NBER WORKING PAPER SERIES

CONSUMER SPENDING DURING UNEMPLOYMENT: POSITIVE AND NORMATIVE IMPLICATIONS

Peter Ganong
Pascal J. Noel

Working Paper 25417
http://www.nber.org/papers/w25417

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 01238
January 2019

This paper was previously circulated as "How Does Unemployment Affect Consumer Spending?" We thank Joe Altonji, Scott Baker, John Campbell, Raj Chetty, Gabe Chodorow-Reich, David Cutler, Stefano DellaVigna, Manasi Deshpande, Xavier Gabaix, Ed Glaeser, Je Grogger, Nathan Hendren, Erik Hurst, Simon Jäger, Damon Jones, Larry Katz, Rohan Kekre, David Laibson, Annie Levenson, Je Liebman, Mandy Pallais, Jonathan Parker, Bruce Meyer, Mikkel Plagborg-Møller, Marilyn Newman, Martin Rotemberg, Jesse Shapiro, Andrei Shleifer, Dan Shoag, Rob Vishny, Constantine Yannelis and three anonymous referees for helpful comments. Thanks to Ari Anisfeld, Guillermo Carranza Jordan, Chanwool Kim, Lei Ma, and Xian Ng for outstanding research assistance. Thanks to the Heterogeneous Agent Resource ToolKit (HARK, https://github.com/econ-ark/HARK) for model 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 de-identified data assets that are selectively available to be used for academic research. More information about JPMCI de-identified data assets and data privacy protocols are available at www.jpmorganchase.com/institute. All statistics from JPMCI data, including medians, reflect cells with multiple 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 compensated for providing research advice on public reports produced by JPMCI's research team.We gratefully acknowledge funding from the Washington Center for Equitable Growth, the Charles E. Merrill Faculty Research Fund at the University of Chicago Booth School of Business, 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. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Consumer Spending During Unemployment: Positive and Normative Implications
Peter Ganong and Pascal J. Noel
NBER Working Paper No. 25417
January 2019
JEL No. E21,E24,J65

ABSTRACT

Using de-identified bank account data, we show that spending drops sharply at the large and predictable decrease in income arising from the exhaustion of unemployment insurance (UI) benefits. We use the high-frequency response to a predictable income decline as a new test to distinguish between alternative consumption models. The sensitivity of spending to income we document is inconsistent with rational models of liquidity-constrained households, but is consistent with behavioral models with present-biased or myopic households. Depressed spending after exhaustion also implies that the consumption-smoothing gains from extending UI benefits are four times larger than from raising UI benefit levels.

Peter Ganong
Harris School of Public Policy
University of Chicago
1155 E. 60th Street
Chicago, IL 60637
and NBER
ganong@uchicago.edu

Pascal J. Noel
University of Chicago Booth School of Business
5807 South Woodlawn Avenue
Chicago, IL 60637
pascal.noel@chicagobooth.edu

A data appendix is available at http://www.nber.org/data-appendix/w25417

1 Introduction

How does consumer spending evolve during an unemployment spell? Understanding the path of spending during this particularly stressful period in a household's financial life is important for explaining fundamental household decisions and for designing an appropriate policy response. However, this question has been difficult to answer precisely in prior work because the existing data sources on spending during unemployment are mostly small, infrequent surveys. The goal of this paper is to document the monthly path of spending during unemployment in the U.S. and to assess its positive and normative implications.

Using de-identified bank account data, we find that spending is highly sensitive to income, both at the onset of unemployment and even at the large and predictable decrease in income arising from the exhaustion of unemployment insurance (UI) benefits. From a positive perspective, we use our empirical results to develop a new test to distinguish between the two kinds of theories that have been used to explain the excess sensitivity of spending to income. The behavior we document is inconsistent with theories based on benchmark rational models of forward-looking but liquidity-constrained households because liquidity constraints are unable to explain why households fail to save in anticipation of predictable income declines. However, the path of spending we document is consistent with models where some households exhibit present-biased or myopic behavior (Laibson 1997; Campbell and Mankiw 1989). From a normative perspective, we examine the implications of our results for UI policy. Because spending is so much lower after benefit exhaustion than during UI receipt, we find that the consumption-smoothing gains from extending UI benefits are four times greater than from increasing the level of UI benefits.

To document the path of spending, we worked with the JPMorgan Chase Institute (JPMCI) to construct a de-identified panel dataset with measures of monthly income and spending. The dataset is based on the universe of Chase consumer checking and credit card accounts, aggregated for households to the monthly level. Spending is measured from debit and credit card transactions, cash withdrawals, and electronic transactions captured through the bank account. Our analysis sample includes 182,000 households who received direct deposit of UI benefits into their bank accounts in 20 U.S. states between January 2014 and June 2016. Since UI recipients have substantial prior labor force attachment and 90 percent of UI recipients have a bank account, our sample is likely to be representative of UI recipients. In fact, UI recipients in the JPMCI data are similar to external benchmarks for total household income, spending on nondurables, checking account balances, and age. These data have four strengths relative to most prior work on spending during unemployment: a monthly frequency, detailed spending categories, substantial policy variation across 20 U.S. states, and the availability of monthly employment histories constructed from direct deposit paychecks.

In the first part of the paper, we exploit these four strengths of the JPMCI data to document four novel empirical findings. First, spending drops sharply at UI benefit exhaustion.

Previously-available data measured spending too infrequently to capture the drop in the precise month of benefit exhaustion. When we trace out the monthly path of nondurable spending, we find a drop of six percent at unemployment onset, a drop of less than one percent per month during UI receipt, and a 12 percent drop at UI benefit exhaustion. The drop at exhaustion is larger than the drop in all previous months combined. This finding is surprising because the exhaustion of UI benefits occurs predictably after six months for every state in our baseline sample. We explore the implications of this result for models of household consumption behavior below.

Second, although one might have thought that the drop in spending reflected a financial management strategy of no longer paying bills and continuing to buy necessities, in fact the opposite is true. Spending drops sharply on necessities such as groceries, medical copayments, and drugstores. The drop in spending persists for at least five months after exhaustion, so households are not simply eating previously-purchased nonperishable food and delaying medical payments by a month.

Third, variation in the path of income during unemployment in Florida and New Jersey illustrates that high-frequency changes in income cause the sharp changes in spending described above. Compared to other states, income is unusually low during unemployment in Florida because UI benefits replace a smaller share of household income and benefits last for only four months instead of six. We document a larger-than-average drop in spending at the start of UI receipt and a sharp drop when benefits expire after four months rather than after six months, as in the states with longer benefits. We use the path of spending in Florida as an out-of-sample test of the theoretical models examined later in this paper.

The most remarkable evidence of excess sensitivity comes from New Jersey, which has unusually generous UI benefits. UI payments there begin quickly, such that many workers receive their last paycheck and their first UI check in the same week. This induces a sawtooth pattern in average income, which falls as some workers become unemployed, then rises by 10 percent in the month with the extra check, and then falls to a stable level thereafter. Spending around onset follows the same sawtooth pattern: falls, rises, and then falls to a stable level, suggesting that much of the extra check is spent immediately, even though households know their income is likely to fall sharply in the following month. This pattern shows that the excess sensitivity we document holds for all UI recipients, not just UI exhaustees. Taken together, our evidence suggests that the sensitivity of spending to UI benefits is robust to different settings, different samples, different sources of policy variation, and different points in the spell.

Our fourth empirical finding is that job-finding spikes modestly at UI benefit exhaustion. The job-finding hazard spikes from 19 percent the month before exhaustion to 26 percent in the last month of benefits. Although spikes at exhaustion have been documented in European data, to the best of our knowledge this is the first estimate of the spike in the

job-finding hazard around exhaustion in the U.S.[1] By estimating the path of spending and job-finding for the same set of households, our work provides a single unified set of empirical targets for theoretical work investigating household behavior during unemployment. In the model section, we show that the observed spike in job-finding at benefit exhaustion is far too small to quantitatively explain the magnitude of the drop in spending.

The second part of our paper develops a novel test exploiting high-frequency data to distinguish between the two types of theories economists have developed to explain the empirical excess sensitivity of spending to predictable income changes: rational (or near-rational) models of forward-looking agents and behavioral models with present-biased or myopic agents. The rational models exhibit excess sensitivity because some agents happen to be close to their liquidity constraint when an economic shock arrives (Kaplan and Violante 2014; Carroll et al. 2017). In the behavioral models, agents have a self-control problem (Laibson 1997) or rule of thumb (Campbell and Mankiw 1989). Some of the strongest empirical evidence for excess sensitivity comes from tax rebates.[2] Unfortunately, because tax rebates are an *increase* in income, they do not distinguish between these two types of models: spending can increase in response to predictable income gains either because households are myopic or because liquidity constraints prevent them from smoothing consumption by increasing borrowing before they receive the income.

Our test is informative because these two types of theories make different predictions about the path of spending around a predictable *decrease* in income. To smooth consumption around an income decline, consumers only need a saving technology such as a bank account to build up a buffer of assets. Although liquidity constraints can explain why consumers fail to borrow, they cannot explain why consumers fail to save. A forward-looking but liquidity-constrained consumer will exhibit no excess drop in spending in the month when income falls. In contrast, a myopic consumer will sharply cut her spending once the predictable income decline is realized.

The path of spending around a predictable income decline is thus an "identified moment" in the language of Nakamura and Steinsson (2017). As they describe, estimates of the marginal propensity to consume (MPC) from tax rebates have successfully served as an identified moment that has convincingly ruled out classical models with complete markets and standard preferences in favor of alternative models with either liquidity-constrained or behavioral households. In this spirit, we propose the magnitude of the monthly consumption

[1] Prior work using U.S. data suggests there may be a spike in job-finding at UI benefit exhaustion. Meyer (1990) shows a spike in UI exit hazards in the weeks leading up to UI benefit exhaustion. However, there is considerable uncertainty in the prior literature about the magnitude of such a spike because Card et al. (2007) demonstrate using European data that job-finding hazards spike much less than UI exit hazards at UI exhaustion.

[2] Examples using tax rebates and refunds include Hsieh (2003), Johnson et al. (2006), Parker et al. (2013), Baugh et al. (2018), Souleles (1999), Shapiro and Slemrod (2009), and Kueng (2015). Baker and Yannelis (2017) and Gelman et al. (2015) examine the spending response to an unanticipated, temporary loss of income: the federal government shutdown. Pagel and Vardardottir (2016) examine the spending response to paycheck receipt.

drop in response to a predictable income decline as a new identified moment allowing us to further refine the class of potential models, distinguishing models with liquidity constraints from a class of behavioral models. While prior work seeking to distinguish between these theories has reached inconclusive results by testing for asymmetry in the low-frequency response to increases and decreases in income, our strategy instead examines the magnitude of the spending drop in the exact month of a sharp predictable income decline.[3]

The spending of the unemployed, particularly around the predictable decline in income at UI benefit exhaustion, is a useful laboratory to study consumption models for two reasons. First, unemployment–just like the receipt of tax rebates–is ubiquitous; over 90 percent of baby boomers in the National Survey of Longitudinal Youth have experienced unemployment. Second, focusing on UI exhaustion enables us to overcome a limitation of prior work that analyzes the predictable income decline arising from retirement, whose interpretation is ambiguous because of the increased time available for home production (Banks et al. 1998; Bernheim et al. 2001; Aguiar and Hurst 2005). Because UI benefit exhaustion is a change in income without a change in a household's time budget, this ambiguity does not apply to our research design.

To formalize this test, we examine a predictable income decline in a structural model of consumption and job search and compare various parameterizations of the model to the data. We study a model with endogenous consumption in the tradition of Deaton (1991) and Carroll (1997) and endogenous job search as in Mortensen (1977) and DellaVigna et al. (2017). We calibrate the model to match liquid assets, income levels, and job-finding hazards in the JPMCI data, and we estimate the discount rate, borrowing limit, and job search cost to best fit the observed path of spending and job search in the JPMCI data. To the best of our knowledge, we are the first to study the effects of a predictable income decline in a quantitative structural model.[4] Our quantitative model results confirm the intuition described above: the benchmark version of the model that allows for liquidity constraints (but not present-bias) predicts that agents cut spending gradually with no excess drop in spending at exhaustion, unlike the data.

In contrast, we show that models incorporating present-biased or myopic agents are able to closely match the JPMCI data. We examine a model where agents may differ in their present-bias preference parameter β (as in Laibson 1997). Agents with a sufficiently low β who are also liquidity-constrained set their consumption close to income in each month. The inclusion of a low-β type enables this model to reproduce the "spender-saver" behavior in Campbell and Mankiw (1989), albeit with a micro-foundation relying on heterogeneity in β. At exhaustion, we estimate that about one-third of agents are the low β types. A model with consumers whose exponential time preference parameter δ is so low that they are

[3]See Jappelli and Pistaferri (2010) for a summary of this prior work, and see Altonji and Siow (1987), Shea (1995), Bowman et al. (1999), and Garcia et al. (1997) for more details.

[4]Both Banks et al. (1998) and Bernheim et al. (2001) make the qualitative point that because retirement is a predictable income decline, a rational agent should not cut her spending.

essentially myopic also makes a similar prediction, albeit with slightly lower goodness-of-fit. These hybrid models are able to match the path of spending because the present-biased or myopic consumers sharply cut consumption around onset and exhaustion, while the presence of the non-present-biased buffer stock agents best matches the gradual decline in spending during UI receipt.

The final part of the paper examines the normative implications of our results for UI policy. The consumption implementation of the canonical Baily (1978)-Chetty (2006) formula for the optimal level of UI benefits requires the spending drop during unemployment as an input. Relative to the prior literature, we empirically contribute the first estimate of the drop at UI benefit exhaustion and a more precise and comprehensive measure of the average drop in consumption during UI receipt.[5]

These results enable us to estimate the welfare gain from extending the duration of UI benefits. Although most prior work on optimal UI benefits has focused on the *level* of UI benefits as the key policy parameter, there is substantial policy variation across countries, states, and time in the potential *duration* of benefits. In response, a newer strand of the literature has examined the optimal path of benefits (Schmieder and von Wachter 2017; Kekre 2017; Kolsrud et al. 2018). Although there is substantial research estimating the fiscal cost of extensions (summarized by Schmieder and von Wachter 2017), we are not aware of any paper that has estimated the consumption-smoothing gains from extensions.[6]

We find that the welfare gains from improved consumption-smoothing due to extending the duration of UI benefits are four times as large as from raising the level of UI benefits in a generalized Baily-Chetty formula. The economic intuition for our result is that long-term unemployment is a state of the world where consumption is much lower and marginal utility is much higher; benefit extensions target this high marginal utility state of the world effectively. These large consumption-smoothing gains imply that–even after incorporating the fiscal externality from increases in UI generosity–extensions raise welfare more than benefit increases of equivalent cost.

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 shows that spending is highly sensitive to the level and duration of UI benefits, including a sharp drop at benefit exhaustion. Section 4 compares predictions from different consumption models to the data.

[5]Researchers calibrating the Baily-Chetty formula have typically relied on the Panel Study of Income Dynamics (PSID), which has two limitations for this purpose. The first limitation is that the PSID is administered infrequently and has an ambiguous reference period, making it difficult to assess if survey responses about spending coincide with unemployment (Chodorow-Reich and Karabarbounis 2016). The second limitation is that until recently the survey covered only food and housing spending. Researchers have raised concerns that the measured drop in food expenditure fails to capture a larger drop in unmeasured consumption categories (Chetty and Szeidl 2007). Our estimates address these twin concerns by using monthly data and decomposing household expenditure into 27 detailed categories.

[6]In related work, Kolsrud et al. (2018) analyze the consumption-smoothing gains from changes to the path of UI benefits in Sweden, where exhaustion is not relevant because there is no time limit on benefit receipt.

Section 5 evaluates the consumption-smoothing gains from UI benefits. Section 6 concludes.

2 JPMCI Data and External Validity

We study the evolution of consumption during unemployment, building on work by Kolsrud et al. (2018). Our data are novel in two ways. First, we observe a monthly panel of spending, so we can track how the same household discretely changes their spending around events such as UI benefit exhaustion. Second, most work on spending during unemployment relies on surveys which are only able to cover a few categories and suffer from recall bias. Using de-identified bank records drawn from the universe of households with a Chase bank account, we are able to construct high-frequency estimates of spending by category during unemployment.

2.1 Analysis Sample

Our analysis sample is drawn from the 27 million households with a checking account in the JPMCI data. The unit of observation is household-by-month, from September 2013 through June 2016. We study households who did not receive any UI payments in 2013 and received at least one month of UI benefits between January 2014 and June 2016.

We restrict the sample to the 20 U.S. states where Chase has physical branches and UI benefits are paid by direct deposit.[7] From January 2014 onward, the potential UI benefit duration was six months in 17 of these states and less than six months in Florida, Michigan, and Georgia. Our primary analysis sample uses the 17 states that offered exactly six months of benefits. However, we use spending data from Florida in Section 3.3.1, and we use all 20 states in our cross-state analysis in Section 3.3.2.

In the JPMCI data we observe UI recipients who were paid by direct deposit. Nearly all states offer UI recipients a choice between receiving benefits by direct deposit or prepaid debit cards. Data from Saunders and McLaughlin (2013) show that the share of UI recipients who are paid using direct deposit in our 20 states is 45 percent. We show below in Section 2.3 that UI recipients paid by direct deposit have similar income to all UI recipients.

Our primary analysis sample imposes two further sample restrictions. The first restriction is motivated by the fact that 28 percent of banked households have checking accounts at multiple banks (Welander 2014) and our inability to capture spending by UI recipients out of non-Chase bank accounts. To limit the sample to households who use Chase as their primary bank account, we restrict the analysis sample to households with at least five monthly checking account "outflows." An outflow is any debit from a checking account including a cash withdrawal, an electronic payment, a paper check, or a debit card transaction. We select households who have five outflows in each month from three months prior

[7] Online Appendix Figure 1 shows a map of the 20 states and Appendix Table 1 provides basic summary statistics on the sample.

to their UI spell to three months after their UI spell. This criterion retains 65 percent of household-months and is conservative in that it probably drops some households who do bank primarily with Chase. (Our key empirical result that spending drops sharply at benefit exhaustion does not depend upon this sampling restriction.)

The second restriction is to limit the sample to households with a single contiguous UI spell. This restriction is necessary because we define UI exhaustees based on the number of weeks of benefits received and benefit duration measured in weeks is only available in the JPMCI data for contiguous UI spells. This criterion retains 82 percent of remaining household-months for a final sample size of 182,361 households and 5.4 million household-months. See online Appendix B for additional details on how we measure UI receipt and exhaustion. As a robustness check, we verify that the spending drop at the onset of unemployment is similar for households with one UI spell and households with multiple UI spells.

In some cases, we observe spending for a Chase account belonging to a household member (e.g., a spouse) who is not receiving UI, which is useful for understanding whether households substitute spending toward other bank accounts during unemployment. The JPMCI data can be aggregated in two different ways to reflect two definitions of households. The first, more narrow definition considers a household to include all bank accounts that are administratively linked. Most people link their bank accounts administratively when they get married, making it easy for spouses to access each others' accounts (TD Bank 2014). The second, broader definition includes other adults in the family whose accounts are not administratively linked. If two Chase customers married one another and did not administratively link their bank accounts (such that they constituted two separate households under the narrow definition), the JPMCI data would still usually classify them as part of a single household under the broader household definition. We use the narrow definition in most of our analysis (i.e., if according to this definition a household has multiple checking accounts with Chase, we analyze the sum of debits and credits across these multiple accounts), but we analyze our most important results using the broad definition as well. In any case, the definitions coincide for 79 percent of households.

2.2 Variables: Constructing Spending, Income, Assets, and Liabilities

Spending – Spending is measured from debit and credit card transactions, cash withdrawals, and electronic transactions captured through the bank account. We divide spending into "strict nondurables," "other nondurables," and "durables" using the the standard expenditure taxonomy for survey data (Lusardi 1996). Examples of "strict nondurable" spending categories according to this taxonomy include groceries, food away from home, fuel, utilities, and haircuts. Examples of "other nondurables" include clothing, medical copays, and payments at drugstores. We provide a comprehensive description of our spending categorization procedure in online Appendix C and provide a summary here.

We allocate spending via credit card, debit card, and electronic payment between these three categories using transaction metadata. To categorize a debit or credit card purchase, we crosswalk Merchant Category Codes (MCCs) that merchants report to the Internal Revenue Service on what types of goods and services they sell to the spending categories from Lusardi (1996) mentioned above. We then use Lusardi's taxonomy to categorize each MCC as strict nondurable, other nondurable, or durable. For a subset of electronic payments, the JPMCI data have a categorization of the transaction that is indicativeof its economic purpose.

Following the standard approach in the consumption excess sensitivity literature (Lusardi 1996), our analysis focuses primarily on the nondurable subset of spending. This variable–the sum of strict nondurables, other nondurables, cash withdrawals, and miscellaneous nondurables–accounts for 44 percent of checking account outflows.[8] Durable spending accounts for another seven percent of outflows. The remaining outflows we can categorize include debt payments (14 percent of outflows) and transfers to other accounts (six percent of outflows). The final 29 percent of outflows cannot be categorized because they are either paper checks or electronic payments that could not be categorized.[9]

We compare spending levels in the JPMCI data to external benchmarks in order to evaluate the fraction of household spending that is captured by the JPMCI data. For this analysis, we use the crosswalk to the Consumer Expenditure (CE) Survey described above, and we also crosswalk to the Bureau of Economic Analysis' Personal Consumption Expenditures (PCE). We summarize the conclusions in Table 1 and provide additional detail in online Appendix Tables 2, 3, and 4.

Among different types of spending, the JPMCI data provide the best coverage of spending on nondurables. Mean spending on nondurables in JPMCI is 139 percent of the CE Survey benchmark and 66 percent of the PCE benchmark.[10] For comparable durables, JPMCI is 31 percent of CE Survey and 24 percent of PCE. As we discuss in online Appendix C, JPMCI covers nondurables better than other types of spending because these goods are typically paid for using debit cards, credit cards, and cash. This is useful in our context, since good coverage of spending on nondurables is essential for studying high-frequency fluctuations in consumption.

Income – Our analysis focuses on a subset of checking account inflows: paychecks paid

[8]We treat all cash withdrawals as nondurable spending because the Federal Reserve's Diary of Consumer Payment Choice (DCPC) shows that 75 percent of cash is spent on nondurables. See online Appendix C for an explanation of the "miscellaneous nondurables" category, which accounts for five percent of checking outflows.

[9]In a robustness check in online Appendix E.1, we construct a measure of nondurable spending which assumes that all uncategorized spending (including cash) is spent on nondurables in proportion to that payment method's nondurable expenditure share in the DCPC.

[10]It is well known that the CE Survey understates consumption expenditures. Passero et al. (2014) carefully crosswalk CE Survey and PCE expenditure categories and found the ratio of CE Survey to PCE was 0.60 across all categories and 0.77 across comparable categories. Online Appendix Table 2 indicates that the CE Survey's coverage of nondurable categories is even worse relative to PCE, with a ratio of 0.48.

by direct deposit and UI benefits. Paychecks account for 59 percent of checking account inflows prior to UI onset, while UI benefits account for 25 percent of inflows during a UI spell. The remainder of inflows prior to onset are accounted for by transfers from outside savings and investment accounts (12 percent), other income (three percent), government income (four percent) and uncategorized inflows (22 percent). Uncategorized inflows, which are often paper checks, cannot be assigned to income or dissavings.

We construct individual employment histories (also known as "labor market biographies") using an anonymized identifier associated with the employer named on each direct deposit paycheck. To the best of our knowledge, we are the first to construct such biographies for U.S. workers at a monthly frequency using administrative data. We describe the construction of these biographies in more detail in online Appendix D, and validate their use by showing that job separations spike at UI onset and that job starts are highly correlated with UI exit. We use these biographies to measure nonemployment duration after UI benefit exhaustion and to measure the monthly job-finding hazard during unemployment. One weakness of this methodology is that we are unable to detect re-employment when wages are paid using paper checks.

Assets and Liabilities – The JPMCI data do not directly measure total liquid assets or total liabilities, but do contain coarse measures of liquid assets, as well as a measure of Chase credit card liabilities. We use an estimate of the household's total liquid assets based on an internal JPMC statistical model in our heterogeneity analysis. This model uses a combination of checking account activity and third-party data sources to construct an estimate of total liquid assets. For households with Chase credit cards, we observe revolving balances on Chase credit cards, new purchases on Chase credit cards, and credit limits on Chase credit cards.

2.3 Representativeness

Our results in the JPMCI data seem likely to generalize to the broader population of UI recipients. Program rules require UI recipients to have substantial prior labor force attachment and therefore they are very likely to have a bank account. In fact, we calculate that 90 percent of households reporting UI income in the past year in the SCF had a bank account at the time of the survey.

Two pieces of evidence suggest that UI recipients in the JPMCI data have similar pre-separation income to UI recipients overall in the U.S. First, online Appendix Figure 2 shows that state-level monthly UI benefits in the JPMCI sample are very similar to benefit levels in Department of Labor administrative data. Because UI benefits are determined by pre-separation earnings, this implies that the JPMCI sample has similar pre-separation income to all UI recipients. Second, we measure the distribution of household income in the twelve months prior to UI receipt using the Survey of Income and Program Participation (SIPP). In the JPMCI data, we rescale checking account income into pre-tax dollars. The distribution

of income in these two datasets is quite similar, as shown in online Appendix Figure 3 and summarized in Table 1.

Beyond income, Table 1 also shows that JPMCI UI recipients are similar to external benchmarks in terms of age and checking account balances.[11] Mean age is 41.1 years in JPMCI compared to 44.0 years in the SIPP. The median checking account balance in the JPMCI sample is $1,250, which is 83 percent of the SCF balance in the household's primary checking account.[12]

3 How UI Affects Household Behavior

In this section, we show that spending is highly sensitive to UI benefits, both at the onset of unemployment and even at the large and predictable decrease in income arising from UI benefit exhaustion. We also provide category-level evidence suggesting that the large drop in expenditure at exhaustion reflects an actual drop in consumption, that state-level variation in spending tracks state-level variation in the path of income, and that job-finding spikes modestly at UI benefit exhaustion.

3.1 Path of Spending During Unemployment

3.1.1 Methodology

Monthly data enable a more granular view of income and spending during unemployment for a given household than has been previously available. Figure 1 nonparametrically shows the monthly evolution of income and spending around UI receipt. We stratify income and spending groups by the number of months of UI benefits received. For income, we focus on the two components which we measure most reliably: labor income paid by direct deposit and UI benefits. For spending, we use the definition of nondurable spending from Lusardi (1996).

We develop a methodology for summarizing the path of income and spending at a monthly frequency throughout a UI spell. Figure 2 presents this summary of the estimated path for the sample that stays unemployed each month. Our approach takes the set of

[11] In addition, as evidence for external validity, we estimate that the share of U.S. households receiving UI via direct deposit is close to the share of households in the data. Across the U.S., an average of 2.6 million people received UI benefits each week in 2015. We estimate that in an average week in 2015, 0.9 percent of families in the U.S. received UI benefits via direct deposit. In the bank data, the average monthly UI recipiency rate in 2015 was 0.7 percent.

[12] Online Appendix Tables 5 and 6 provide additional statistics comparing income and checking account balances in the two samples. It is not possible to directly compare the total liquid assets in our sample to an external benchmark because we do not have a complete lens on a household's total liquid assets, but a back-of-the-envelope calculation suggests that these two numbers are similar. We estimate that UI recipients in the JPMCI data have a median of $4,088 in liquid assets by multiplying the median checking account balance of $1,250 in the JPMCI data by the ratio of median liquid assets to median checking account balances of 3.27 from the SCF reported in online Appendix Table 6. This is quite similar to Chetty's (2008) estimate that median liquid assets for job losers is $3,336 ($1,763 in 1990 dollars) using the SIPP. The algorithm for rescaling post-tax dollars into pre-tax dollars is described in online Appendix Table 5.

households that are still unemployed each period and compute their change in spending relative to the prior period. This series is the conceptual counterpart of the job-finding hazards commonly computed in the job search literature, which take the set of jobseekers that are still unemployed each period and compute the share who find a job that period. This measure is useful because it enables us to reduce dimensionality in a way that facilitates comparison to the model in Section 4.

We construct Figure 2 as the cumulative sum of a series of one-month changes. Formally, let y be the outcome of interest, either income or consumption. We define i as a household, t as months since first UI receipt, and n as the number of households in the analysis group at time t. Prior to onset, $\bar{y}_t = \frac{1}{n}\sum_i y_{i,t}$ and all future UI recipients are included in the sample. In month $t = 0$, everyone who gets one full month of UI is included, in month $t = 1$, everyone who gets two full months of UI is included, and so on.[13] Each point in the figure is estimated as

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

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

Unemployment duration is measured using UI checks through the sixth month of UI receipt and labor market biographies afterwards.[14] Our choice of an evolving sample is a convenient way to summarize the data, but is not quantitatively important for our empirical results. A spending series where the sample is limited to UI exhaustees yields quite similar results, as shown in online Appendix Figure 4.

The unit of observation in the JPMCI data available to us is the calendar month, and this time aggregation poses an empirical challenge because UI benefits may be exhausted in the beginning or middle of the calendar month.[15] To address this issue, we limit the sample of six-month UI recipients in our figures to households who received their last UI check on the 21st of the month or later. These households experience a sharp drop in UI payments because in one calendar month they receive nearly a full month of UI benefits, while in the next month they receive none at all. As a robustness check, we compare our results for this sample to the results for a two-month exhaustion time window for all UI recipients in online Appendix Figure 4.[16]

[13] Due to variation in the frequency of UI check payments, we implement this screen by sampling households who receive at least one UI check in the subsequent month.

[14] One potential concern about using labor market biographies to infer continued unemployment after UI exhaustion is that this strategy will mislabel as unemployed people who actually found a job, but either had that paycheck deposited to a non-Chase account or were paid by paper check. To ameliorate this concern, we limit the sample to people who are eventually reemployed according to the labor market biography data and obtain very similar results. Relatedly, we show that the spending drop of UI benefit exhaustees is large regardless of continued post-exhaustion unemployment duration. The results of both analyses are shown in online Appendix Figure 4.

[15] In some recent work JPMCI has analyzed daily transaction-level data, but transaction data on the universe of UI recipients were not available when we conducted this research.

[16] We use this larger sample and analyze the change in a two-month time window at exhaustion to improve

We take two additional steps to clean the spending and income data. First, to eliminate seasonality, inflation, secular trends, and business cycle fluctuations, all results for income and spending are presented relative to a comparison group. The comparison group is households in the analysis sample that did not exhaust UI, analyzed for months when they were employed and not receiving UI.[17] Second, to reduce the influence of outliers, we winsorize each variable at the 95th percentile of positive values for that variable.[18]

3.1.2 Key Result: Spending Drops at Onset and Exhaustion

The top panel of Figure 2 demonstrates that households experience two discrete drops in income during unemployment. First, income drops one month before UI payments begin. This is because it usually takes a few weeks to start receiving UI payments after a job separation. The drop in average income is not entirely concentrated in a single month because of heterogeneity in time between last paycheck and first UI check. Second, income drops by an even larger amount at UI benefit exhaustion.[19] The difference in income drops emerges because households that exit UI in fewer than six months have found a job, while households that exhaust their UI benefits are often unemployed for more than six months. Income does not drop to zero at exhaustion because many households have other transfer income, another member who continues to earn labor income, or capital income.

The bottom panel of Figure 2 provides high-frequency evidence that spending is highly sensitive to income, both at the onset of unemployment and even at the large and predictable decrease in income at UI benefit exhaustion. The path of spending approximately mirrors the path of income, with two discrete drops. Monthly spending drops by 6 percent at the start of unemployment, falls by less than one percent per month while receiving UI, and then drops by 12 percent at UI benefit exhaustion. Because the onset of unemployment is a largely unexpected shock, the drop in spending at onset is not surprising from the perspective of economic theory. In contrast, the sharp drop in spending after the large and predictable income drop at exhaustion is the key empirical fact that we use in Section 4 to differentiate benchmark models of rational forward-looking agents from a class of

precision in some of our analysis below.

[17] Formally, with i as a household, t as a calendar month, and $y_{it,raw}$ as the original data, we define $\bar{y}_t^{control} \equiv \frac{1}{n_t}\sum_i y_{i,t}^{control}$, $\bar{y}^{control} \equiv \frac{1}{nT}\sum_t \sum_i y_{i,t}^{control}$, and the analysis variables in this paper are $y_{it} = y_{it,raw} - \left(\bar{y}_t^{control} - \bar{y}^{control}\right)$.

[18] Online Appendix Table 7 shows that our winsorization procedure leads to conservative estimates of the spending drop at UI benefit exhaustion. Without it, our estimates would be even larger.

[19] The fact that this second drop is larger has the counterintuitive implication that UI benefits have a replacement rate of 66 percent. This apparent 66 percent replacement rate is larger than typical statutory UI pre-tax replacement rates, which are around 45 percent in the U.S. Two factors explain nearly all of the gap in measured replacement rates: differential tax treatment of UI benefits and labor income payment method. First, UI checks are not subject to withholding, whereas a typical paycheck will have 7.65 percent deducted in payroll taxes and 15 percent in income tax withheld. Second, we are only able to detect labor income paid by direct deposit in the JPMCI data; we have calculated using the SCF that about 15 percent of labor income is paid by paper checks and pre-paid debit cards rather than by direct deposit. After adjusting for these two issues, our income estimates are similar to those reported for a representative sample of UI recipients by Rothstein and Valetta (2017).

behavioral models with present-biased or myopic agents. We report regression coefficients with standard errors in online Appendix Table 8 and explore the robustness of this finding in the next subsection.

3.1.3 Heterogeneity analysis and robustness

To provide further evidence that changes in income are driving changes in spending, we analyze differences in household-level UI replacement rates. We assign households to terciles using the ratio of monthly UI benefits to monthly income prior to unemployment. This measure combines cross-state variation in replacement rates, within-state variation in the UI replacement rate by prior income, and cross-household variation in the share of income accounted for by the household member who lost her job. Unfortunately, because it is difficult to estimate annual household income in our data, our variable is a noisy proxy for true household-level replacement rates. Nevertheless, we show that spending drops more at exhaustion when UI benefits account for a larger share of household income in Figure 3. To analyze variation in monthly UI benefits that is not driven by individual-level heterogeneity and is less vulnerable to measurement error, we complement the analysis here with state-level analysis in Section 3.3.

A key determinant of the magnitude of the spending drop is a household's assets prior to the onset of unemployment. We again assign households to terciles, this time using the ratio of the estimated liquid assets measure described in Section 2.2 to monthly spending prior to unemployment. Figure 3 also shows that spending drops more at onset and at exhaustion for low-asset households, consistent with Johnson et al. 2006. It appears that households that have liquid assets use them as a buffer to avoid cutting spending during unemployment. This empirical finding provides motivation for a model of agents with heterogeneous patience in Section 4.2.

We also examine heterogeneity in the drops at onset and exhaustion using eight other covariates:checking account balances, joint checking accounts (a coarse proxy for marriage), employment of other household members, exhaustion date, income, age, mortgage payments (a proxy for home ownership), and having a Chase credit card (online Appendix Tables 9 and 10, as well as online Appendix Figure 5).

The drop documented above may not reflect a drop in total household nondurable spending because of spending on credit cards, spending from non-Chase bank accounts, uncategorized payments, and selection into direct deposit of UI benefit receipt. In this section we investigate the first two concerns, while we discuss the two latter concerns in online Appendix E.1, where we show that our findings are robust to alternative assumptions about the composition of uncategorized bank account outflows (online Appendix Figure 6) and are likely to generalize to most UI recipients who do not receive direct deposit of UI benefits (online Appendix Figure 7).

One potential concern is that households might be substituting spending to their credit

cards, but we find no economically significant use of credit cards to smooth consumption during unemployment. On average, revolving balances on Chase credit cards rise by about $20 per month during UI receipt, as shown in online Appendix Figure 8. Put otherwise, new credit card borrowing finances less than 0.5 percent of monthly consumption during unemployment. Although our estimate does not capture borrowing on non-Chase credit cards, our finding of limited credit card borrowing during unemployment is consistent with three other studies that capture utilization across all credit cards. The most comparable estimate is Herkenhoff et al. (2016), who find a revolving debt increase of $28 per month using a sample of credit bureau records matched to employment records. Even more striking, other researchers have found *decreased* credit card borrowing (Bethune 2017) associated with unemployment. We think that understanding the limited use of credit card spending during unemployment is an important area for future research.

A second potential concern is that the drop at exhaustion reflects substitution of spending to other bank accounts, but we do not find substitution to accounts of other household members (e.g., spouses). In Section 2.1, we explain how the JPMCI dataset is able to capture spending for two customers who bank with Chase and form a household unit without administratively linking their bank accounts. Spending out of other accounts is roughly constant at benefit exhaustion, with an increase of $40 (see online Appendix Figure 9).[20] Because only about one-quarter of households have accounts at multiple institutions, the average drop in household spending is $252, rather than $263, after incorporating spending from outside accounts.[21] Nevertheless, it would be useful in future work to replicate our analysis in a dataset from a financial aggregator where we can be confident of observing all outside bank accounts and credit cards.

Another piece of evidence that may address concerns about missing spending data is that at a low frequency, our estimates of the average drop in spending during unemployment are quantitatively in line with prior estimates using survey data from the U.S. For example, Gruber (1997) reports that food spending falls by 6.8 percent.[22] Although Gruber's empirical specification does not generate a standard error for this statistic, Hendren (2017) reports

[20] One might have expected that households would cut spending out of *every* bank account at exhaustion, not just the account which stopped receiving UI payments. The absence of a decline in spending in other household accounts could arise from several channels, including a convenience motive or a form of mental accounting. The convenience hypothesis is that cutting spending out of both accounts would eventually require a transfer from the second account to the primary account to cover the continued spending out of the account with fewer inflows, so it is more convenient to just cut spending out of one account. The mental accounting hypothesis is that when someone has a decline in income, she cuts outflows from the account with diminished income.

[21] In percentage terms, the two-month drop in spending falls from 10.7 percent to 8.7 percent, as shown in online Appendix Table 12. See online Appendix Table 11 for a similar calculation of how incorporating outside bank accounts affects the drop in spending at onset.

[22] Several authors have replicated this estimate: Chetty and Szeidl (2007), Kroft and Notowidigdo (2016), East and Kuka (2015), Chodorow-Reich and Karabarbounis (2016), Saporta-Eksten (2014), and Hendren (2017). For examples of non-U.S. estimates, see Browning and Crossley (2001) and Kolsrud et al. (2018) for estimates for Canada and Sweden, respectively. These estimates are difficult to compare to ours because UI benefits are more generous in these countries.

a comparable estimate with standard error of 0.5 percent. When we construct comparable estimates to Gruber, we find that spending drops by 6.4 percent on all nondurables and 6.2 percent on food.[23] Both of these estimates have a standard error of 0.1 percent.

3.2 Spending Categories

In this section we leverage the detailed spending categories available in the JPMCI data to document that the drop in spending at exhaustion appears to reflect a change in a household's actual consumption bundle from the prior month. Although we have shown that *expenditure* drops sharply at exhaustion, a remaining concern is that this would not necessarily imply that *consumption* drops sharply if the categories of expenditures that are falling are those with weak links between the timing of expenditure and the timing of consumption flows. Table 2 decomposes the drop in spending into 27 mutually exclusive and comprehensively exhaustive categories.[24] Categories linked to necessities exhibit sharp drops. For example, grocery spending drops by 16 percent, medical out-of-pocket spending drops by 14 percent, and drug store purchases drop by 15 percent. This evidence suggests that consumption falls at exhaustion; however, one important limitation of using bank account data to measure consumption is that we do not capture in-kind transfers such as free food or medical care.

The sharp drop in grocery spending at exhaustion probably reflects a deterioration in diet quality. Aguiar and Hurst (2005) find that unemployment decreases both spending and consumption. Specifically, they compare the diets of employed and unemployed people, controlling for a wide variety of observables, and report a similarly-sized gap in spending on groceries between the employed and unemployed (9–15 percent) to the drop we see within the unemployed at exhaustion. They estimate that a drop of this size causes a nine percentage point decrease in the share of households consuming any fresh fruit and a five percentage point increase in the share of households consuming any hot dogs or processed lunch meat.

In contrast, households appear to prioritize their most important financial commitments at exhaustion and these categories show much smaller declines. Table 2 shows that the drop in spending is quite small for utilities, insurance payments, installment debt, and credit card bills. (In online Appendix E.2 and online Appendix Figure 10, we analyze the evolution of spending on durables in more detail.) There is little evidence to suggest that benefit exhaustion does immediate damage to a household's long-term financial health. These empirical results are consistent with the presence of consumption commitments, as suggested by Chetty and Szeidl (2007).

[23] It is ambiguous what time horizon respondents use when they describe their food consumption in the PSID. The estimates above assume the reference period is unemployment onset. Alternatively, if we assume the reference period is an annual time horizon then our comparable estimates are 6.9 percent on nondurables and 4.3 percent on food (online Appendix Table 13).

[24] Online Appendix Table 14 provides comparable statistics for the onset of unemployment.

To complement our analysis of outflow categories, we also analyze inflow by category throughout the spell. Other than UI benefits, there is little increase in other types of checking account inflows that could offset the loss of labor income (online Appendix Figure 8). In addition, checking account outflows modestly exceed inflows during unemployment, such that the average cumulative drawdown of assets among households that exhaust UI benefits is $1,389 from the start of unemployment through the exhaustion of benefits. This evidence is broadly consistent with Kolsrud et al's (2018) finding that households in Sweden show limited drawdown of liquid assets during unemployment.

3.3 State-Level Policy Variation

To provide further evidence that spending is highly sensitive to UI benefits, we exploit the cross-state policy variation available in the JPMCI data. We first analyze event studies of spending in one state with unusually limited UI benefits (Florida) and one state with unusually generous UI benefits (New Jersey). Then, we use all 20 states in our sample to estimate an MPC out of UI benefits.

3.3.1 Event studies in Florida and New Jersey

Figure 4 shows that Florida has unusually limited UI benefits. First, UI benefits lasted four months in Florida from January 2014 through June 2015, which is the time period that we analyze for Florida UI recipients. Second, the drop in household income at onset in Florida is larger than in states where benefits last for six months ("six-month states"). Two factors contribute to this drop: UI benefits replace 40 percent of pre-onset income in Florida, compared to 45 percent in six-month states, and UI recipients in Florida were responsible for 72 percent of pre-onset household income, compared to 64 percent in six-month states.[25] Altogether, household income falls by 32 percent at onset in Florida compared to 19 percent in six-month states.

Figure 4 also shows that spending in Florida is much lower during unemployment, and drops sharply at month four, coincident with the expiration of UI benefits. Spending drops more at the onset of unemployment in Florida than in six-month states. Spending drops again at UI benefit exhaustion after four months in Florida and is stable after six months of unemployment. In contrast, the six-month states exhibit the opposite pattern: spending is stable after four months and only drops sharply after benefits run out at six months. The second drop indicates that the drop in spending that we observe coincident with benefit exhaustion for households in six-month states is caused by the end of benefits, rather than by some external negative shock which happens to occur after six months of unemployment.

[25]Florida has strict eligibility criteria for claiming UI benefits, such that only 11 percent of unemployed households received UI benefits in 2015, compared to 27 percent nationally (McKenna and McHugh 2016). Florida's strict UI eligibility rules may explain why UI recipients there are more likely to be the household breadwinner.

At the opposite end of the spectrum from Florida, New Jersey has unusually generous UI benefits, as shown in the bottom panels of Figure 4. The drop in household income is small because UI benefits replace 49 percent of pre-onset income (rather than 45 percent on average) and because UI recipients were responsible for 60 percent of pre-onset household labor income (rather than 64 percent on average). The most unique aspect of New Jersey's benefit structure from a research perspective is that it is one of the few states which does not have a required waiting period before beginning UI payments. This absence of a waiting period leads to the unusual phenomenon that many workers are still receiving paychecks when their first UI check arrives. This induces a sawtooth pattern in average income, which falls as some workers become unemployed, rises by 10 percent in the month with the extra check, and then declines to a stable level thereafter.

Spending has the same sawtooth pattern as income among UI recipients in New Jersey, suggesting that much of the extra check is spent immediately even though households know their income is likely to fall sharply in the following month. This sawtooth pattern is useful for two reasons. First, it indicates that the sensitivity of spending to income is widespread among UI recipients, rather than being limited to exhaustees. Second, it holds for income increases, rather than just decreases, which echoes the prior literature on excess sensitivity that has typically studied income increases.

3.3.2 Marginal propensity to consume (MPC) out of UI benefits

To estimate an MPC, we move from analyzing Florida and New Jersey to exploiting all the cross-state variation in the JPMCI data. We show how the spending change in each state from the calendar month before UI payments begin ($t = -1$) to the first calendar month of full UI benefits ($t = 1$) varies with the UI household replacement rate, defined as the ratio of UI benefits to pre-onset labor income paid by direct deposit. We find that the sensitivity of spending to UI benefits at onset that we documented in New Jersey holds across all 20 states.

We estimate that each additional dollar at the start of unemployment leads to 27 cents of additional spending on nondurables. The top panel of Figure 5 shows the UI household replacement rate on the x-axis and the spending drop by state on the y-axis. The plot shows a strong positive correlation. The plot includes a fit line from a state level regression $\Delta Spend = a + \beta UI$ where the β coefficient indicates that a 10 percent increase in the UI household replacement rate leads to a 3.4 percent spending increase on nondurables. Evaluated at the mean value for UI benefits and spending on nondurables, this estimate implies an MPC of 27 cents with a standard error of seven cents. This empirical design captures the *immediate one-month* spending response to additional UI benefits. We focus on the high-frequency response at onset as best-suited for calculating an MPC because households have time to prepare for benefit exhaustion, though we show in the bottom panel of Figure 5 that a similar pattern holds across states at exhaustion.

Our headline estimate of a nondurables MPC of 27 cents captures only a portion of the spending response to additional UI benefits for two reasons. First, because our empirical design captures the immediate spending response, it will understate the *nondurables* MPC if UI causes households to raise their nondurables spending prior to receipt of the first UI check or if UI benefits in month one raise spending in subsequent months. Second, there is likely some spending response *outside of nondurables*. There are two back-of-the-envelope ways to estimate the total spending response shown in online Appendix Figure 12. For every dollar of UI, we estimate that total bank account outflows rise by 83 cents (this is an upper bound on the MPC because it includes transfers to savings accounts). Alternatively, food spending rises by 10 cents; under homothetic preferences and using a 13 percent food expenditure share from the CE Survey, this implies an MPC of 77 cents.

Although differences in the UI household replacement rate are not exogenous, an institutional feature of the UI system offers a placebo test to evaluate the potential bias in this identification strategy. There is usually a one-month lag between the onset of unemployment and the receipt of UI benefits, as documented in Figures 1 and 2. Online Appendix Figure 12 shows that spending drops an equal amount in high- and low-benefit states before UI payments begin. A cross-state regression of the spending drop from three months before UI receipt to one month before UI receipt on the UI household replacement rate is not economically or statistically different from zero. The fact that spending diverges in high- and low-benefit states only after UI payments begin suggests that our empirical strategy captures the causal impact of UI benefits on spending.

The MPC out of UI benefits is of interest to three distinct literatures. First, there is a literature which argues that if unemployed households have a higher MPC than employed households then this means that increasing UI generosity is an effective macroeconomic stabilization tool (McKee and Verner 2015; Kekre 2017).The high MPC for unemployed households that we estimate is consistent with an MPC for unemployed households that is substantially higher than for employed households.

Second, our MPC estimate helps distinguish among alternative explanations for why spending falls at the start of unemployment. Browning and Crossley (2001) describe three factors which could contribute to the spending drop at the start of an unemployment spell: a temporary drop in cashflows, a decrease in work-related expenses, and a permanent income loss. The cross-state variation provides an empirical strategy for assessing the role of the temporary income loss. The unemployed in states with high UI benefits likely have a similar change in time budget and loss in permanent income to the unemployed in states with low UI benefits. However, for each additional lost dollar of household income, UI recipients in low benefit states cut spending by an additional 27 cents. The total drop in income at the start of unemployment is \$470. Using the MPC estimate to extrapolate linearly, we estimate that the temporary drop in cashflows should cause spending to fall by $(0.27 \times \$470 =)$ \$127. The temporary cashflow channel can explain nearly all of the \$158 drop in nondurable

spending at the start of unemployment. See online Appendix E.3 for further analysis of the role of home production in explaining the drop in spending at the onset of unemployment.

Third, there is a social insurance literature which interprets the MPC as a measure of the consumption-smoothing benefits of UI. Although to the best of our knowledge our estimate is the first for the MPC on *nondurables* out of UI benefits, the prior literature has been able to estimate the MPC on *food*. Our estimated response of food spending to UI benefits is within the range of these prior estimates. Gruber (1997) and East and Kuka (2015) estimate that a 10 percentage point increase in the UI replacement rate raises food spending by 2.7 percent and 1.1 percent respectively using the PSID.[26] Our comparable statistic is that a 10 percentage point increase in the UI replacement rate would raise a household's food spending by 2.7 percent.

3.4 Job-Finding

We use the JPMCI data to estimate the monthly job-finding hazard of UI recipients, which is of interest for two reasons. First, we use the job-finding hazard as an empirical target in estimating our structural model in Section 4. Second, to the extent that there is a spike in the probability of finding a job at benefit exhaustion, it means that not finding a job in this month is particularly bad news. The arrival of this bad news could potentially explain the drop in consumption documented above.

Data limitations have prevented prior work from being able to estimate the job-finding hazard around benefit exhaustion in the U.S. Besides JPMCI, there is no other U.S. dataset we are aware of in which one can (1) observeUI receipt, (2) observeemployment, and (3) observe these outcomes at a monthly (or higher) frequency.[27] In contrast, most European UI systems collect this data and so job-finding around benefit exhaustion has been calculated in at least eight European countries.

We construct a job-finding hazard using the paycheck data described in Section 2.2.

[26]These two papers do not report a MPC out of UI benefits. Their analysis focuses on unemployment by household heads rather than by all household members. UI replaces a larger fraction of *household* income when the head is unemployed. Because the JPMCI sample includes all UI receipt (and not just household heads), the estimated MPC on food in the JPMCI sample may be larger than the implied MPC based on the PSID estimates. Empirical results from McKee and Verner (2015) and DiMaggio and Kermani (2016) also suggest a large MPC out of UI benefit levels. However, Browning and Crossley (2001) estimate that a 10 percentage point increase in the UI replacement rate raises total spending by only 0.8 percent.

[27]The Survey of Income and Program Participation has (1) and (2), the Current Population Survey (CPS) and the Panel Study of Income Dynamics (PSID) have (2) and (3), and UI administrative data have (1) and (3). Prior work using U.S. data suggests there may be a spike in job-finding at UI benefit exhaustion. Kroft et al. (2016) and Rothstein (2011) analyze job-finding for unemployed workers in the CPS and find a spike of about 8 percent in the job-finding hazard after six–seven months of unemployment. Meyer (1990) shows a spike in UI exit hazards from 7.5 percent two weeks prior to benefit exhaustion to 16 percent in the week prior to benefit exhaustion using UI administrative data. Katz and Meyer (1990) show spikes in the job-finding hazard at 26 and 39 weeks among households who received UI benefits at some point in their unemployment using the PSID during a time period when most respondents were eligible for either 26 or 39 weeks of benefits. However, the PSID does not include high-frequency data on UI benefit receipt nor each individual's potential benefit duration.

When used for measuring job starts, paycheck data contain false positives, which may arise from an employer changing the transaction description on their paychecks, and false negatives, which arise from people paid by paper check. To address these issues, we use the pattern of job starts for people who exit UI prior to benefit exhaustion, and hence are highly likely to exit to a job, to infer the job-finding hazard at UI benefit exhaustion for those who may or may not be exiting to a job. We describe our methodology for constructing a job-finding hazard in more detail in online Appendix D.

We show in Figure 6 that job-finding spikes modestly at UI benefit exhaustion, but we find that the spike is far too small quantitatively to explain the observed drop in consumption at exhaustion. Note that in contrast to the consumption figures where the x-axis measures months since first UI check, the x-axis in Figure 6 measures months since last potential UI check. This allows us to circumvent challenges raised by the monthly time aggregation problem discussed in Section 3.1.1 and provide a clean measure of the spike in the job-finding hazard in the last potential month. Our key finding is that the hazard rate is 18.8 percent with two months of UI benefits remaining and 25.9 percent in the last month when UI benefits are paid, which is about 38 percent higher than the previous month. We report bootstrap standard errors and show that this increase in the hazard is statistically significant. Although the spike we document is economically and statistically significant, we document below in the model section that the spike would need to be more than twice as large in order to rationalize the observed drop in consumption for those failing to find a job in the last month.

Our estimate of a 38 percent spike in the monthly job-finding hazard at benefit exhaustion falls in the middle of the range of estimates from the prior literature studying European countries. For example, Card et al. (2007) estimate a 17 percent increase at 30 weeks and 25 percent increase at 20 weeks in Austria; DellaVigna et al. (2017) estimate a 36 percent increase at 270 days in Hungary; and Schmieder and von Wachter's (2016) figures imply a 46 percent increase at 18 months and 61 percent at 12 months in Germany. Similar to this prior research, we also find that the job-finding hazard gradually falls after UI benefit exhaustion. This pattern might be consistent with unobserved heterogeneity (Paserman 2008), duration dependence (Kroft et al. 2016) or reference dependence (DellaVigna et al. 2017). We analyze a model with unobserved heterogeneity in job search costs in Section 4.

Household search behavior appears highly sensitive to cross-state variation in the generosity of UI benefits. We document this by comparing job-finding hazards in Florida and New Jersey to job-finding hazards in our baseline sample. This parallels our analysis of spending in these two states in Section 3.3.1. As shown in online Appendix Figure 13, the job-finding hazard in Florida is substantially higher than in more generous six-month states in the early months of unemployment and spikes at four months, coincident with the earlier UI benefit exhaustion. In contrast, the same figure shows that the job-finding hazard in New Jersey, which has some of the most generous UI benefits, is significantly lower during

UI receipt but then has a larger spike in the month prior to UI benefit exhaustion when households are about to lose these more generous benefits.

4 Positive Implications for Models of Consumption

We use our empirical setting as a new test to distinguish between the two kinds of theories economists have developed to explain the excess sensitivity of spending to predictable income *increases*: rational (or near-rational) models with forward-looking agents and behavioral models with present-biased or myopic agents. Although a variety of rational models can explain the large response to predictable income increases studied in other contexts, we show that these models cannot explain the large response to predictable *decreases* in income at UI benefit exhaustion. In contrast, a model with present-biased or myopic agents is able to match the JPMCI data.

We describe the setup of our baseline model, which combines the canonical buffer-stock model of consumption in Carroll et al. (2018) with a model of endogenous job search. Agents choose their level of consumption each month, c_t, and their job search effort if unemployed, s_t, to maximize their expected discounted flow of lifetime utility. We assume agents have diminishing marginal utility over consumption $u(c_t)$ and that exerting search effort is associated with a strictly increasing and convex disutility cost $\psi(s_t)$, as in DellaVigna et al. (2017). Agents earn a monthly return of R on their beginning of month assets a_t. The only risk to income z_t comes from unemployment; this risk is partially mitigated by UI benefits, which expire after seven months.[28] Income follows a Markov process Π based on exogenous separations from employment and endogenous job search during unemployment. The agent's problem in month t can be written as

$$\max_{\{c_t, s_t\}} u(c_t) - \psi(s_t) + \mathbb{E}\left[\beta \sum_{n=1}^{T-t} \delta^n (u(c_{t+n}) - \psi(s_{t+n}))\right] \tag{3}$$

$$\text{subject to } c_t + a_{t+1} = Ra_t + z_t \tag{4}$$

$$c_t \geq 0 \tag{5}$$

$$a_{t+1} \geq \underline{a} \tag{6}$$

$$Ra_T + z_T - c_T \geq 0 \tag{7}$$

where δ is the exponential discount factor, β is the quasi-hyperbolic discount factor, $u(c) = \frac{c^{1-\gamma}}{1-\gamma}$, $\psi(s) = k\frac{s^{1+\xi}}{1+\xi}$ when unemployed and $\psi = 0$ when employed, z_t evolves according to transition matrix $\Pi(s_{t-1})$, T is the number of months in the agent's life, and $\underline{a}$ is the

[28] We analyze UI recipients eligible for six months of benefits. In Section 3.1.2, we documented that the decline in household 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 last seven months rather than six months.

borrowing limit. The preference parameter γ reflects risk aversion over consumption and $1/\xi$ is the elasticity of search effort with respect to the value of effort. The model's setup nests the canonical rational forward-looking representative agent model with $\beta = 1$ and allows for naive present-bias behavior in the $\beta < 1$ case. We allow for unobserved heterogeneity in job search costs k. In some parameterizations, we also allow for unobserved heterogeneity in β. We solve the household problem recursively by induction.

Where possible, we use data from the JPMCI sample to parameterize the economic environment. We normalize monthly income to 1.00 in the employed state. Household income is 0.83 while receiving UI benefits and 0.54 after UI benefit exhaustion in the JPMCI data.[29] For other aspects of the environment and preferences, we draw on external sources described in online Appendix Table 15. We estimate the remaining parameters of the model (denoted collectively by θ) which best fit the spending and job search data from Section 3. The moments $\hat{m}$ are the average monthly consumption levels in Figure 2 and job-finding hazards during unemployment in Figure 6. The estimator chooses the parameters

$$\hat{\theta} = \arg\min_{\theta} GOF(\theta) = (m(\theta) - \hat{m})' W (m(\theta) - \hat{m}) \tag{8}$$

to minimize the goodness of fit (GOF) measure. The weighting matrix W is the inverse of the estimated variance-covariance matrix of the sample moments. It is block diagonal, so the GOF measure can be separated into two additive components: consumption GOF and search GOF. The exact moments are reported in online Appendix Table 16.

4.1 Standard Model

We first examine the model's parameter estimates and predictions for rational forward-looking agents by restricting the quasi-hyperbolic present-bias parameter β to 1. We refer to this as the standard model. In this class of models, agents with liquidity constraints build a buffer of assets to protect against future income risk (Carroll 1997). Prior research summarized in Jappelli and Pistaferri (2010) has found that this type of model can explain excess sensitivity of spending to predictable income increases because households are either unable to borrow or unwilling to deplete their buffer in anticipation of future income increases.

The best fit of this model to the data generates consumption parameter estimates for δ and $\underline{a}$ that are broadly consistent with the prior literature. The first column of Table 3 shows the parameter estimates. We estimate a monthly discount factor δ of 0.99, which implies an annual discount factor of 0.89. This estimate is on the lower end of the prior

[29] Our household income concept includes labor income from all household members, capital income, and government transfers. The estimates of household income during unemployment here are larger than those shown in Figure 2 because Figure 2 shows the evolution of direct deposit labor income and UI benefits. In contrast, the estimates here include non-labor income and an adjustment for labor income paid by paper check.

literature but in line with estimates from some lifecycle models (e.g., Berger et al. 2018). Our estimate of the exogenous liquidity constraint $\underline{a}$ of 4.5 months of income is about twice the magnitude of unsecured credit lines reported in the SCF (Kaplan and Violante 2014).

The model also generates job search patterns consistent with the prior literature in two ways. First, the bottom panel of Figure 7 shows that average search effort rises prior to exhaustion and falls afterward, similar to the pattern in the data. This model result echoes prior work by Paserman (2008), who shows that a model with unobserved heterogeneity in job search costs is able to match the well-known empirical pattern of rising and then declining job-finding rates. In our model, search effort rises in anticipation of UI benefit exhaustion, especially among the types with low search costs. Then dynamic selection shifts the composition of the sample toward high search cost types thereby lowering the average job-finding hazard, as shown in online Appendix Figure 14.[30] Second, our model's implied responsiveness of search to changes in UI benefit generosity is close to the median estimates in a literature review by Schmieder and von Wachter (2017), as shown in online Appendix Table 17.

Although the standard model's parameter estimates are consistent with prior work, the model fails to predict the sharp drop in spending at benefit exhaustion. The top panel of Figure 7 visually compares the path of spending in the model to the data. The model does a good job of matching the spending drop at the onset of unemployment. However, with each additional passing month of unemployment, adverse information about the path of future income is gradually revealed. The agents gradually cut spending each month in order to have more assets left in the event that UI benefit exhaustion does occur. Because there is no particular surprise in the month when UI benefits run out, there is no excess drop in spending.

4.1.1 Why the drop at exhaustion is hard to fit with any rational forward-looking model

The results above show that one reasonable model parameterization cannot generate a drop in spending at exhaustion. In this section, we explain why the consumption drop at exhaustion is difficult to fit with *any* rational or near-rational model. We describe three rational theories which have been used in the consumption literature to explain excess sensitivity—liquidity constraints, near-rationality and home production. In each case we show how the behavior in our setting is inconsistent with these potential explanations, leaving us to appeal instead to behavioral explanations with present-biased or myopic agents. In online Appendix F.1 and online Appendix Figure 16, we show that two other theories based on prior labor economics research—permanent income loss and optimism about job search prospects—are also unable to account for the sharp drop in spending at exhaustion.

[30] In online Appendix Figure 15, we show that a model without unobserved heterogeneity in search costs is unable to match these dynamics and instead, contrary to the data, features rising search effort throughout the unemployment spell.

The discussion in this section clarifies why our empirical setting is particularly useful for distinguishing between the rational and behavioral explanations that have previously been used to explain excess sensitivity.

One popular rational forward-looking theory used to explain excess sensitivity is liquidity constraints, but liquidity constraints cannot explain why agents fail to save. In models with liquidity constraints, spending increases in response to predictable income *increases* because borrowers near their liquidity constraints are unable to smooth consumption by increasing borrowing before they receive the income (Gourinchas and Parker 2002, Kaplan and Violante 2014). However, agents do not need to borrow to prepare for a predictable income decrease such as UI exhaustion, they only need to save. Hence, although these theories can explain why spending would rise after the predictable arrival of a tax rebate, they cannot explain why spending would fall at the predictable *decrease* in income at UI benefit exhaustion.

Our model's results from a corner case further illuminate why the liquidity constraints theory cannot explain the behavior we observe. We examine the behavior of a representative agent who is completely liquidity constrained, with no ability to borrow ($\underline{a} = 0$) and no assets at the onset of unemployment ($a_0 = 0$). We find that even this agent will gradually cut her spending in anticipation of benefit exhaustion, as shown in online Appendix Figure 17.[31] To smooth consumption around benefit exhaustion, agents only need a saving technology to build up a buffer of assets to prepare for the drop in income. Households receiving direct deposit of UI into their bank account have such a technology by construction and will save to prepare for exhaustion.

Two other prior explanations for excess sensitivity with forward-looking agents are near-rationality and home production, but these theories are also not applicable to this setting. One account of the spending response to tax rebates is that it reflects nearly-rational behavior, in the sense that the decision to spend a large fraction of a tax rebate has little welfare loss relative to the optimal spending path (Fuchs-Schuendeln and Hassan 2016). The key prediction of the near-rational theory is that the estimated MPC is *smaller* when the change in household income is larger (Kueng 2015). This prediction does not hold in the JPMCI data, where the MPC is slightly *larger* among households where UI accounts for a larger share of income (see Figure 3 and Appendix Table 10). A final strand of the literature has suggested that the drop in spending at retirement may be attributable to increased home production (Aguiar and Hurst 2005).[32] Because benefit exhaustion is a change in income without a change in the agent's time budget, the observed drop in spending cannot be explained by a change in home production.

We conclude that none of these prior explanations based on rational forward-looking

[31] In the Kaplan and Violante (2014) model, the availability of a high-return illiquid asset with a transaction cost leads agents to hold relatively few liquid assets. It seems likely that an agent in their model with few liquid assets who did not access her illiquid asset would behave similarly to the agent with zero assets at onset in our model, while an agent who did access her illiquid asset would not dramatically cut spending at exhaustion.

[32] Stephens and Toohey (2018) re-examine this result using a wide range of dietary intake surveys.

models and labor market expectations are able to explain excess sensitivity to a predictable drop in income at UI benefit exhaustion. In the next section, we explore whether the second class of models previously used to explain excess sensitivity, those with present-biased or myopic households, are better able to explain the empirical behavior we observe.

4.2 Behavioral Model

Can a model with behavioral households explain the sharp drop in spending at UI benefit exhaustion? By behavioral, we mean models where households are either much more impatient than standard exponential discount rates would imply, or are not fully forward-looking. A classic example of this class of models is the hyperbolic discounting model of Laibson (1997). We begin our examination of these models by relaxing our assumption that $\beta = 1$ in equation (3), and allow β to be estimated to capture naive present-bias behavior.

We find that a model with only present-biased agents is unable to fit the data *quantitatively*, but that a present-biased agent helps to *qualitatively* capture the behavior we observe. The results from allowing β to be estimated on a grid between 0 and 1 are shown in column 2 of Table 3. We find that the parameter estimates are unchanged, i.e. that even when we relax the assumption that $\beta = 1$, the model that best fits the data is one with no present bias such that the estimated β is also equal to 1. To see why present-bias cannot explain the data by itself, we consider a corner case where the household has a degree of present-bias similar to that estimated in prior work (Laibson et al. 2007; DellaVigna et al. 2017). For this case, we set $\beta = 0.5$ and leave the other parameters unchanged from the standard model parameterization in column 1 of Table 3.

A household that is highly present-biased and liquidity-constrained does indeed have a sharp drop in spending at exhaustion and little drop prior to exhaustion, as shown in online Appendix Figure 17. This behavior can be thought of as a micro-founded version of Campbell and Mankiw's (1989) hand-to-mouth consumer in their "spender-saver" model, because she sets consumption very close to income each period.[33] However, both the drop at onset and the drop at exhaustion are far too large relative to the average path we observe in the JPMCI data. These counterfactually large drops are why we find little present-bias when we estimate our standard model. Nevertheless, the path of spending we find in this corner case suggests that the behavior we observe is qualitatively well matched by this type of model, generating a sharp drop at exhaustion with little drop prior to exhaustion in a way that none of the rational forward-looking models explored in the prior section were able to generate.

Since we find that a model with only present-biased households cannot quantitatively fit the data, but that including such households is crucial for generating a qualitatively sharp

[33]We evaluate the model for the present-biased agent starting 12 months before unemployment. This is enough time for the agent to exhaust her initial assets, so that she has no assets remaining and sets consumption equal to income before unemployment begins.

drop at exhaustion, we next examine a model where households may differ in their present-bias parameter. Specifically, we add two more parameters to the standard model by allowing for heterogeneity in the β parameter, which can be either β_{low} or β_{high} and estimating the share of each in the data. This model thus has a total of four types: low β with low search cost, low β with high search cost, high β with low search cost, and high β with high search cost. Our parameter estimates are shown in column 3 of Table 3. We estimate substantial present-bias among some consumers, with $\hat{\beta}_{low} \approx 0.5$. For the remaining consumers, we estimate $\hat{\beta}_{high} \approx 0.9$, but we obtain very similar results if we re-estimate the model under the assumption $\beta_{high} = 1$ (see online Appendix Table 18) and therefore refer to the high β consumers as "standard agents."

In contrast to models with only patient forward-looking households, we find that the model with a mix of present-biased agents ("spenders") and standard agents ("savers") generates a predicted spending path that closely tracks the empirical path of mean spending during unemployment. This mix of agents echoes the "spender-saver" model in Campbell and Mankiw (1989). Figure 8 shows that this model is able to match both the gradual drop in spending during UI receipt as well as the sharp drop in spending at UI benefit exhaustion.[34] Allowing for some agents with low β enables the goodness-of-fit measure for consumption to fall from 350 for the standard model to 99 in the model with two types of β. The intuition for why the model fits well is that the mean path of spending in the data is well-approximated by four line segments: (1) a sharp drop at onset, (2) a gradual decline during UI receipt, (3) a sharp drop at exhaustion, and (4) a gradual decline after exhaustion. The low β agents enable the model to match the sharp drops (1) and (3), while the high β agents enable the model to match the gradual declines (2) and (4).

We generate hand-to-mouth behavior using naive present-bias–rather than a low exponential discount parameter δ–because of the plausibility of the parameter values needed to generate a discrete drop at exhaustion. Even a monthly discount factor δ of 0.9 (corresponding to an annual discount factor of 0.28) generates both an anticipatory drop before exhaustion and a further drop in spending after exhaustion. This is a high degree of impatience relative to the prior literature: nine times more impatient than the most impatient agent in Carroll et al. (2017) and 12 times more impatient than the most impatient agent in Krusell and Smith (1998). Only by lowering δ as far as 0.6 is it possible to generate a drop primarily in the month of benefit exhaustion, with roughly constant spending in the prior and subsequent months. Online Appendix Figure 17 shows these results visually and online Appendix Figure 18 shows the evolution of average consumption and search in a model where some agents have very low δ.

Table 3 column 4 shows that the model with δ heterogeneity has two shortcomings. The model's goodness-of-fit for consumption is 148, which is slightly worse than the goodness-

[34]The path of consumption and search by each of the four types of agents are shown in online Appendix Figure 19.

of-fit estimate of 99 with heterogeneity in β (but far better than the fit measure of 350 for the standard model). More problematic is that the model has an implied annual discount factor of 0.002, such that the agents are *myopic*, placing essentially no weight on future consumption. In contrast to our results with $\beta = 0.5$, the exponential discount factor needed to generate hand-to-mouth behavior is inconsistent with prior work (Frederick et al. 2002).

We find estimates in line with prior work in terms of the composition of consumer types. We estimate that 25 percent of consumers in the population are present-biased in a way that generates hand-to-mouth behavior. This estimate is in line with recent estimates of the share hand-to-mouth used to interpret the consumption response to tax rebates (Kaplan and Violante 2014).[35] One interesting feature of our results is that there is dynamic selection of β types, similar to the dynamic selection in job search costs discussed above. Present-biased types have time-inconsistent preferences over job search–they search too little today because they incorrectly expect that they will search more tomorrow–and so the share of present-biased types among the unemployed rises over time. By month five—the last month of UI benefits—the present-biased types are 33 percent of the population.

4.3 Out-of-Sample Tests and Summary

One common critique of behavioral models is that even if they do a good job of capturing the facts at hand, introducing behavioral elements is a form of over-fitting (Gabaix and Laibson 2008). In many empirical contexts there is one benchmark rational model, but there are many possible behavioral models. If we select which behavioral model to estimate after seeing the data, this model is likely to make worse out-of-sample predictions than a more parsimonious rational model. In our context, we introduce present-biased agents in order to better fit the drop at exhaustion. A skeptical reader might expect that the rational model will do better when applied to a new economic environment (DellaVigna 2017).

Motivated by this critique, we compare the predictions of the models from the previous sections in a new economic environment: unemployment in Florida. As discussed in Section 3.3.1, Florida has some of the least generous UI benefits in the U.S. UI benefits replace only 40 percent of household income in Florida and last only four months. We take agents with the preferences estimated in Table 3 and compute optimal spending and search decisions for an environment with the household income process of UI recipients in Florida. The heterogeneous β model does a better job of fitting the spending data than the standard model, as shown visually in Figure 9. The improvement in fit is confirmed

[35] Our estimate is also qualitatively consistent with estimates based on time series data. Campbell and Mankiw (1989) estimate that aggregate data on *annual* consumption are consistent with about 50 percent of agents being hand-to-mouth consumers and 50 percent being permanent income consumers. Unlike Campbell and Mankiw, our second consumer type is a buffer stock agent with an exogenous liquidity constraint rather than a permanent income consumer, and this buffer stock type exhibits excess sensitivity behavior between a permanent income type and a hand-to-mouth type. Hence it is not surprising that we would estimate a larger share of buffer stock agents than Campbell and Mankiw estimate of permanent income consumers.

quantitatively by comparing the GOF measures for the consumption data, which are 169 and 406 respectively.[36]

We also explore using the distribution of the spending drop to test another prediction of the model in Section 4.2, but this test is inconclusive. Evaluated at our parameter estimates, that model features some agents who are very impatient and the rest are rational. This model has a testable prediction that there should be a bimodal distribution of consumption drops at exhaustion with a large mass close to zero and a second point mass equal to the value of UI benefits lost. We investigate this distribution empirically in an attempt to test this prediction. Unfortunately, because household income and spending are so volatile on a month-to-month basis, the JPMCI data around exhaustion are not informative for this question.[37] In future work, it would be interesting to use other types of heterogeneity—such as by unemployment duration (Figure 1) or by asset levels (Figure 3)—to further test these models. For example, it may be difficult for the heterogeneous β model to capture the empirical pattern of a drop in spending among high-asset types if those types have high β.

We propose a new test to distinguish between rational (and near-rational) models of consumption and behavioral alternatives with present-bias or myopia. We conclude that any model which can explain the behavior we document must include some agents who behave in a hand-to-mouth fashion, such that the model generates a large spending drop from a predictable income decline. Prior theories which sought to explain excess sensitivity using liquidity constraints, exponential impatience, near-rationality, and home production do not generate a large drop in spending from a predictable income decline at UI benefit exhaustion. A model with a significant share of agents that are present-biased matches the behavior of people eligible for six months of UI benefits. It also performs well in one out-of-sample test: predicting spending in Florida, which has much less generous UI benefits.

Finally, it is important to note that our new test distinguishes between rational models of consumption and a *class* of behavioral alternatives with present-bias or myopia. We implement a model with heterogeneity in present-bias, but as discussed above we could have generated the same results with some myopic agents who placed essentially no weight on utility from consumption one year in the future. Another intriguing possibility is that a large fraction of agents cut their spending at exhaustion by less than the drop in income, but more than the very small drop predicted by the rational model.[38]

[36] Online Appendix Figure 20 shows that the models make similar predictions for search effort in Florida.

[37] Although there is a sharp point mass in the *income* change distribution due to lost UI benefits, the distribution of the change in checking account *inflows* is much more diffuse, as shown in online Appendix Figure 21. Checking account inflows are more diffuse because they include non-UI deposits such as paper checks, transfers between accounts and other uncategorized transactions. Because the change in checking account inflows is diffuse, the fact that the spending drop at benefit exhaustion is also diffuse is not informative for distinguishing between models of consumer behavior.

[38] An example of a model where all agents cut spending at exhaustion by more than predicted by our baseline model is Gabaix (2016). In this model, agents have a cost of gathering all the information necessary to plan future consumption. They underestimate the size of the income drop associated with UI benefit exhaustion, are surprised by the size of the actual income drop at UI benefit exhaustion, and cut spending at exhaustion. We describe the logic underlying the model and our estimation procedure in online Appendix

5 Normative Implications for UI Policy

In this section, we use our empirical results to estimate the consumption-smoothing gain from extending the duration of UI benefits. Although most prior work on optimal UI benefits has focused on the *level* of UI benefits, there is substantial policy variation across countries, states, and time in the potential *duration* of benefits. A newer strand of the literature has developed theory to examine the welfare implications of extending the duration or changing the path of UI benefits (Kekre 2017; Schmieder and von Wachter 2017; Kolsrud et al. 2018). Implementing these theories requires estimates for both the fiscal costs and the consumption-smoothing gain of extending the duration of UI benefits. Although there is a rich literature on the former, we are not aware of any paper with estimates of the latter. As we show below, our estimates of the drop in consumption at benefit exhaustion help identify these consumption-smoothing gains.

We find that the welfare gains from improved consumption-smoothing due to extending the duration of UI benefits are four times as large as from raising the level of UI benefits. The economic intuition for our result is that long-term unemployment is a state of the world where consumption is much lower and marginal utility is much higher. Duration extensions target this high marginal utility state of the world more effectively than level increases. Finally, we show that extensions have positive welfare impacts even after taking into account the fiscal externality from UI benefits.

We calculate the consumption-smoothing gains from level increases and duration extensions using a sufficient statistic formula.[39] Specifically, we use the formula developed by Schmieder and von Wachter (2017) (henceforth SvW), who generalize the two-state Baily-Chetty formula to allow for an integrated treatment of the welfare consequences of changing benefit levels and durations in an environment where UI benefits last for a limited number of periods. We briefly describe their approach here and refer readers to their paper for further details. These formulas are also similar to those in Kolsrud et al. (2018), who develop a methodology to evaluate the welfare implications of small changes to the benefit level at any point in the unemployment spell.[40]

First, we consider a benefit level increase db. This is financed by a tax increase $d\tau$ on employed households. In states $j \in \{1 \dots 7\}$ the household is unemployed and receiving UI benefits. We calculate average consumption during each month of UI receipt $\bar{c}_j$ using our

F.2 and show the results in online Appendix Figure 22. The model is qualitatively successful in that it produces a larger drop in spending at exhaustion than our baseline model of fully rational forward-looking households. However, our implementation of this model is unable to quantitatively match the size of the drop at exhaustion.

[39] In Online Appendix G and online Appendix Table 20 we compare the gains in the context of the structural model from Section 4.2 and find very similar results.

[40] They empirically analyze the optimal path of UI benefits in Sweden, where exhaustion is not relevant because there is no time limit on benefit receipt (though the benefit level becomes less generous after 20 weeks of unemployment). The Schmieder and von Wachter (2017) formula can be thought of as a special case of Kolsrud et al. (2018) applied to evaluate a larger change in UI benefits—increasing the benefit level in the first month of exhaustion from zero to the level prior to exhaustion.

estimate of the average consumption path in Figure 2. We calculate π_j, the fraction of time that households are in each state, using our estimates of the job-finding hazards reported in Figure 6. Equation (1) in SvW implies that, in this environment, the welfare gain of a benefit level increase can be approximated as

$$\frac{dW}{db} \approx \left[\sum_{j=1}^{7} u'(\bar{c}_{ui,j})\pi_j\right] db - u'(\bar{c}_{emp})\pi_{emp} d\tau. \tag{9}$$

The intuition for this formula is that the first term values the transfer db to each month of UI receipt using the marginal utility for the average household in that month (weighted by the share of households who reach that month of UI receipt), and the second term values the cost of financing this transfer by raising taxes $d\tau$ in the employed state using the marginal utility of the average employed household. When we implement this formula in Table 4, we report this welfare change normalized by a Lucas-type money metric: $\frac{d\tilde{W}}{db} = \frac{dW}{db}/u'(\bar{c}_{emp})$.

Second, we consider an extension in the potential duration of benefits by dP periods. This raises income in each of these periods by the benefit level b, and is similarly financed by a tax increase $d\tau$ on employed households. Since we will focus on one-month extensions, we calculate average consumption of exhaustees (who receive this benefit extension) $\bar{c}_{exhaust}$ using our estimate for average consumption in the first month of exhaustion in Figure 2. Equation (2) in SvW shows that the welfare gain from this potential benefit duration extension can be approximated as

$$\frac{dW}{dP} \approx u'(\bar{c}_{exhaust})\pi_{exhaust} b \cdot dP - u'(\bar{c}_{emp})\pi_{emp} d\tau. \tag{10}$$

The intuition for equation (10) is similar to equation (9), except that the transfer of benefits ($b \cdot dP$) is going to the first month after exhaustion and so is valued according to the average marginal utility of consumption for that group.

To implement the formulas above quantitatively, we consider benefit increases and extensions of equivalent cost. In the absence of job search distortions (which we return to below), the tax increase needed to finance a one-month benefit extension dP is also sufficient to fund a 1.8 percentage point increase in household income db during UI receipt. As in Section 4, we assume constant relative risk aversion with a risk aversion parameter of two. Implementing equation (9) for such a benefit increase, Table 4 shows in column 1 that private welfare is 0.021 percent higher when using the JPMCI nondurable consumption estimates and 0.019 percent higher when using the Gruber (1997) PSID food consumption estimates. This is unsurprising because Section 3.1 shows that our estimates of the spending drop during UI receipt are similar to Gruber (1997). However, Gruber is unable to measure the consumption-smoothing value of extensions because the PSID only contains annual consumption data.

We fill this gap using the SvW generalization of the Baily-Chetty formula and find that

duration extensions have consumption-smoothing gains that are four times larger than level increases. Implementing equation (10) indicates that welfare is 0.082 percent higher under a one-month benefit duration extension, which is four times larger than the gains from level increases discussed in the previous paragraph. This result is not driven by our choice of the risk aversion parameter. The ratio of the gains from a duration extension to the gains from a level increase varies from 3.6 to 4.8 as the risk aversion parameter rises from around one to four (online Appendix Table 19).

A full evaluation of the welfare gains from increasing benefit levels and extending benefit durations requires incorporating the fiscal externality due to job search distortions from more generous UI. Our data are not well-suited for analyzing job search distortions, so we use estimates from the SvW literature review.[41] They propose a metric which is the ratio of behavioral cost (BC)—the total cost to the government of increasing UI generosity, including the extra spending induced because UI recipients will respond by taking longer to find a job—to the mechanical cost (MC) of increasing generosity absent any change in behavior. They call this statistic the "BCMC" ratio. In their literature review, they report a median BCMC ratio of 1.32 across 12 studies of increases and a median BCMC ratio of 1.52 across 11 studies of extensions. This implies that spending $1 on extensions is 15 percent more expensive than spending $1 on level increases. The specific concern here is that because the BCMC ratio is larger for extensions than for level increases and each policy is funded with a tax on the employed state, then the welfare cost from the increased taxes needed to fund the extensions may outweigh the extra consumption-smoothing gains documented above. We continue to find a welfare increase from UI benefit extensions, even after incorporating the fiscal externality from job search distortions. To incorporate the fiscal externality, we re-evaluate equations (9) and (10), adjusting taxes for the median BCMC ratios from SvW. Table 4 shows that increasing UI benefits *lowers* welfare by 0.023 percent, while extending benefits *raises* welfare by 0.016 percent.

A promising area for future research is to examine the consequences of UI benefit increases and extensions in a richer policy space and with richer models of job search. On the policy side, the presence of households who sharply cut spending in response to a sharp cliff in benefits may mean that the time-limited benefit schedule observed in the U.S. is itself sub-optimal. It would be interesting to understand the welfare consequences of dramatically

[41] The problem is that our data only have *cross-sectional* variation in UI benefit levels, while SvW review *quasi-experimental* estimates of the distortions from changes in UI generosity. Although most of the estimates in the SvW literature review come from European countries, we obtain similar estimates using our cross-sectional variation in the fiscal cost of alternative UI policies in Florida and New Jersey, which suggests that the median of the SvW estimates is a plausible benchmark for U.S. data. In New Jersey, the mechanical cost of higher monthly benefits is 7.8 percent higher UI spending, compared to other states where UI benefits last six months. Using the central estimate of BCMC ratios from SvW, we project that total expenditures per UI recipient will be 10.4 percent higher. In fact, they are 11.0 percent higher in the JPMCI data. In Florida, the mechanical savings from lower benefit levels and shorter durations is 38.5 percent lower UI spending. Using the central estimates of BCMC ratios from SvW, we project that total expenditures per UI recipient would be 49.1 percent lower. In fact, they are 45.7 percent lower in the JPMCI data.

reforming the benefit schedule by allowing for permanently stable, gradually declining, or gradually increasing benefits. In addition, it would be useful to conduct a welfare analysis which embeds a model with reference-dependent job search (e.g. DellaVigna et al. 2017) and/or human capital depreciation (Acemoglu 1995). It is possible that introducing these forces would reduce the apparent attractiveness of benefit extensions (Lindner and Reizer 2016).

6 Conclusion

This paper documents the paths of spending and job search during unemployment using high-frequency bank account data. Spending is highly responsive to the level of UI benefits and drops sharply at benefit exhaustion. Job search is also responsive to the level of UI benefits, rising at benefit exhaustion and falling afterward. The drop in spending in response to a predictable income decline is an identified moment allowing us to distinguish between two kinds of theories economists have developed to explain the excess sensitivity of spending to income. We show that this drop is inconsistent with the rational buffer stock model. Instead, a model where some liquidity-constrained agents exhibit present-bias or myopia is necessary to match the path of spending during unemployment. Low spending after exhaustion implies that the consumption-smoothing gains from extending UI durations are four times larger than the gains from increasing the level of UI benefits.

We highlight three interesting directions for future research. First, measuring the high-frequency consumption responses to other predictable income declines will continue to improve our understanding of consumption models. For example, in subsequent work Jorring (2018) studies the consumption response to the upward resetting of payments on home equity lines of credit and Baugh et al. (2018) study the consumption response to income tax payments.

Second, it would be interesting to further explore our finding that households do not seem to borrow much during unemployment. For example, households only borrow an average of about $20 per month on Chase credit cards during unemployment, despite having large unused credit lines. Because unemployment is a mostly temporary shock to income, the rational buffer stock model predicts a large increase in credit card utilization during unemployment. The absence of credit card borrowing we observe among unemployed households is particularly striking against the backdrop of widespread credit card borrowing by U.S. households overall (Laibson et al. 2007).

Finally, our estimated model implies that the distribution of the spending drop is bimodal, with many households cutting spending by an amount equal to the income loss at benefit exhaustion, and a small spending drop for the remainder of households. Another theory consistent with the data is that a large share of households are slightly myopic, cutting spending by some amount less than the lost income. As we discuss in Section 4.3, it is difficult to test between these hypotheses using the JPMCI data because of volatility

in checking account inflows. With data on total household liquid assets to better identify which households are hand-to-mouth or an alternative statistical methodology for studying income and spending, it might be possible to distinguish between these hypotheses.

References

Acemoglu, D. (1995). Public Policy in a Model of Long-term Unemployment. *Economica*, 62(246):161–178.

Aguiar, M. and Hurst, E. (2005). Consumption versus Expenditure. *Journal of Political Economy*, 113(5):919–948.

Altonji, J. G. and Siow, A. (1987). Testing the Response of Consumption to Income Changes with (Noisy) Panel Data. *The Quarterly Journal of Economics*, 102(2):293–328.

Baily, M. N. (1978). Some aspects of optimal unemployment insurance. *Journal of Public Economics*, 10(3):379–402.

Baker, S. R. and Yannelis, C. (2017). Income Changes and Consumption: Evidence from the 2013 Federal Government Shutdown. *Review of Economic Dynamics*, 23:99–124.

Banks, J., Blundell, R., and Tanner, S. (1998). Is There a Retirement-Savings Puzzle? *American Economic Review*, 88(4):769–788.

Baugh, B., Ben-David, I., Park, H., and Parker, J. A. (2018). Asymmetric Consumption Response of Households to Positive and Negative Anticipated Cash Flows. Working Paper 25086, National Bureau of Economic Research.

Berger, D., Guerrieri, V., Lorenzoni, G., and Vavra, J. (2018). House Prices and Consumer Spending. *The Review of Economic Studies*.

Bernheim, B. D., Skinner, J., and Weinberg, S. (2001). What Accounts for the Variation in Retirement Wealth among U.S. Households? *American Economic Review*, 91(4):832–857.

Bethune, Z. (2017). Consumer Credit, Unemployment, and Aggregate Labor Market Dynamics. *Working Paper*.

Bowman, D., Minehart, D., and Rabin, M. (1999). Loss aversion in a consumption-savings model. *Journal of Economic Behavior & Organization*, 38(2):155–178.

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

Campbell, J. Y. and Mankiw, N. G. (1989). Consumption, Income and Interest Rates: Reinterpreting the Time Series Evidence. In *NBER Chapters*, pages 185–246. 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., Palmer, N., White, M. N., Kazil, J., and Low, D. (2018). econ-ark/HARK: 0.8.0 (Version pre).

Carroll, C., Slacalek, J., Tokuoka, K., and White, M. N. (2017). The Distribution of Wealth and the Marginal Propensity to Consume. *Quantitative Economics*, 8(3):977–1020.

Carroll, C. D. (1997). Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis. *The Quarterly Journal of Economics*, 112(1):1–55.

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. (2016). The Cyclicality of the Opportunity Cost of Employment. *Journal of Political Economy*, 124(6):1563–1618.

Deaton, A. (1991). Saving and Liquidity Constraints. *Econometrica*, 59(5):1221–48.

DellaVigna, S. (2017). Structural Behavioral Economics. *Forthcoming in Handbook of Behavioral Economics, eds. David Laibson, Douglas Bernheim, and Stefano DellaVigna*.

DellaVigna, S., Lindner, A., Reizer, B., and Schmieder, J. F. (2017). Reference-dependent Job Search: Evidence from Hungary. *The Quarterly Journal of Economics*, 132(4):1969–2018.

DiMaggio, M. and Kermani, A. (2016). The Importance of Unemployment Insurance as an Automatic Stabilizer. Working Paper 22625, National Bureau of Economic Research.

East, C. N. and Kuka, E. (2015). Reexamining the consumption smoothing benefits of Unemployment Insurance. *Journal of Public Economics*, 132(C):32–50.

Frederick, S., Loewenstein, G., and O'Donoghue, T. (2002). Time Discounting and Time Preference: A Critical Review. *Journal of Economic Literature*, 40(2):351–401.

Fuchs-Schuendeln, N. and Hassan, T. A. (2016). Natural Experiments in Macroeconomics. In *Handbook of Macroeconomics*, volume 2a, pages 923–1012. Elsevier.

Gabaix, X. (2016). Behavioral Macroeconomics Via Sparse Dynamic Programming. Working Paper 21848, National Bureau of Economic Research.

Gabaix, X. and Laibson, D. (2008). The Seven Properties of Good Models. In *The Foundations of Positive and Normative Economics*. Oxford University Press.

Garcia, R., Lusardi, A., and Ng, S. (1997). Excess Sensitivity and Asymmetries in Consumption: An Empirical Investigation. *Journal of Money, Credit and Banking*, 29(2):154–176.

Gelman, M., Kariv, S., Shapiro, M. D., Silverman, D., and Tadelis, S. (2015). How Individuals Smooth Spending: Evidence from the 2013 Government Shutdown Using Account Data. Working Paper 21025, National Bureau of Economic Research.

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.

Hendren, N. (2017). Knowledge of Future Job Loss and Implications for Unemployment Insurance. *American Economic Review*, 107(7):1778–1823.

Herkenhoff, K., Phillips, G., and Cohen-Cole, E. (2016). How Credit Constraints Impact Job Finding Rates, Sorting & Aggregate Output. Working Paper 22274, National Bureau of Economic Research.

Hsieh, C.-T. (2003). Do Consumers React to Anticipated Income Changes? Evidence from the Alaska Permanent Fund. *American Economic Review*, 93(1):397–405.

Jappelli, T. and Pistaferri, L. (2010). The Consumption Response to Income Changes. *Annual Review of Economics*, 2(1):479–506.

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.

Jorring, A. (2018). The Costs of Financial Mistakes: Evidence from U.S. Consumers. *Working Paper*.

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). The impact of the potential duration of unemployment benefits on the duration of unemployment. *Journal of Public Economics*, 41(1):45–72.

Kekre, R. (2017). Unemployment Insurance in Macroeconomic Stabilization. *Working Paper*.

Kolsrud, J., Landais, C., Nilsson, P., and Spinnewijn, J. (2018). The Optimal Timing of Unemployment Benefits: Theory and Evidence from Sweden. *American Economic Review*, 108(4-5):985–1033.

Kroft, K., Lange, F., Notowidigdo, M. J., and Katz, L. F. (2016). Long-Term Unemployment and the Great Recession: The Role of Composition, Duration Dependence, and Nonparticipation. *Journal of Labor Economics*, 34(S1):7–54.

Kroft, K. and Notowidigdo, M. J. (2016). Should Unemployment Insurance Vary with the Unemployment Rate? Theory and Evidence. *Review of Economic Studies*, 83(3):1092–1124.

Krusell, P. and Smith, A. (1998). Income and Wealth Heterogeneity in the Macroeconomy. *Journal of Political Economy*, 106(5):867–896.

Kueng, L. (2015). Explaining Consumption Excess Sensitivity with Near-Rationality: Evidence from Large Predetermined Payments. Working Paper 21772, National Bureau of Economic Research.

Laibson, D. (1997). Golden Eggs and Hyperbolic Discounting. *The Quarterly Journal of Economics*, 112(2):443–77.

Laibson, D., Repetto, A., and Tobacman, J. (2007). Estimating Discount Functions with Consumption Choices over the Lifecycle. Working Paper 13314, National Bureau of Economic Research.

Lindner, A. and Reizer, B. (2016). Frontloading the Unemployment Benefit: An Empirical Assessment. IEHAS Discussion Paper 1627, Institute of Economics, Centre for Economic and Regional Studies, Hungarian Academy of Sciences.

Lusardi, A. (1996). Permanent Income, Current Income, and Consumption: Evidence from Two Panel Data Sets. *Journal of Business & Economic Statistics*, 14(1):81–90.

McKee, G. and Verner, E. (2015). The Consumption Response to Extended Unemployment Benefits in the Great Recession. SSRN Scholarly Paper ID 2630790, Social Science Research Network, Rochester, NY.

McKenna, C. and McHugh, R. (2016). Share of Unemployed Receiving Jobless Aid Remained at Record Low in 2015. *National Employment Law Project Blog*.

Meyer, B. D. (1990). Unemployment Insurance and Unemployment Spells. *Econometrica*, 58(4):757–782.

Mortensen, D. T. (1977). Unemployment Insurance and Job Search Decisions. *ILR Review*, 30(4):505–517.

Nakamura, E. and Steinsson, J. (2017). Identification in Macroeconomics. Working Paper 23968, National Bureau of Economic Research.

Pagel, M. and Vardardottir, A. (2016). The Liquid Hand-to-Mouth: Evidence from a Personal Finance Management Software. In *2016 Annual Meeting of the Society for Economic Dynamics*.

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):2530–2553.

Paserman, M. D. (2008). Job Search and Hyperbolic Discounting: Structural Estimation and Policy Evaluation*. *The Economic Journal*, 118(531):1418–1452.

Passero, W., Garner, T. I., and McCully, C. (2014). Understanding the Relationship: CE Survey and PCE. In *Improving the Measurement of Consumer Expenditures*, pages 181–203. University of Chicago Press.

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. (2017). Scraping By: Income and Program Participation After the Loss of Extended Unemployment Benefits. *Journal of Policy Analysis and Management*, 36(4):880–908.

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. and von Wachter, T. (2017). A Context-Robust Measure of the Disincentive Cost of Unemployment Insurance. *American Economic Review*, 107(5):343–348.

Schmieder, J. F., von Wachter, T., and Bender, S. (2016). The Effect of Unemployment Benefits and Nonemployment Durations on Wages. *American Economic Review*, 106(3):739–777.

Shapiro, M. D. and Slemrod, J. (2009). Did the 2008 Tax Rebates Stimulate Spending? *American Economic Review*, 99(2):374–79.

Shea, J. (1995). Union Contracts and the Life-Cycle/Permanent-Income Hypothesis. *American Economic Review*, 85(1):186–200.

Souleles, N. S. (1999). The Response of Household Consumption to Income Tax Refunds. *American Economic Review*, 89(4):947–958.

Stephens, M. and Toohey, D. (2018). Changes in Nutrient Intake at Retirement. Technical Report 24621, National Bureau of Economic Research, Inc.

TD Bank (2014). TD Bank Survey Finds Many Couples Maintain Separate Bank Accounts. Technical Report, TD Bank.

Welander, T. (2014). Trends in Consumer Payments and Retail Banking: Report 1 of 4. Technical report, GC Insights Marketing Research Services.

Figure 1: Event Study by UI Duration

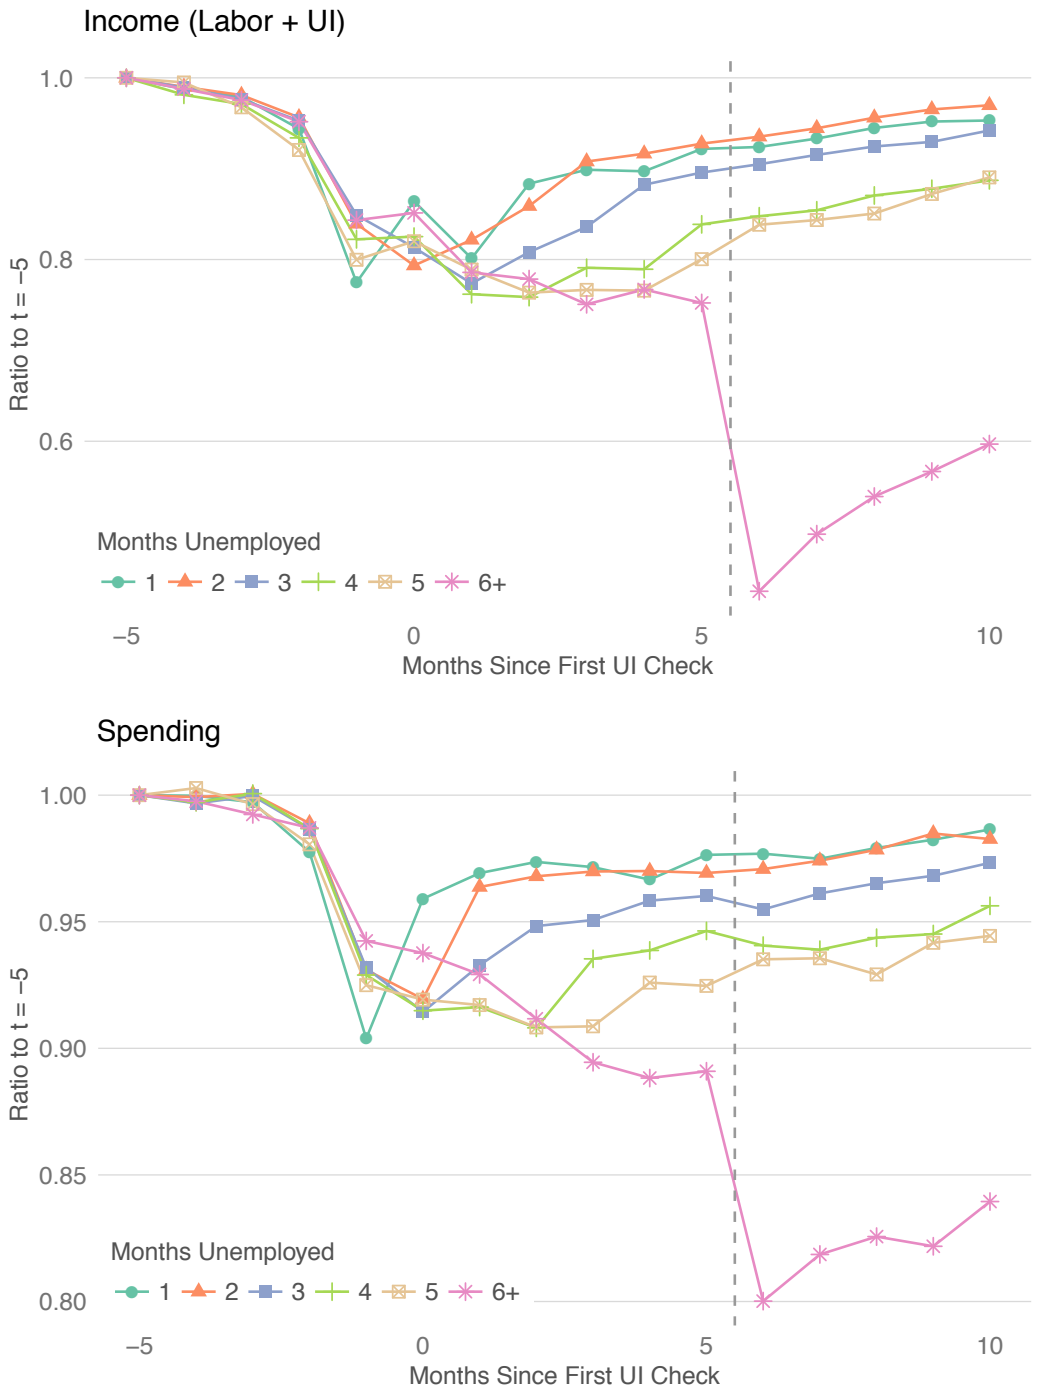

Notes: This figure plots household labor income plus UI benefits (top panel) and spending on nondurables (bottom panel) as a function of completed UI duration. The vertical line marks UI benefit exhaustion. Income is positive after UI benefit exhaustion because of labor income of other household members. The 6+ group includes all households who receive six months of UI benefits and, unlike the other lines, is a composite of households with different non-employment durations. In online Appendix Figure 4 we further stratify the spending path for the 6+ group to show results separately for exhaustees with non-employment durations of 6, 7, 8, 9, and 10+ months. Sample is households that receive UI benefits and meet the sampling criteria described in Section 2.1.

Figure 2: Income and Spending If Stay Unemployed

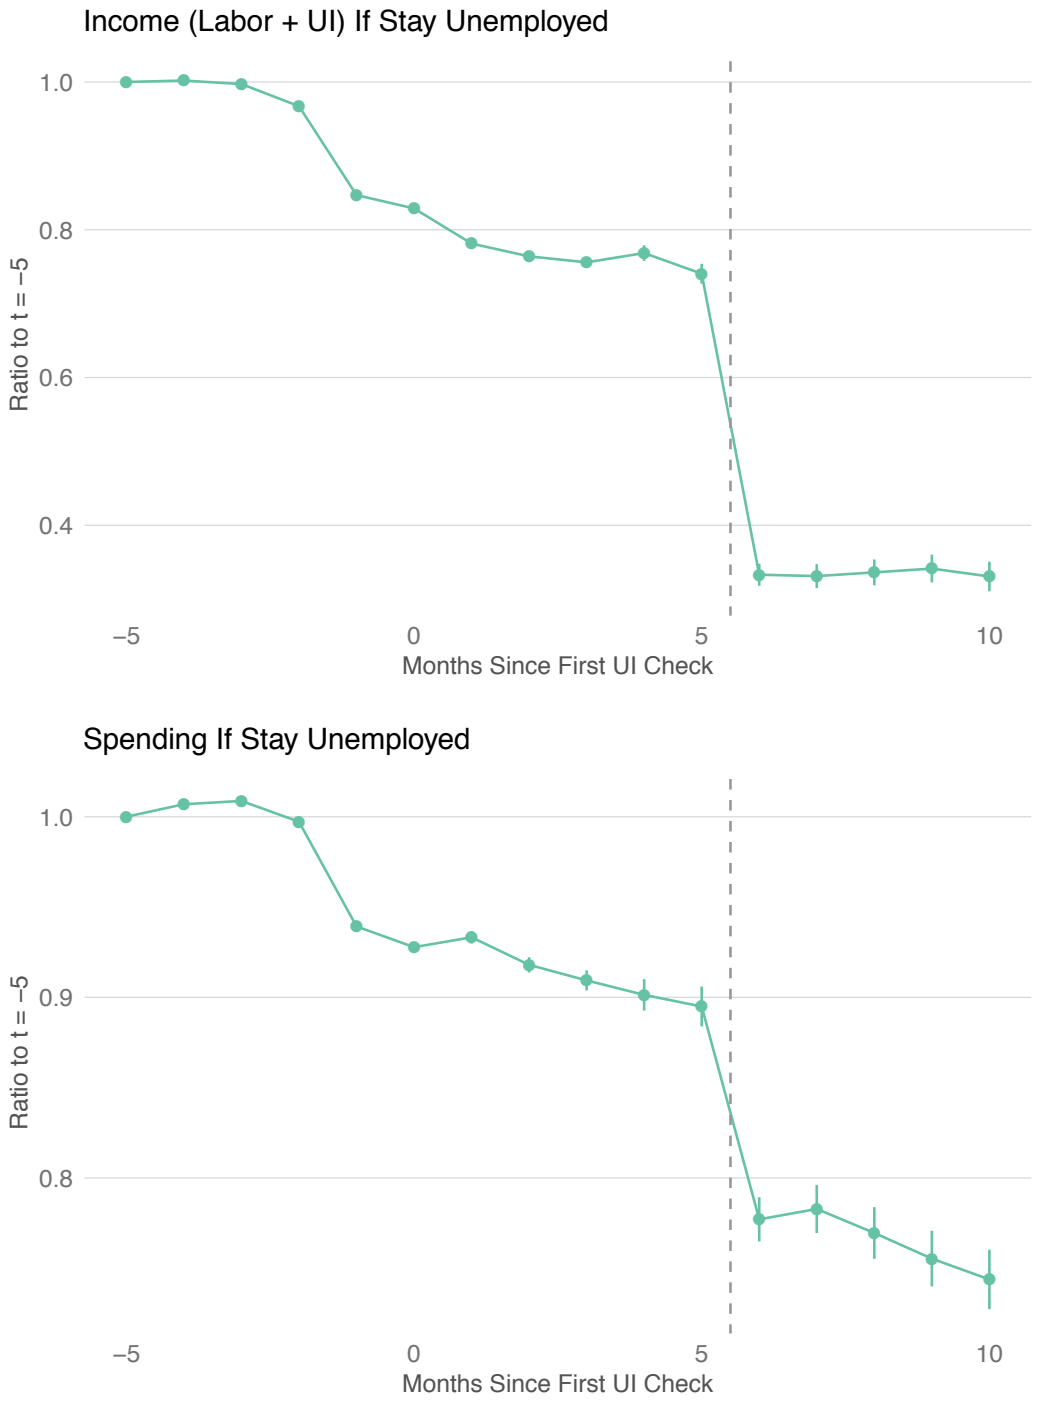

Notes: This figure plots income and spending for the sample that stays unemployed. In months $t = \{-5, -4, -3, -2, -1, 0\}$, this includes everyone who receives UI at date 0 and meets the sampling criteria described in Section 2.1. In month $t = 1$, this includes only households who continue to receive UI and excludes households who receive their last UI check in month 0. In month $t = 2$, this excludes households who receive their last UI check in month 0 or month 1, and so on. Employment status after UI exhaustion is measured using paycheck deposits. The vertical line marks UI benefit exhaustion. Income is positive after UI benefit exhaustion because of labor income of other household members. Vertical lines denote 95 percent confidence intervals for change from the prior month. See Section 3.1.1 for details.

Figure 3: Heterogeneity in Income and Spending If Stay Unemployed

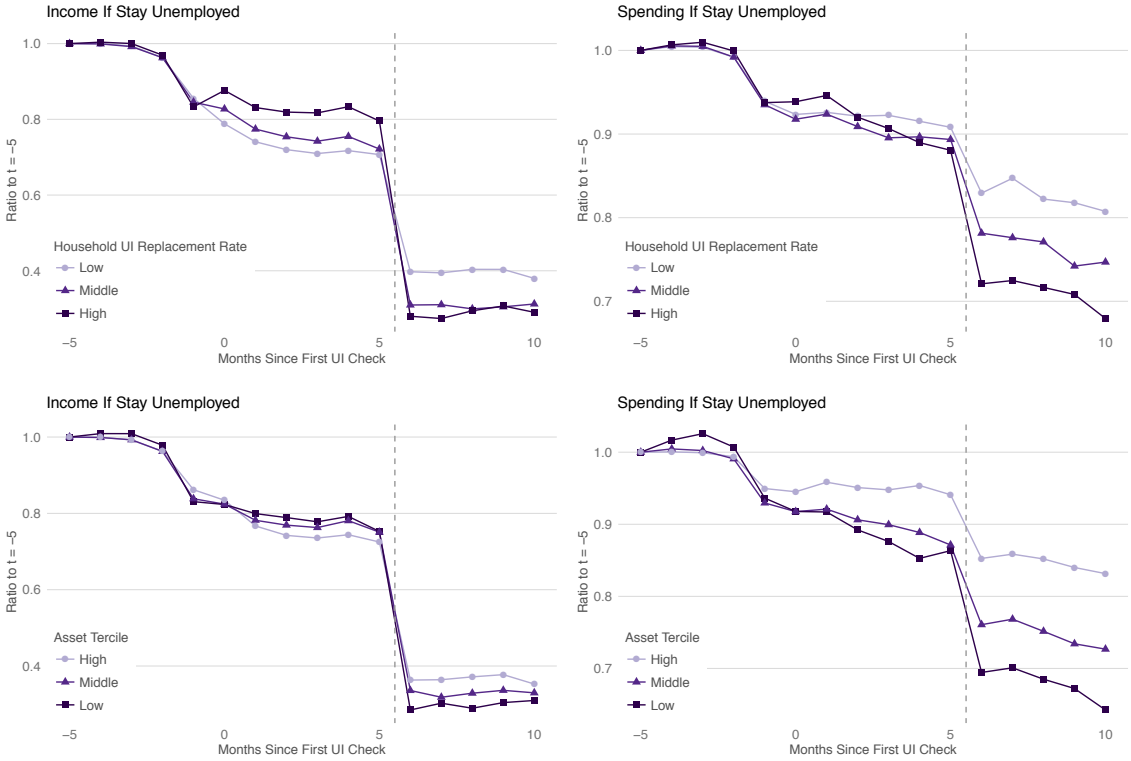

Notes: This figure shows heterogeneity in income and spending by the ratio of UI benefits to estimated household annual income and the ratio of estimated total liquid assets (a measure described in Section 2.2) to consumption prior to the onset of unemployment. The sample is households that receive UI and stay unemployed, as described in the note to Figure 2.

Figure 4: Spending and State-Level Variation in UI Potential Duration and Timing

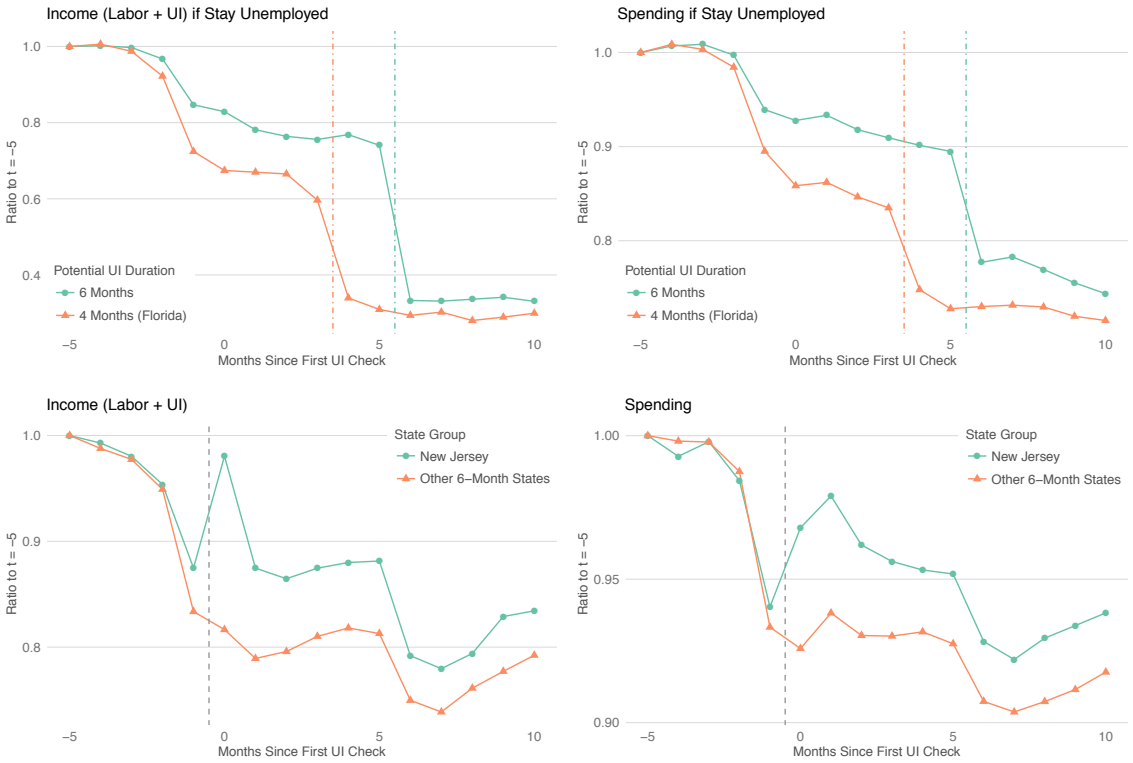

Notes: This figure analyzes income and spending in Florida (top panels) and New Jersey (bottom panels). See Section 3.3.1 for details.

Top panels: Although most states offer up to six months of UI benefits, Florida offered a low UI benefit for up to four months from January 2014 through June 2015. This figure compares income and spending for people who stay unemployed in Florida to six-month states. The vertical lines mark exhaustion. See Section 3.3.1 for details. The sample is households that receive UI and stay unemployed, as described in the note to Figure 2.

Bottom panels: Compared to other states offering six months of benefits, New Jersey has a high benefit level and begins benefit payments quickly after unemployment onset. On average, income spikes at the beginning of UI receipt in New Jersey because there is a week where the household receives both a paycheck and a UI check. This figure plots income and spending for all UI recipients that meet the sampling criteria described in Section 2.1 (not just those who stay unemployed) in order to emphasize the spike in income and spending at onset. Online Appendix Figure 11 plots the same series for the stay-unemployed sample. The vertical line marks the beginning of UI payments.

Figure 5: Marginal Propensity to Consume Out of UI Benefits

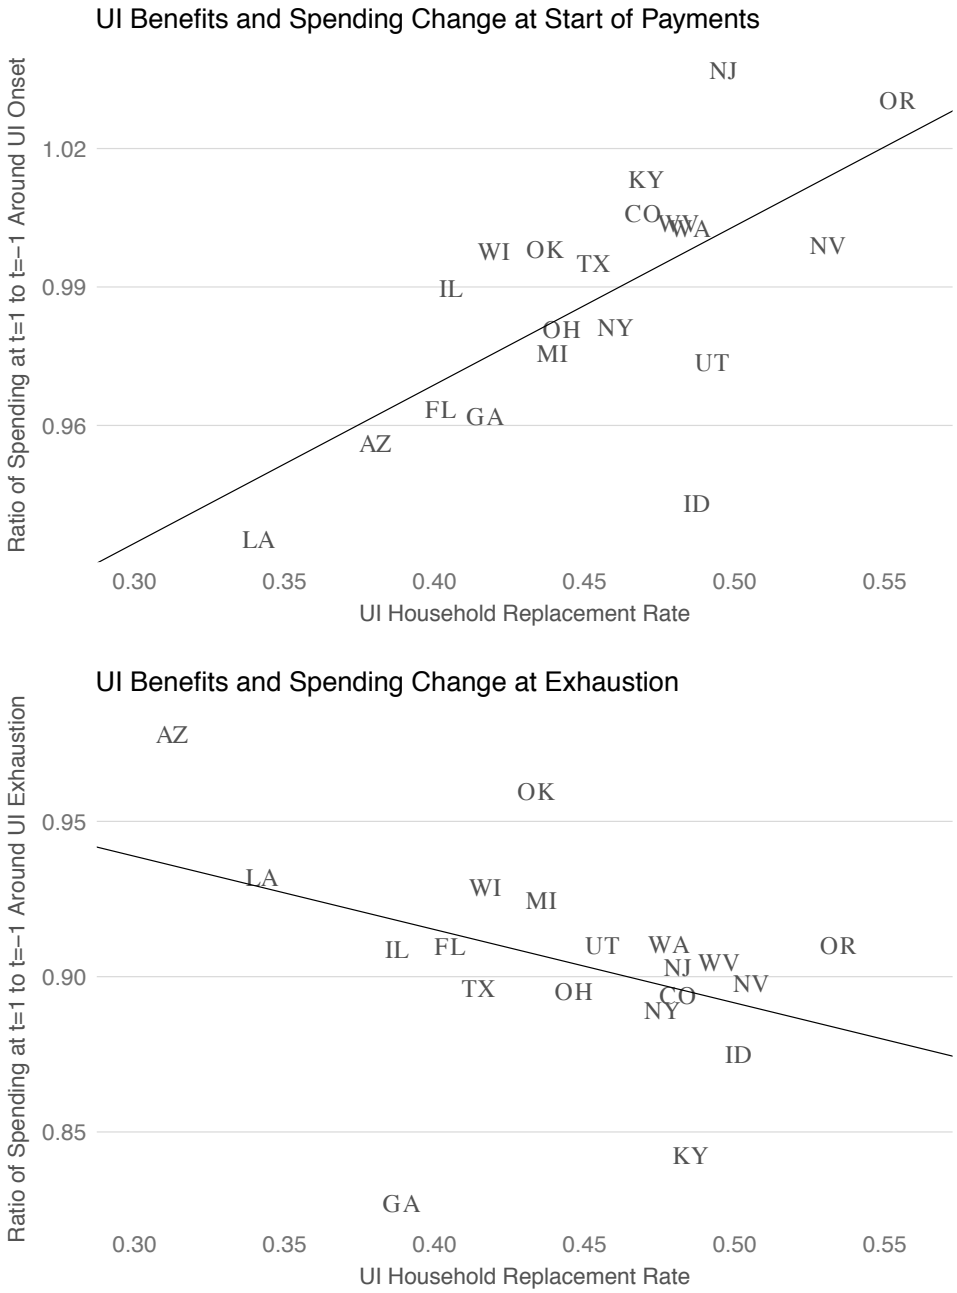

Notes: This figure compares the state-level UI household replacement rate and spending change around UI onset (top panel) and UI exhaustion (bottom panel). The estimated slope in the top panel implies a marginal propensity to consume of 27 cents out of UI benefits. We analyze households that receive UI benefits and meet the sampling criteria described in Section 2.1. The top panel includes households that received a full calendar month of benefits at $t = 1$ and the bottom panel includes households that exhausted UI benefits. The x-axis values differ between the two panels because the bottom panel calculates the replacement rate for exhaustees only. See Section 3.3.2 for details. The two states that are outliers in the bottom panel—Arizona and Georgia—account for 0.6 percent and 1.3 percent of the sample of exhaustees, respectively.

Figure 6: Job-Finding Hazard

Note: This figure plots job-finding hazards around UI benefit exhaustion. See Section 3.4 for details on how we construct a job-finding hazard from paycheck-inferred job starts by UI recipients. Vertical lines denote 95 percent confidence intervals computed by bootstrap. Among households eligible for 26 weeks of benefits, 37 percent will receive benefit checks in six calendar months and the other 63 percent will receive benefits for seven calendar months. This figure pools these groups and reports the job-finding hazard by months until the last potential UI check is received. Online Appendix Figure 26 reports job-finding hazards separately for these two groups.

Figure 7: Standard Model

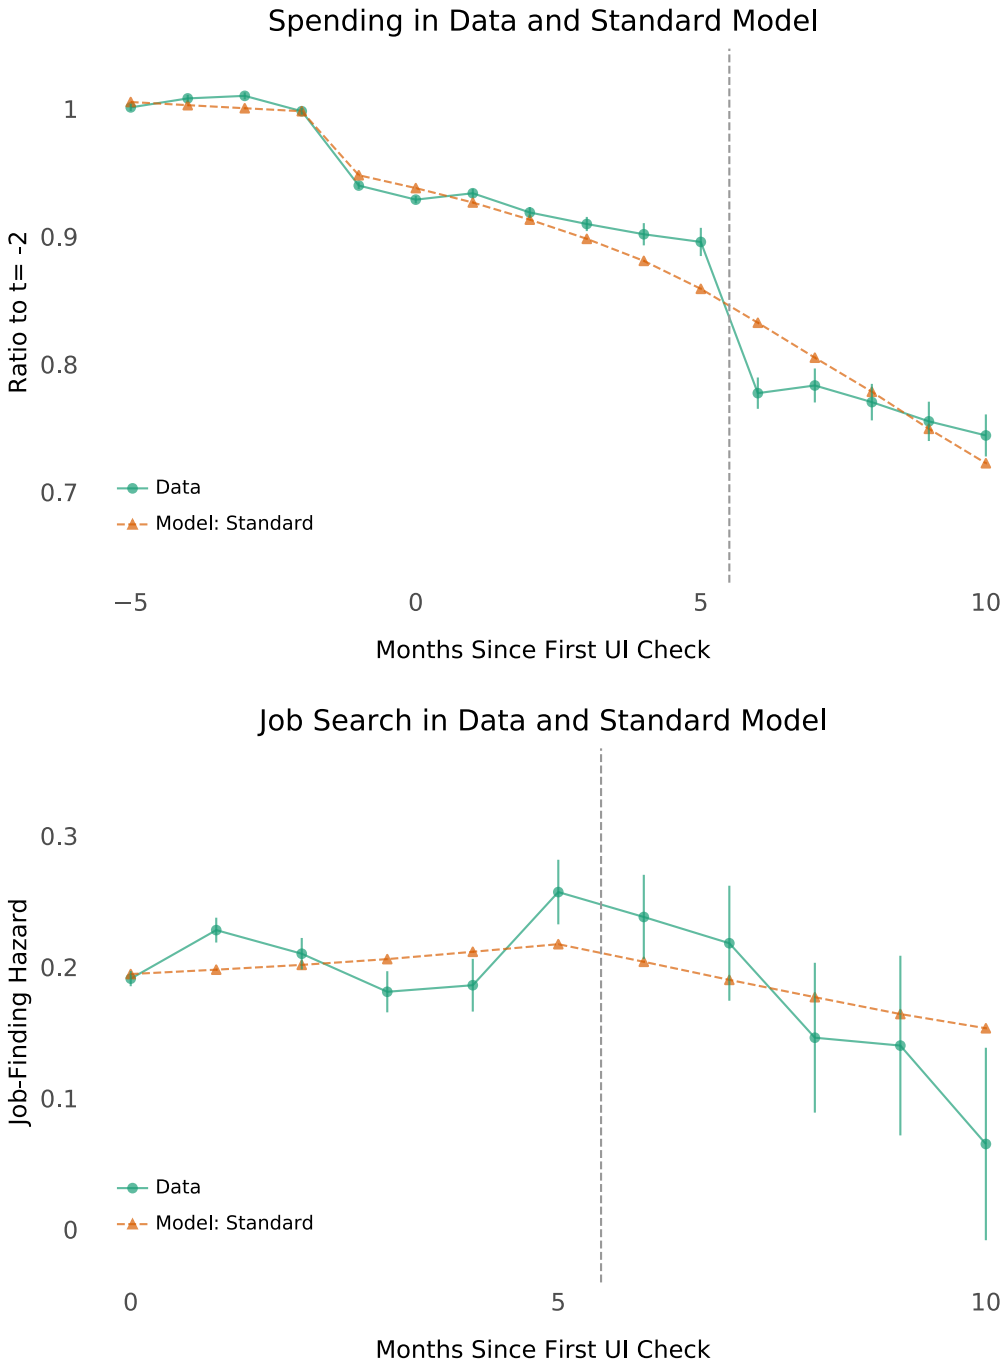

Notes: This figure plots the predictions of the "standard model" which features rational forward-looking agents with a liquidity constraint described in Section 4.1 alongside data on spending and job search during an unemployment spell from Figures 2 and 6, respectively. The figure shows average predicted behavior aggregating over two types of agents with different job search costs. Predictions and data are for an environment where UI benefits last six months and the UI recipient stays unemployed for ten months. See online Appendix Figure 14 for spending and search behavior of each type in the model.

Figure 8: Heterogeneity in Beta

Notes: This figure plots the predictions of the model with heterogeneous agents described in Section 4.2 alongside data on spending and job search during an unemployment spell from Figures 2 and 6 respectively. This model features heterogeneity in the patience parameter β–which provides a micro-foundation for "spender-saver" dynamics in aggregate spending–and job search cost parameter k. Predictions and data are for an environment where UI benefits last six months and the UI recipient stays unemployed for ten months. See online Appendix Figure 19 for spending and search behavior of each type in the model.

Figure 9: Out-of-Sample Test of Models

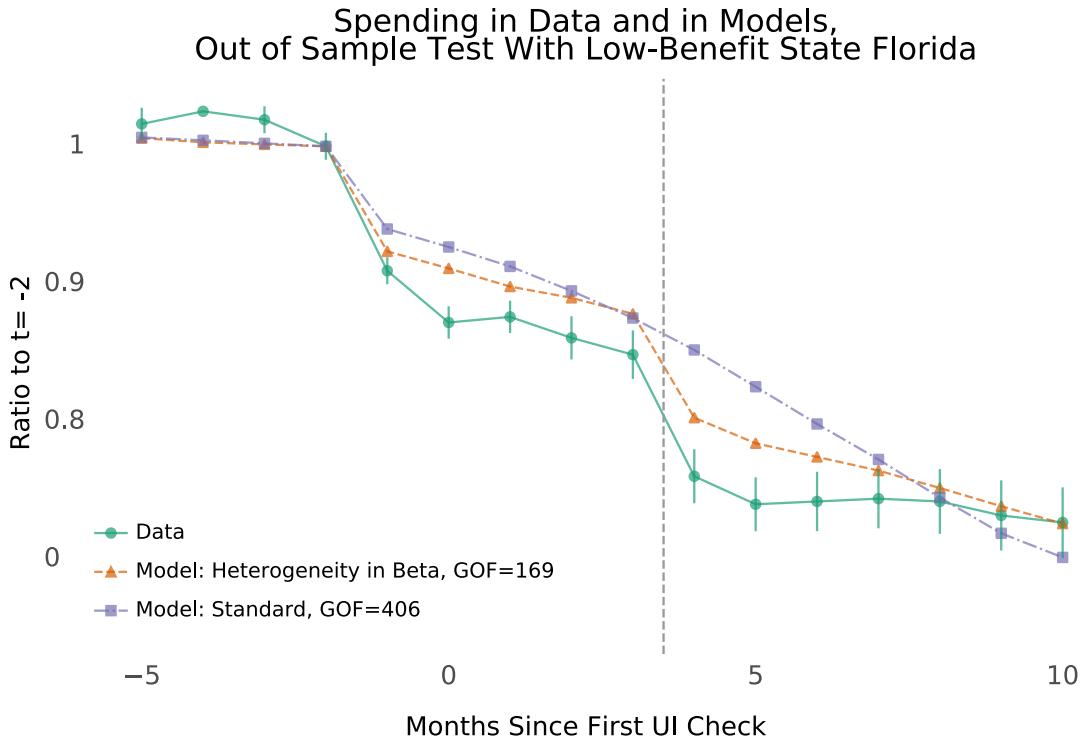

Notes: We use the standard model estimated in Figure 7 and the heterogeneous β model estimated in Figure 8 on data from states where UI benefits last six months to predict spending in Florida, a state where UI benefits last four months and are much lower than other states. This figure plots those predictions compared to the data series from Figure 4. The legend includes the consumption goodness-of-fit (GOF) statistic for each model series. See Section 4.3 for details.

Table 1: Representativeness: JPMCI Data Compared to External Benchmarks

Category	JPMCI (1)	Benchmark (2)	Ratio (1) / (2) (3)	Benchmark Source (4)
Mean Spending[a]				
Nondurables	$2,317	$1,671	139%	CE Survey
Nondurables	$2,317	$3,490	66%	PCE
Durables	$395	$1,280	31%	CE Survey
Durables	$395	$1,643	24%	PCE
Mean Income (Pre-tax Direct Deposit + Paper Checks)[b]				
Labor Earnings	$5,002	$5,750	87%	SIPP
Total Income	$6,334	$6,290	101%	SIPP
Mean Age	41.1	44.3	93%	SIPP
Median Checking Account Balance	$1,250	$1,500	83%	SCF
Number of U.S. States	20	50	--	--

Notes: This table compares the representativenessness of UI recipients in the JPMCI data that meet the sampling criteria described in Section 2.1 to external benchmarks from the Consumer Expenditure (CE) Survey for 2015, Bureau of Economic Analysis' Table 2.4.5U for 2015 (PCE), the Survey of Income and Program Participation for 2004 (SIPP), and the Survey of Consumer Finances for 2013 (SCF). All income and spending varibles are monthly.

a. Definition of nondurable and durable spending from Lusardi (1996). See Section 2.2 for details.

b. Labor income is adjusted for the fact that some earnings are paid by paper check rather than direct deposit (see notes to online Appendix Table 5 for details).

Table 2: Spending Change at UI Exhaustion

		Pre Onset	Pre Exhaustion	Post Exhaustion	Change (col 3 - col 2)	Change (col 4 / col 2)
Spending		$	$	$	$	%
Type[a]	Category[b]	(1)	(2)	(3)	(4)	(5)
Durable	Home Improvement	48.7	46.5	37.2	-9.4	-20.2%
Other ND	Discount Stores	57.7	58.1	47.1	-11.0	-18.9%
Other ND	Department Stores	19.4	16.5	13.6	-2.9	-17.7%
Durable	Miscellaneous Durables	27.1	26.3	21.8	-4.5	-17.1%
Other ND	Other Retail	148.0	137.0	114.4	-22.6	-16.5%
Strict ND	Food Away From Home	193.4	164.3	138.2	-26.1	-15.9%
Strict ND	Groceries	302.3	293.7	247.4	-46.3	-15.8%
Other ND	Drug Stores	39.5	35.4	30.0	-5.4	-15.3%
Durable	Retail Durables	48.3	43.3	36.7	-6.6	-15.3%
Nondurable	Cash	703.7	584.1	495.9	-88.2	-15.1%
Other ND	Medical Copay	35.4	29.3	25.3	-4.0	-13.6%
Durable	Entertainment	29.4	27.0	23.4	-3.6	-13.4%
Durable	Auto Repair	40.4	36.3	31.6	-4.7	-12.9%
Other ND	Online	42.6	38.8	34.1	-4.7	-12.1%
Strict ND	Transportation	155.6	127.6	114.0	-13.6	-10.6%
Durable	Hotels & Rental Cars	27.0	21.4	19.2	-2.2	-10.3%
Strict ND	Professional & Personal Services	55.4	50.0	45.0	-5.0	-10.0%
Strict ND	Telecom	111.6	106.6	97.4	-9.2	-8.7%
Strict ND	Utilities	190.1	182.4	173.3	-9.2	-5.0%
Strict ND	Flights	32.5	24.5	23.5	-0.9	-3.9%
Nondurable	Miscellaneous Nondurables	308.6	276.6	268.5	-8.1	-2.9%
Durable	Insurance	151.6	159.0	154.6	-4.4	-2.8%
Other Bank Account Outflows						
	Transfer to External Account	356.1	271.6	237.3	-34.3	-12.6%
	Uncategorizable Electronic	635.2	485.4	441.9	-43.6	-9.0%
	Paper Checks	1,057.6	968.9	923.7	-45.2	-4.7%
	Non-Chase Credit Card Bill	436.8	365.2	351.1	-14.1	-3.9%
	Installment Debt	380.9	348.7	335.3	-13.3	-3.8%

Notes: n=27,740 households who exhausted UI benefits and meet the sampling criteria described in Section 2.1. This table decomposes the drop in spending during unemployment into 27 categories. Column 1 is three months prior to the first UI payment, column 2 is the month before UI exhaustion and column 3 is the month after UI exhaustion.

a. Spending categories of strict nondurable, other nondurable, and durable from Lusardi (1996). Cash withdrawals and miscellaneous nondurables are included in the headline nondurables series.

b. See online Appendix B for additional details.

Table 3: Model Estimates

	Standard Model	Standard Model	Heterogeneity in β	Heterogeneity in δ
	(1)	(2)	(3)	(4)
Number of Types	2 types	2 types	4 types	4 types
Calibrated Consumption Parameters				
Risk Aversion γ[a]	2	2	2	2
Naive Hyperbolic Discount Factor β	1.000	--	--	1
Estimated Consumption Parameters				
Monthly Exponential Discount Factor δ	0.9898	0.9898	0.9951	{0.6003, 0.9894}
	(0.0003)	(0.0003)	(0.0001)	(0.0158, 0.0003)
Naive Hyperbolic Discount Factor β	--	1.000	{0.522, 0.899}	--
		(0.007)	(0.025, 0.026)	
Borrowing Limit a	4.5	4.5	6.1	7.8
	(0.2)	(0.2)	(0.6)	(0.3)
Impatient/Myopic Population Share	--	--	0.25	0.17
			(0.01)	(0.01)
Estimated Search Parameters				
Cost of Job Search k	{9.0, 129.5}	{9.0, 129.5}	{4.7, 53.6}	{4.5, 55.0}
	(1.8, 42.3)	(1.8, 42.3)	(1.0, 20.3)	(0.8, 12.9)
Convexity of Job Search Cost ξ	1.4	1.4	1.1	1.4
	(0.1)	(0.1)	(0.1)	(0.1)
Low Job Search Cost Population Share	0.67	0.67	0.79	0.59
	(0.04)	(0.04)	(0.03)	(0.05)
Model Fit				
N Moments	27	27	27	27
N Estimated Parameters	6	7	9	8
Consumption Goodness of Fit	350	350	99	148
Search Goodness of Fit	81	81	86	97
Total Goodness of Fit	431	431	186	246

Notes: This table presents parameter estimates of models of consumption and job search during unemployment. The model is described in Section 4.1 and is fit using equation (8) to the data on spending and job search during an unemployment spell from Figures 3 and 7, respectively. Columns 1 and 2 allow for unobserved heterogeneity in job search costs. Columns 3 and 4 allow for unobserved heterogeneity in time preference parameters by allowing for differences in β and δ respectively. In column 1, β is fixed at 1, while in columns 2 and 3, β is estimated and is constrained to be between 0 and 1. Similarly in column 4, δ is estimated and is constrained to be between 0 and 1. Goodness of fit total may not be sum of components due to rounding. Standard errors of estimated parameters in parentheses.

a. Calibrated from Carroll (1997)

Table 4: Welfare Impact of Changes in UI Generosity

Welfare Change as an Equivalent Increase in Lifetime Income

	ΔWelfare - UI Benefit Increase	ΔWelfare - UI Duration Extension	Difference (col 2 - col 1)	Ratio (col 2 / col 1)
	(1)	(2)	(3)	(4)
Consumption-Smoothing Gains Only				
JPMCI Nondurables	0.021%	0.082%	0.061%	3.94
Gruber (1997) Food	0.019%	--	--	--
Consumption-Smoothing Gains and Fiscal Externality Loss				
JPMCI Nondurables	-0.023%	0.016%	0.039%	--
Gruber (1997) Food	-0.025%	--	--	--

Notes: We evaluate the welfare impact of budget-neutral tax-financed changes in the generosity of UI benefits as a percentage of lifetime income for CRRA utility with risk aversion of 2. We use a sufficient statistic formula which generalizes the Baily-Chetty formula to allow for finite duration of UI benefits. See Section 5 for details.
Rows 1 and 2 show the welfare gains in the absence of a fiscal externality from UI benefits. We compare a UI benefit increase of 1.77 percent of monthly employed income with a one-month extension of UI benefits. These changes have the same fiscal cost and require a tax increase of 0.14 percent of monthly employed income to fund.
Rows 3 and 4 show the welfare gains with taxes adjusted for the fiscal externality arising from moral hazard in response to changes in UI benefits. Fiscal externalities are the median estimates in the literature review in Schmieder and Von Wachter (2017). We compare a UI benefit increase of 2.03 percent with a one-month extension of UI benefits. These change have the same fiscal cost and require a tax increase of 0.21 percent to fund.