	3415	-8.9%
Card: Work-Related	637	557	498	-10.6%
Card: Non-Work-Related	752	728	632	-13.2%
Cash Withdrawal	588	499	417	-16.4%
General Bills	316	321	301	-6.2%
Credit Card Bills	280	272	255	-6.2%
Installment Debt	402	373	359	-3.8%
Paper Checks	503	452	434	-4.0%
Unclassified	379	375	365	-2.7%
Saving	203	170	154	-9.4%
C. Selected Categories Ranked By Size of Drop				
Food At Home	295	287	250	-13.0%
Retail	352	329	287	-12.8%
Any Medical Copay	0.249	0.222	0.197	-11.6%
Food Away From Home	177	155	140	-9.9%
Any Entertainment	0.415	0.414	0.377	-9.0%
Any Flights/Hotels	0.146	0.128	0.117	-8.9%
Any Student Loan Pay	0.12	0.091	0.083	-8.8%
Telecom	107	107	98	-8.6%
Transport	177	150	138	-8.0%
Utilities	178	175	165	-6.0%
Any Auto Loan Pay	0.175	0.167	0.158	-5.2%
Any Mortgage Pay	0.164	0.162	0.156	-3.5%
Any Credit Card Pay	0.561	0.575	0.563	-2.0%

Notes: n=35,507 families who exhausted UI benefits and had potential benefit duration of 26 weeks or fewer. Pre Onset is three months prior to first UI payment, Pre Exhaustion is the month before UI Exhaustion and Post Exhaustion is the month after UI exhaustion. The top two panels present a decomposition of the change in inflows and outflows at onset. To make this decomposition less sensitive to outliers, we drop observations with inflows above the 95th percentile in the pre or post period for panel A and outflows above the 95th percentile for panel B. In panel C, we winsorize each continuous outcome variable at the 95th percentile.

Table 8 -- Model Parameters

Parameter	Value
Income and Assets Matched to Bank Data	
Income z_t	1 Employed
	0.84 Unemp ≤ 7 months
	0.53 Unemp > 7 months
Transition Matrix Π	
	u e
	e ₋₁ 0.0325 0.9675
	u ₋₁ t $\neq 7$ 0.75 0.25
	u ₋₁ t=7 0.7 0.3
Preferences & Environment	
N Months of Life	240
Monthly Discount Factor β	0.996
Risk Aversion γ	2
Monthly Interest Rate R	1.0025

Notes:

Income: z includes UI benefits, labor income from other family members, and non-labor income. Levels calibrated to match bank data.

Transition Matrix: Matches the transition rates in bank data during employment and first six months of unemployment. Surge in job-finding at exhaustion matches Card, Chetty, and Weber (2007). Benefits last six months, but because labor income declines prior to onset, as shown in Figure 2, we assume a seven-month duration.

Appendix Table 1 -- Summary Statistics Prior to Onset

Category	Mean	Median	Std Dev	Share > 0
	(1)	(2)	(3)	(4)
A. Total Inflows	4371	3520	3069	1
Labor Direct Deposit	2656	2160	2473	0.81
Govt: IRS, SS, DI, SSI	192	0	770	0.12
Paper Checks	928	140	1634	0.62
Other Income	154	0	498	0.48
Other Inflows	8	0	114	0.23
Dissaving	433	0	1087	0.42
B. Total Outflows	4360	3540	2961	1
Card: Drops at Retirement	689	480	795	0.95
Card: Stable at Retirement	780	600	690	0.96
Cash Withdrawal	589	300	879	0.82
General Bills	308	160	1008	0.81
Credit Card Bills	302	0	663	0.35
Installment Debt	454	20	1023	0.52
Paper Checks	545	60	990	0.55
Unclassified	443	60	907	0.67
Saving	252	0	745	0.39

Notes: n= 235,245. This table presents summary statistics on the analysis sample three months prior to the onset of UI. To make this decomposition less sensitive to outliers, we drop observations with inflows above the 95th percentile in the pre or post period for panel A and outflows above the 95th percentile for panel B. Medians are for data to the nearest \$20 bin to prevent disclosure of individual observations.

Appendix Table 2 -- Additional Borrowing Outcomes

	Two-Month Drop at Onset (t=-3 to t=-1) ^a	Monthly Drop During UI Receipt ^b	Two-Month Drop at Exhaustion ^c
Chase Credit Cards			
Credit Limit (\$)	30 (6)	-3 (2)	-11 (15)
Credit Score (sd) ^d	0.0025 (0.0012)	-0.0006 (0.0006)	-0.0015 (0.0133)
Credit Bureau Records			
All Credit Cards -- Credit Limit (\$)	14 (0)	6 (0)	39 (0)
Share with Any Delinquency (sd) ^d	0.0074 (0.0019)	0.0084 (0.0007)	0.0264 (0.0053)
n	91,779	269,223	14,496

Notes:

- Changes at the onset of unemployment. We define this as from three months before the first UI payment to one month before the first UI payment. Each observation is a family.
- Monthly changes while receiving UI. Each observation is a family-month. Standard errors in this column are clustered at the family level.
- Changes at the exhaustion of UI benefits. We define this as from one month before the last UI payment to one month after the last UI payment for benefit exhaustees. Sample is exhaustees eligible for 26 weeks of benefits or less. Each observation is a family.
- To prevent disclosure of sensitive data, this variable has been rescaled to have mean 0 and standard deviation 1.

Appendix Table 3 -- Spending Recovery After Reemployment

	Low Asset	Medium Asset	High Asset
	(1)	(2)	(3)
Baseline Characteristics			
Labor Direct Deposit + Govt (Mean, \$)	2766	3140	4521
Inflows (Mean, \$)	3454	4138	6872
Nondurables Spending (Mean, \$)	2382	2520	2954
Outflows (Mean, \$)	3477	4049	5911
Liquid Asset Holdings (Median, \$) ^a	965	5000	21522
Months of Assets (Liquid Assets / Inflows)	0.28	1.21	3.13
Cumulative Sum While Unemployed, t= -3 to t= 2 ^b			
Income Reduction (Months Lost)	0.58	0.61	0.66
Spending Reduction (Months Lost)	0.31	0.27	0.2
Implied Decumulation (Income Lost - Spending Lost)	0.27	0.34	0.46
Percent of Liquid Assets Decumulated	97%	28%	15%
Cumulative Sum After Reemployment, t= 3 to t=12			
Spending Reduction (Months)	0.43	0.27	0.11
Percent of Initial Assets Rebuilt	154%	22%	4%

Notes: We divide the sample into terciles on the basis of their total liquid assets held both inside the bank and outside the bank.

a. We estimate liquid asset holdings as the median within each tercile in the Survey of Consumer Finances.

b. The cumulative loss of income is equal to about 0.6 months of baseline income for each group. The cumulative loss of spending during unemployment is equal to 0.2-0.3 months of baseline spending. Implied decumulation is income lost - spending lost. Percent of Liquid Assets Decumulated is the ratio of Implied Decumulation to Assets at onset.