Title
Essays in Empirical Macroeconomics

Author
Flamang, Niklas

Publication Date
2023

Peer reviewed|Thesis/dissertation

Essays in Empirical Macroeconomics

by

Niklas Flamang

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Economics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Emi Nakamura, Chair
Professor Yuriy Gorodnichenko
Associate Professor Amir Kermani

Summer 2023

Essays in Empirical Macroeconomics

Abstract

Essays in Empirical Macroeconomics

by

Niklas Flamang

Doctor of Philosophy in Economics

University of California, Berkeley

Professor Emi Nakamura, Chair

In this dissertation, I explore the consumption smoothing behavior of households. A key problem at the core of much of economics is how households allocate consumption over time and across states of the world: Do they borrow in lean times and save in good times? Do they save when young and dis-save when old? Do they account for future changes in income and adjust consumption right away? Studying consumption behavior at retirement and delinquencies during unemployment, I find that households cut their consumption substantially at the onset of retirement and that the sensitivity of delinquencies to unemployment appears to mostly be driven by a lack of liquidity during unemployment. Together, these results suggest that households exhibit a large degree of sensitivity to current income, a component largely ignored by state-of-the-art consumption models.

To Jael, Ella, and Noah

Thank you for your patience even when I did not have any.

Contents

List of Figures

List of Tables

Acknowledgments

This thesis would not have been possible without the continuous advice and support of Emi Nakamura. I am indebted to Yuriy Gorodnichenko for reading many iterations of parts of this dissertation and to Amir Kermani for countless invaluable conversations. I am thankful to Jón Steinsson for advice and encouragement in times good and bad. Special thanks to Sree Kancherla for oh-so-many conversations about economics and everything else. I am grateful to Pedro Pires for many distractions and an excellent taste in music. For valuable comments and suggestions, I thank Hadar Avivi, Stefano DellaVigna, Peter Ganong, Erik Hurst, Pat Kline, Quan Le, Michael Love, Enrico Moretti, Cristobal Otero, Jonathan Parker, Adam Rosenberg, Jesse Rothstein, Emmanuel Saez, Benjamin Schoefer, Damian Vergara, Chris Walters, Harry Wheeler and Danny Yagan as well as numerous seminar participants. I am thankful to Moritz Schularick for setting me on a path to pursuing a PhD in economics and to countless faculty for their help and encouragement over the years. I am grateful to the many people who commiserated with me along the way. Most of all, I thank my family for always bearing with me, even through the long nights and short weekends.

Chapter 1

Introduction

One of the core questions of macroeconomics is how households smooth their consumption in light of changes to their income. This holds for variations in income both at high frequency and at low frequency. For example, do households with highly seasonal incomes accumulate savings in good times when income is high and then dis-save in leaner months? How do households deal with the fact that their life-time income has a distinct hump shape, with little earnings early in their (working) lives, very high earnings in middle age, and relatively low incomes later in life.

Fundamentally, the question of household smoothing behavior will depend on three pieces. First, what does the household's utility function look like? For utility functions with little concavity, allocating consumption across time becomes much less important as it makes little difference whether a household consumes a lot in some periods and relatively little in other periods. On the other hand, a lot of concavity in the utility function will imply that households really want to keep consumption very steady across periods, as volatile consumption would lead to large losses in utility relative to a more stable consumption path.

Second, what is the prevailing interest rate (or are the prevailing interest rates) on saving and borrowing? If interest rates are high, financing consumption today by borrowing against future income is very costly. Households will save more (or borrow less), thereby moving consumption from today into the future (relative to a scenario with lower interest rates). On the other hand, if the interest rate is very low, consumption can be moved around in time almost costlessly.

The third key parameter when thinking about the household's consumption problem is the discount factor. How do households value consumption at different points in time? Maybe they have a strong preference for consumption today over consumption at some point in the future, so they save relatively little (or borrow a lot) so that overall movements in consumption track movements in income quite closely. On the other hand, households may be very patient and value consumption in the future as highly as they value consumption today.

These three components are captured in the standard Euler equation that economics students learn in an intermediate macro class:

$$u'(c_t) = (1 + r)\beta u'(c_{t+1}) \tag{1.1}$$

Here, $u'(c_t)$ is the marginal utility of consumption c in period t, $(1 + r)$ is the gross interest rate, and β is the discount factor. More sophisticated versions will include expectations operators for future periods, but the general intuition remains the same: A household's intertemporal consumption problem boils down to the concavity of the utility function, the interest rate, and the discount factor.

There are obvious modifications to the above problem: Do households have access to financial markets that allow them to freely save and dis-save without prohibitive costs that make doing so impossible? For example, if borrowing against future income growth is very hard (and costly) households will have a much harder time to smooth consumption over the early part of the life-cycle. They will under-consume relative to an ideal path of consumption, solely because they cannot borrow against future increase in income that are all but guaranteed. Of course, poorly developed financial markets might even impede *saving* so that households could have a hard time building up savings to buffer against future declines in income. Similarly, it might be harder to borrow in bad times (e.g., after losing a job) than in good times (e.g., after getting a promotion). But in a stylized world, borrowing during unemployment and saving after large gains in income is exactly what we would expect.

A similar problem arises with respect to idiosyncratic income risk. For example, if job losses do not only imply temporary income losses but *permanently* lower earnings, households face a degree of income risk that is likely to be hard to insure against. In this case, households' consumption problem gets more complicated as it would have to account for the potential of permanent income losses in the future.

Given that the question of households' consumption behavior is one of the core questions of economics, it is not surprising that the questions discussed above have received considerable attention. It was Modigliani and Brumberg (1954) who first formulated the life-cycle hypothesis, arguing that households accumulate assets when they are young and draw them down when they are old. Friedman (1957) formulated a theory of life-cycle consumption that rests on the idea that households consume *permanent* rather than *current* income, so that households attempt to equalize the marginal utility of consumption in every period. Later work by Hall (1978) showed that an important implication of the life-cycle permanent income hypothesis is that consumption will only respond to unanticipated changes in income, thereby representing a martingale process.

These theories notwithstanding, empirical research has suggested over and over that the life-cycle profile is hump-shaped (e.g., Carroll and Summers 1991, Gourinchas and Parker 2002. At higher frequency, we also have many empirical results at odds with the theory of households equalizing marginal utilities in all periods. If households smoothed consumption the way the life-cycle model of consumption predicts, then we would expect the marginal propensity to consume out of small and temporary income shocks to be roughly zero. However, empirical research suggests that the MPC out of these shocks is far larger than zero

(Ganong and Noel (2019), Ganong et al. (2020a), Gerard and Naritomi (2021), Johnson et al. (2006), Parker et al. (2013).

These results have given rise to a number of theories why the life-cycle permanent income might fail. Early, and prominently, Deaton (1991) and Carroll (1997) proposed the buffer stock saving model. Here, households face borrowing constraints so that they cannot borrow as much as they would like in order to smooth consumption. For low levels of wealth, these buffer-stock savers will be very sensitive to changes in income. Once households have built up their buffer stocks, they will look more like LC-PIH households and smooth fluctuations in income. Another popular theory that can explain high marginal propensities to consume out of current income is that of wealthy hand-to-mouth consumers (Kaplan and Violante 2014): With access to one high-return illiquid asset and one low-return liquid asset, households will rationally "lock up" a lot of their wealth in the illiquid asset. Being relatively illiquid, these housholds will exhibit large marginal propensities to consume out of transitory income shocks. Another possible explanation is behavioral consumers: Maybe households are impatient and do not save as much as a standard model would predict. In fact, the hump shape of the life-cycle profile of consumption was one of the original motivations for Laibson's (1997) model of quasi-hyperbolic discounting.

In this dissertation, I investigate the question of household smoothing behavior by studying two important questions. In Chapter 2, I revisit the consumption-retirement puzzle. A number of studies have found that expenditure falls precipitously at the onset of retirement (e.g., Banks et al. (1998) Bernheim 1987, Hamermesh 1984, Hausman and Paquette (1987), Olafsson and Pagel (2018b), Stephens and Toohey (2018)). Since retirement represents a fully anticipated decline in current income, this finding is starkly at odds with the life-cycle permanent income hypothesis. However, there is still open debate to what extent the post-retirement *expenditure* drop represents a true drop in *consumption*. If households shop for lower prices (cf., Aguiar and Hurst 2007), expenditure will be a systematically biased measure of consumption, thereby overstating the true decline in consumption after retirement. Similarly, households may cut back on work-related expenses and engage in more home production, so that for any given level of expenditure, consumption does not actually decline by much (Aguiar and Hurst 2005, Aguiar and Hurst 2007, Hurd and Rohwedder 2005). I find that neither shopping for lower prices nor substitution towards home production are the main drivers of the decline in expenditures at retirement. Rather, most of the decline in expenditure is a true decline in *consumption*.

In Chapter 3, I study whether large increases in the generosity of unemployment insurance benefits insulate local financial conditions from local business cycle conditions. One of the main motivations for providing unemployment insurance is a failure of credit markets: Households who face unemployment may find themselves unable to borrow in order to smooth consumption. Without this mechanism, providing liquidity to households would be almost "worthless" as households could simply access credit markets to hold them over until they find a new job. However, another interpretation might be that unemployment often leads to permanent losses in income. In this case, it is unclear whether the lack of access to credit is actually a market failure. After all, households may not be able to borrow to finance their pre-unemployment levels of consumption simply because their lifetime income has declined substantially. One way to think about this question is to investigate what happens during large benefits extensions. If the main channel is one of market failures in financial markets, then increased generosity in UI should lead to more smoothing and smaller declines in consumption. On the other hand, if the main channel through which unemployment affects consumption is permanent income, then increases in UI generosity are very small relative to the effects on permanent income. This would imply that consumption still falls precipitously during unemployment. We tackle this problem by looking at the sensivity of county-level delinquencies to county-level unemployment rates during the expansion and subsequent withdrawal of Covid-era unemployment insurance policies. Here, we find that local financial conditions are much less sensitive to local economic conditions in a regime with substantially higher replacement rates. One interpretation of this evidence is that it is unlikely that the effect of unemployment on permanent income is driving the sensitivity of delinquencies to the unemployment rate.

In Chapter 4, I discuss the implications of my findings for current consumption models and also give an outlook for promising future avenues of research. In particular, it seems that while much progress has been made in devising more realistic models of consumption, economists still have not found a tractable model of household consumption behavior that captures many empirical regularities. This is important at the micro level, for example when thinking about the optimal design of social insurance policies such as retirement programs and unemployment insurance. It is also important at the macro level, where much of the transmission channel for monetary and fiscal policy runs through households with marginal propensities to consume that are substantially larger than zero.

Chapter 2

Revisiting the Consumption-Retirement Puzzle

Disclaimer: This chapter includes analyses calculated (or derived) based in part on data from Nielsen Consumer LLC and marketing databases provided through the NielsenIQ Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the NielsenIQ data are those of the researcher(s) and do not reflect the views of NielsenIQ. NielsenIQ is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein.

In this chapter, I revisit an old question: What happens to consumption at the onset of retirement? A number of studies have found that expenditure falls precipitously at the onset of retirement (e.g., Banks et al. (1998) Bernheim 1987, Hamermesh 1984, Hausman and Paquette (1987), Olafsson and Pagel (2018b), Stephens and Toohey (2018)). This is a very sharp test of the validity of many consumption models for the following reasons: First, retirement is almost always an expected decline in income. Households know that they will have access to lower incomes in retirement than in their (peak) earning years. Second, retirement is only a change in *current* income while it leaves permanent income unchanged. While Social Security and other pension income are unlikely to fully replace labor income, this effect should be fully accounted for by a household that is consuming *permanent* rather than current income. In fact, this logic was one of the main motivations for Modigliani and Brumberg (1954) who argue that young households will be net savers while old households will be net dis-savers. Third, failures to smooth the decline of income at retirement cannot be explained by credit frictions. Unlike tax refunds, tax rebates, or lottery winnings, retirement represents a *decline* in income, not a temporary increase. Therefore, failures to smooth consumption around retirement are more likely to be indicative of a failure to save rather than credit market failures that prevent households from borrowing. Retirement also occurs after practically all labor market earnings have already realized, ruling out uninsurable income risk as a candidate explanation for the sensitivity of consumption to income.

2.1 Introduction

Every day, over 10,000 people in the United States reach retirement age. Given the large size of the Baby Boomer generation, this pattern will continue for at least another fifteen years. The aging population has wide-reaching implications for the fiscal position of the United States, its future growth trajectory, and the welfare of the soon-to-be retirees themselves. This latter effect is a source of immediate concern especially because Covid-19 has pushed many people into retirement early. This wave of retirements may have have big welfare consequences as consumer expenditures decline sharply in the first years of retirement (the "post-retirement consumption drop"). Whether this decline is a measurement artifact (e.g., retirees can search for and pay lower prices thus driving a wedge between consumption and nominal spending) or a genuine decrease in consumption still is a matter of debate. Using very rich scanner data covering sixteen years, I document that price effects are relatively small and thus a significant decrease in consumption is likely real.

In this paper, I combine evidence from rich scanner survey data to show that expenditures fall starkly at retirement. In the first six years of retirement, total non-durable expenditures fall by a total of thirteen percent while scanner-data covered expenditures (which mostly spans food at home broadly defined) fall by around seven percent in the same time frame. I then leverage the scanner data to decompose household expenditure patterns. To do so, I define three margins of shopping behavior: price effects, bulk effects, and quality adjustments. Price effects capture purchasing a given barcode-level item at a lower price. This can be achieved through buying items on sale, using coupons or traveling to cheaper stores. Bulk effects are savings that arise from a households' allocation of purchases to bulk quantities (and thus lower per-unit prices). Quality adjustments represent moves along the quality ladder for very similar items (e.g., buying pasture-raised eggs vs. battery-cage eggs). Equipped with these measures of shopping adjustment, I re-estimate the expenditure drop after correcting for price and bulk effects. Three years into retirement, price savings and bulk savings account for a 0.25 percentage point drop in expenditure while uncorrected expenditures have fallen by 2.85 percent. Six years into retirement, some of these savings have faded away so that price and bulk savings contribute 0.2 percentage points to the total expenditure drop of 8 percent. Together, price, bulk, and quality effects contribute 1 percentage point to the expenditure drop at three years and 1.25 percentage points after six years. Nevertheless, even correcting expenditures for all of these margins, seven percentage points of the expenditure drop after six years are the result of changes in *quantities*.

These expenditure and quantity declines are not explained by either work related expenditures or substitution towards home production: Expenditures on goods unrelated to work fall and expenditures on home production inputs also decline, a result that holds both within scanner and survey data. In addition, there is little evidence of expenditure timing: Durable expenditures fall by substantially more than non-durable expenditures while there is no evidence of large durable purchases right before retirement. All of these results point towards the post-retirement expenditure drop representing a true decline in consumption.

To investigate how these patterns vary across households, I perform a variety of tests.

Expenditure and quality drops are smallest for those households that see the smallest post-retirement income drop, pointing towards a significant role for current income in explaining consumption dynamics. Similarly, households with higher wealth are more insulated from the post-retirement income drop and have smaller expenditure drops than households with lower levels of saving, particularly with respect to non-durables. For price and bulk purchasing adjustments, the patterns are less stark. Price adjustments are largest for households with the largest post-retirement income drop, but beyond this group it is not clear if prices fall the most for households for whom expenditures decline the most. Consequently, raw expenditure adjustments look very similar to price-and-bulk corrected expenditure adjustments for most households.

Analyzing whether shopping effects are important in other contexts, I find that households with female heads not working full time pay substantially lower prices, but that much of that effect seems to be driven by fixed differences across households rather than within-household variation over time. Tracing out shopping behavior over the course of an non-employment spell, I find that similar to the case of retirement, the primary margin of shopping adjustment is the quality of the consumption bundle rather than purchasing identical products at lower prices. These results suggest that expenditures are generally a good proxy for consumption expenditures.

My findings have important implications: The large drop of consumption at the onset of retirement is not only inconsistent with the life-cycle permanent income hypothesis, it is also not explained by consumption models that incorporate credit market frictions and uninsurable income risk. While credit frictions can explain why households cannot smooth consumption early in life, the stark drop in consumption at retirement is inconsistent with consumption smoothing. In addition, the consumption drop at retirement is not explained by consumption models with little holdings of liquid assets that result from return differentials between high-return illiquid and low-return liquid assets: First, many illiquid assets become liquid at retirement (e.g., 401k plans and IRAs, pension plans). Second, the decline of *current* income at retirement is permanent, making it hard to justify based on one-time adjustment costs of changing the asset allocation. However, consumption adjustments are smaller for wealthier households, suggesting an important role of differences in wealth accumulation in explaining the drop. Consistent with prior results in the literature, households who see the smallest post-retirement income drop also see consumption fall the least, implying the need for consumption models that can generate sensitivity to current income beyond rational illiquidity or credit market frictions.

Beyond our understanding of consumption behavior, my findings also have important welfare implications. With the baby boomers starting to enter retirement, the United States economy will see four to five million retirements every year over the next fifteen years. To the extent that these retirements go along with large declines in expenditures, this pattern is likely to present a future drag on domestic consumption and hence total output growth. The apparently large sensitivity to current income also suggests that the optimal design of old-age retirement schemes should consider disbursement policies as an important policy lever. Many households do not annuitize their security holdings and they are also unwilling to "eat their house" as evidenced by the relative dearth of reverse mortgage products. Designing private retirement accounts so that asset holdings are easily transformed into constant cash flows seems like a promising way of preventing excessively large expenditure declines in retirement. More generally, my findings further suggests that retirement schemes that guarantee constant cash flows may have substantial welfare benefits relative to systems without recurring payouts.

The main contribution of this paper is to investigate to what extent the post-retirement expenditure drop is a result of measurement error or if it is a real drop in consumption. Starting in the 1980s, research has documented that households appear to under-save for retirement resulting in substantial declines in consumption in old age (Hamermesh 1984, Bernheim 1987, Hausman and Paquette 1987). This finding has sparked a wave of interest in explanations of the post-retirement expenditure drop that are consistent with the life-cycle permanent income hypothesis. These explanations are broadly categorizable into three groups: First, households' may have a lower marginal utility of consumption later in life, for example resulting from changes in household composition and aging itself. Since optimal smoothing of life-cycle consumption predicts the equalization of *marginal utilities*, not consumption levels, this may explain why expenditures drop so much later in life (Banks et al. 1998). Second, changes in expenditures might be explained by declines in work-related expenses. If the post-retirement expenditure drop is driven by these, then utility-relevant consumption has not fallen at all (Hurd and Rohwedder 2005). Third, retired households may exert shopping effort to lower the prices they face and engage in more home production so that lower expenditures translate into constant consumption levels (Aguiar and Hurst 2005, Aguiar and Hurst 2007, Hurst 2008).

Notably, neither of the first two explanations seem sufficient to explain the retirement savings puzzle. Banks et al. (1998) show that expected changes in household composition and mortality are not sufficient to equalize pre-retirement and post-retirement marginal utilities of consumption in British micro data. In addition, while work-related expenses in their data are responsible for a large share of the post-retirement expenditure drop, expenditure in *all* categories of consumption falls. More recently, Olafsson and Pagel (2018b) use data from a personal finance aggregator to show that spending in both leisure and work-related categories of consumption falls while savings actually *increase* in retirement. Last, Stephens and Toohey (2018) leverage forty years of cross-sectional data as well as longitudinal evidence to show that caloric intake falls during retirement.

In addition to the literature on the retirement-savings puzzle, this paper contributes to the literature on household-level prices and the adjustment of shopping behavior in light of shocks. Aguiar and Hurst (2007) were the first to utilize scanner data to construct household-level price indices. In terms of deconstructing deviations from the price of the average bundle, this paper is closely related to Kaplan and Menzio (2015) and Nevo and Wong (2019). Kaplan and Menzio (2015) use the 2004-09 years of the KNCP to decompose prices into a store component, a store-specific goods component, and a transaction component plus three covariance terms. They find that around 90% of the price dispersion across households comes from the store and store-good components, suggesting that at the household level, the choice of retailer is what is driving price dispersion. In related work, Coibion et al. (2015) show that inflation in *effective* prices paid by consumers declines significantly with higher unemployment rates while posted prices remain relatively unchanged. These effects are mostly driven by households switching between retailers rather than purchasing on-sale items. Another closely related paper is Nevo and Wong (2019) who use scanner data to investigate to what extent costly shopping activities lower household-level prices. Here, too, the household-level price index is defined as a cross-sectional measure that compares the cost of the household's bundle at actual prices to the cost of the same bundle at average prices.

An important difference between this paper and both Kaplan and Menzio (2015) and Nevo and Wong (2019) is that both Kaplan and Menzio (2015) and Nevo and Wong (2019) are agnostic about the "correct" level of aggregation for the price index, allowing or increasingly broad comparisons across products. Meanwhile, I follow Aguiar and Hurst (2007) and think of the household-level price index as a measure of prices paid for identical goods. This allows me to interpret the choice of exact item within a narrowly defined consumption category as informative about the quality ladder (after accounting for potential savings from purchasing in bulk). This interpretation of prices as conveying quality information is similar to the logic of Jaravel (2019) who segments each product category into price deciles, interpreting these as representing a quality ladder.

The results of my event studies around non-employment have important implications for interpreting the expenditure drop during unemployment. Gruber (1997a) was the first to empirically investigate the behavior of consumption over an unemployment spell. Recent work by Ganong and Noel (2019) and Landais and Spinnewijn (2021) has leveraged much richer data sets (bank account data and administrative records on earnings and wealth) to infer the expenditure drop around unemployment. The former find that non-durable expenditures at the onset of unemployment fall by around 8%, with declines of 1% for each additional month of unemployment and another 12% drop in expenditures at the exhaustion of unemployment insurance benefits. Relative to these magnitudes, the savings from shopping behavior are quite small.

The rest of this paper is organized as follows: Section 2.2 discusses the data and provides background on the categorization of expenditures in the Kilts Nielsen Consumer Panel. Section 2.3 lays out the economic logic of my shopping adjustment measures and 2.4 details how I map this to the data. In Section 2.5 I trace out the path of expenditures and shopping adjustments after retirement and show that relative to the drop in expenditures, shopping effects are small. Section 2.6 investigates to what extent shopping adjustments matter when comparing across vs. within-households. Section 2.7 discusses policy implications and avenues for future research.

2.2 Data and Variable Definitions

I use data from two panel surveys: the Kilts Nielsen Consumer Panel (henceforth "Nielsen Panel" or KNCP) and the the Health and Retirement Survey (HRS). The Nielsen Panel is a scanner data set that includes information on prices paid at the barcode-by-trip level and includes information on exact consumption bundles at grocery stores broadly defined. The HRS's Consumption and Activities Mail Survey (CAMS) provides information on virtually all Consumer Expenditure Survey categories of consumption and also elicits time use in a similar way to the American Time Use Survey.

Kilts Nielsen Consumer Panel

The main data set for this paper is the KNCP; a data set that covers an annual panel of households from 2004 to 2019. In total, the data set includes 194,551 households and 917,962 household-year observations. The KNCP includes date-exact information about purchases at grocery stores, supermarkets, discount stores, superstores and similar store categories. An annual set of about 60,000 households records purchases by logging each item, providing the exact UPC or providing additional information about data on goods like raw produce that is being sold by weight.[1] Purchases are recorded at the trip level, with information about the price paid, whether an item was recorded as being on sale (associated with a "deal"), whether a coupon was used, the exact number of items purchased and the overall expenditure for the shopping trip. Households are provided with financial incentives for their participation in the KNCP sample, and they may drop out of the sample at the end of a panel year or may continue from one year to the next. A panel year stretches from the last days of December of one year to mid-to-late December of the following year, which implies that a panel year aligns very closely, but not perfectly with the calendar year.

Once a year, in the fourth quarter preceding the data collection of data for a panel year, households are asked a variety of demographic questions. This includes household income, household size, whether a male and female household head are present, questions about household members' ages, education, occupation, a variety of information about living conditions (e.g., the kind of residence a household lives in, the availability of internet and TV service, the presence of a variety of household appliances). The exact date of the collection of this demographic information is not provided to researchers, so I will interpret all of these demographic variables as representing the data for the fourth quarter of the year these data were collected. Households may exit or enter the panel in any given year. About 80% of participants remain in the sample from year to year, and many households remain for the sample for substantial amounts of time. The average number of years in the panel for 4.7 years, and the average number of continuous years in the sample is slightly more than 4 years.

[1]Throughout, I will restrict attention to purchases associated with a UPC code.

Product-Level Information

The main benefit of the KNCP is the level of granularity for purchases. Purchases and prices are recorded at the trip by UPC (barcode) level.[2]. While an increasing number of shoppers is asked to record all purchases (including items without an associated UPC code), throughout I will restrict attention to purchases with an associated UPC. There are about 3.5 million unique UPCs in the KNCP, which are mapped to 1,298 product modules in 110 product groups and 9 departments. Product modules are the lowest level of aggregation and correspond to very fine categories consumption. For example, frozen orange juice, fresh orange juice, fresh apple juice, and sugar-sweetened fruit beverages each are different product modules. Figure 2.A1 presents the product hierarchy in Nielsen, going from departments to product groups to product modules to individual UPCs. For every barcode, Nielsen provides information about the size or weight of a product, and the associated unit of measurement (e.g., ounces, milliliters, square feet, or counts). In addition, Nielsen records whether a UPC corresponds to a multi- or single-pack. This allows me to construct exact quantities for every UPC in the data. In turn, I can construct per-unit prices for each UPC to make prices comparable between products of different sizes. For example, one individual can of Coca-Cola, one two liter bottle of Coca-Cola and one 24-pack of cans of Coca-Cola will all be associated with different UPC codes. Using the quantity information provided by Nielsen, I can compute exact per-ounce of soda prices for each of these. Figure 2.A3 provides an example.

Baseline Sample

For the baseline sample, I only include households who are observable between the ages of 25 and 74 for at least five years so that I can estimate within-household variation. Because I do not know who the "primary shopper" in a given household is, I define the age of a household as the average age of the household heads. I further restrict the sample to households for which I observe shopping trips in at least 11 months out of the year and real spending of at least $250 per year (in 2012 dollars as deflated by the CPI for food at home). This leaves me with 179,703 households and 814,938 household-year observations, where 121,553 households are observed for multiple years. In Column 1 of Table 2.1, I present summary statistics for this sample (weighted using KNCP projection factors). We can see that the weighted KNCP matches the US population quite well on most observables: the sample is broadly representative of the US as a whole in terms of income, household size, and other demographics like race or Hispanic origin (conditional on age).[3] In addition, about nine

[2]Sometimes, the same UPC will correspond to a different product in different years. Most notably, this is true for changes in the size or weight of a product. In these cases, Nielsen constructs a "UPC version" variable that assures that a UPC captures identical items. While I construct all my measures at the UPC-by-UPC version level, I will simply refer to this as "the UPC level" for expositional purposes.

[3]An important caveat is that the unweighted KNCP skews older and substantially over-represents female heads of household.

percent of households experience a retirement while in the panel, with another ten percent of households undergoing non-employment spells while in the Nielsen Consumer Panel.

Retirement Event Study Sample

As noted above, all demographic information is collected in the fourth quarter *preceding* the panel year. That is, if a household is in the Nielsen Consumer Panel for panel year 2004, then the associated demographic information will have been collected during the fourth quarter of 2003. Therefore, employment flows can only observed if the change in employment status covers the fourth quarter of any given year.[4] A second problem is that Nielsen does not actually record information about different kinds of non-employment, but rather reports a single non-employment category that combines unemployment, retirement, and voluntarily staying at home.

To identify retirements, I focus on household heads who have been employed in year $t-1$, who are between ages 60 and 70 in year t, and who are still not employed in year $t+1$.[5] Given these restrictions, all households in my retirement sample will have to be in the data for at least three consecutive years. I impose that household heads cannot "unretire": If I observe multiple retirements for a male or female household head, I only keep the first retirement. For non-concurrent retirement spells (that is, both heads of household retiring at different points in time), I treat the first retirement as the "treatment". Of these households, I then restrict the sample to household-year observations in which the household records at least one shopping trip for each month of the year. This leaves me with a sample of 10,007 households ever entering retirement. Control households are households with at least one working household head between the ages 55 and 70 for which I do not observe either a transition from employment to non-employment. We can think of these household as not yet retired or having at most one household head who retired before entering the Kilts Nielsen Consumer Panel.

In Column 2 of Table 2.1, I present summary statistics for the retirement event study sample. A household experiencing at least one retires undergoes 1.02 retirements on average and remains in the sample for 10.4 years. Otherwise, the retirement sample looks broadly similar to the baseline sample with most differences stemming from the fact that households in the retirement sample are substantially older. In Figure 2.A2, I present the distribution of retirement ages according to my assignment of retirements. Similar to cases with information on actual retirements, retirements spike at age 62—the earliest age at which

[4] A less consequential problem is that the exact date or even month during which the demographic information was collected is unknown.

[5] The logic for the age cutoffs is as follows. At age 59.5, workers are old enough to make penalty-free withdrawals from tax-advantaged retirement accounts (and often retire with pension benefits). From age 70 onward, there is no benefit to delaying claiming Social Security anymore. Therefore, I consider any age outside of this range as "unusual" or potentially driven by labor market shocks that aren't really about the life-cycle. Similarly, I require two years of consecutive non-employment to make sure my measure of retirements does not reflect unemployment spells.

workers become eligible for Social Security benefits. However, by construction my measure does not pick up the asymmetry of retirements present in data sets that explicitly elicit the age at retirement (the vast majority of retirements occur up until age 65 with relatively few retirements thereafter).

The Health and Retirement Survey

In addition to the Kilts Nielsen Consumer Panel, I will also leverage the Health and Retirement Survey and its Consumption and Activities Mail Survey module. The HRS is a biannual longitudinal panel of a nationally representative sample of households with heads ages 50 and older. I combine information on household demographics (household composition, age, race, ethnicity), household income, household assets, and retirement status from the main HRS sample with additional information from the Consumption and Activities Mail Survey (henceforth "CAMS"). CAMS households also provide information on very detailed categories of consumption (CAMS covers virtually all CEX categories of consumption) as well as time use similar to the American Time Use Survey. Given the rich information, this allows me to construct a bi-annual panel of households that includes information on assets, earnings, retirement status, consumption, and time use for the same household. My sample consists of all households in CAMS. In Table 2.2, I present summary statistics for not retired workers, retired workers, and the full sample. We can see that the sample is roughly comparable to the Nielsen Panel although CAMS households have somewhat lower incomes, fewer years of education, and are more likely to be a minority. The outcomes of interest in the HRS sample will be total expenditures, total non-durable expenditures, and time use on market and non-market activities.

2.3 Theoretical Framework

To motivate the empirical approach of this paper, I will present a simple model of shopping behavior that clarifies the empirical objects defined in the next section. Consider a household i that faces the following optimization problem

$$\max_{\boldsymbol{q},\boldsymbol{s},\boldsymbol{a}} u\big(c(\boldsymbol{q})\big) - h(\boldsymbol{s}) \text{ s.t. } e(\boldsymbol{q},\boldsymbol{s}) = \bar{y}, \tag{2.1}$$

where $\boldsymbol{q}$ denotes the households' consumption bundle, $c(\cdot)$ is an increasing concave function, $\boldsymbol{s}$ is shopping behavior with the disutility cost of lowering per-unit prices with $h(\boldsymbol{s})$ being an increasing convex function. In other settings $h(s)$ is often taken to be the opportunity cost of time, but it can principally be more general and capture storage costs or dis-utility from shopping at less nice stores. Lowering per-unit prices can be either achieved through paying less for a given item or by purchasing in bulk. e denotes expenditures and $\bar{y}$

is some exogenous budget constraint.[6] Expenditures are a function of the chosen bundle $\boldsymbol{q}$ and household-level prices (which depend on shopping behavior):

$$e_i = \sum_m \left(\sum_{k \in m} p_{i,k}(s_i) q_{i,k} \right) \tag{2.2}$$

where k denotes varieties of a given consumption category m, $p_{i,k}(s_i)$ is a household-specific price (where higher s_i implies weakly lower prices), and $q_{i,k}$ is household i's quantity of item k. Note that households do not derive utility from paying lower or higher prices outside of the relaxation of the households' budget constraint. However, within a consumption category m, households prefer higher-quality items which have higher prices *on average*. Therefore, $\bar{p}_k$, the average price of item k, is informative about the quality of item k relative to all other items $l \neq k \in m$. The utility from consumption of market goods is given by:

$$u(c_i) = u\big(c(\bar{p}_1 \cdot q_{i,1}, \ldots, \bar{p}_K \cdot q_{i,K})\big) \tag{2.3}$$

Therefore, we can define a measure of "consumption expenditure" that captures only the utility-relevant aspect of expenditures (leaving aside the dis-utility cost of finding the best deals or buying larger bulk quantities).

$$c_i^* = \sum_m \sum_{k \in m} \bar{p}_k q_{i,k}^*, \tag{2.4}$$

where $q_{i,k}$ are components of the consumption bundle $\boldsymbol{q}_i^*$ which is defined as

$$\boldsymbol{q}_i^* = \arg\max_{\boldsymbol{q}_i} u\big(c(\boldsymbol{q}_i)\big)$$

Note that equipped with this measure of consumption expenditure we can re-write 2.2 to get

$$\begin{aligned} e_i &= \sum_m \sum_{k \in m} \big(p_{i,k}(s_i) - \bar{p}_k\big) \cdot q_{i,t} + \sum_m \sum_{k \in m} \bar{p}_k \cdot q_{i,k} \\ &= c_i^* + \sum_m \left(\sum_{k \in m} \big(p_{i,k}(s_i) - \bar{p}_k\big) \cdot q_{i,k} \right) \end{aligned} \tag{2.5}$$

[6]This budget constraint could refer to rules of thumb such as consuming a constant fraction of income every period, but it may also differ across states of the world (e.g., retired vs. working, unemployed vs. employed).

where $\sum_m \Big(\sum_{k \in m} \left(p_{i,k}(s_i) - \bar{p}_k\right) \cdot q_{i,k} \Big)$ is a wedge between consumption-relevant expenditure and expenditure. Household i may exert higher shopping effort so that it faces lower prices for each item k. In that case $p_{i,k}(s_i)$ will be less than zero and measured expenditure will understate the true amount of consumption expenditure for household i. Similarly, household i may allocate a large share of the budget to bulk purchases, thereby lowering overall costs. Here, too, $p_{i,k}(s_i)$ will be small, and measured expenditure for household i will understate consumption relevant expenditure for household i. The reverse holds for low levels of shopping for low prices or small bulk allocations.

Overall, this structure implies that a maximizing household will set $\boldsymbol{q}^*$ and $\boldsymbol{s}^*$ so that the marginal gain from relaxing the budget constraint will equal the dis-utility costs of increasing shopping effort and higher bulk allocations. For example, a household entering retirement may have lower opportunity cost of time so that $h(\boldsymbol{s})$ is lower for any level of s. This would result in that household exerting more shopping effort and paying lower prices. As children move out, a household may also have more available space and hence lower storage costs, again implying lower $h(\boldsymbol{s})$ for any $\boldsymbol{s}$. In that case, the household will buy more in bulk.

Additionally, when picking bundle $\boldsymbol{q}^*$, the household will pick quantities such that the marginal utility of improving the quality of the consumption bundle is equal to the marginal utility from increasing quantities.[7] This is particularly helpful when thinking about consumption of goods captured by the Nielsen Consumer Panel. Much of the covered items represent necessities so that we may not expect to see much adjustment in terms of total quantities. Movements in the quality of the consumption bundle are therefore informative about overall consumption adjustments (including in cases where we would expect a larger fraction of the adjustment to be accounted for by changes in *quantities*).

2.4 Household-Level Prices and the Quality of Consumption

As outlined in the previous section, the main interest of this paper is to decompose expenditures into consumption expenditures c_i^* (themselves composed of quantity and quality effects) and shopping behavior that keeps the quantity and quality of the consumption bundle unchanged (but may lower nominal expenditures). In Section 2.9, I show that we can express log expenditures as:

$$\ln e_i = \text{Quantity Effects}_i + \text{Price Effects}_i + \text{Bulk Effects}_i + \text{Quality Effects}_i \qquad (2.6)$$

Going back to the framework of the previous section, only Quantity Effects$_i$ and Quality Effects$_i$ are utility-relevant while Price Effects$_i$ and Bulk Effects$_i$ are sources of a wedge between ob-

[7]Note that this need not hold for allocating varieties k within a category of consumption m. More likely, the household will be indifferent between spending the marginal dollar on increasing the quantity of some category m and improving the quantity of some other category $n \neq m$.

served expenditure and utility-relevant expenditure. Below, I will discuss how I measure each of these components in the scanner data.

Price Effects: Paying Lower Prices for Identical Goods

My way of measuring shopping effort builds upon the logic laid out by Aguiar and Hurst (2007): With UPC-level information, I can observe whether households pay lower prices for *identical* goods. I will construct cross-sectional price indices of consumption by comparing the realized prices a household actually paid to the identical bundle (at the UPC level) at average prices. Therefore, shopping effort for household i at time t is given by:

$$\text{Price Effects}_{i,t} = \underbrace{\ln\left(\sum_l P_{i,t,k} \cdot Q_{i,t,k}\right)}_{\text{Actual Cost of Bundle}} - \underbrace{\ln\left(\sum_l \frac{\sum_j P_{j,t,k} Q_{j,t,k}}{\sum_j Q_{j,t,k}} \cdot Q_{i,t,k}\right)}_{\text{Cost of Actual Bundle at Mean Prices}},$$

where $P_{i,t,k}$ denotes the price household i paid for UPC k at time t. In the data, purchases are recorded at the shopping trip level, which I aggregate to monthly, quarterly, or annual frequency, so t will correspond to monthly, quarterly, or annual date, with $P_{i,t,l}$ defined as the average price a household paid at that frequency. $Q_{i,t,k}$ is the quantity of UPC k household i purchased in period t. The second term computes the quantity-weighted average price for each UPC k in period t times the actual quantity of UPC k household i purchased in period t. I then normalize this variable to be centered at 0 every period by subtracting by its period-mean:

$$\text{Price Effects}^*_{i,t} = \text{Price Effects}_{i,t} - \frac{1}{N}\sum_j^N \text{Price Effects}_{j,t} \tag{2.7}$$

Intuitively, this measure of adjustment *only* captures prices paid relative to the average. Looking at Figure 2.A4 exerting shopping effort implies paying the pay prices for either store-brand conventional eggs or name-brand pasture raised eggs, not purchasing store-brand conventional eggs instead of name-brand pasture-raised eggs. However, since the price index is defined cross-sectionally, a household need not purchase the same bundle every period. For example, a household could pay exactly the average price for store-brand conventional eggs one period, and then pay exactly the average price for name-brand pasture-raised eggs in the next period. In both cases, the corresponding effort margin would be equal to 0. In Column 1 of Table 2.3, I present moments of the distribution of the effort margin. The standard deviation is around 8.1% while the inter-quartile range is approximately 7.9%. This suggests that there is substantial room to exert shopping effort in order to lower prices.

Bulk Effects: Purchasing Larger Quantities

Another potential way to reduce the per-unit cost of their purchases is for households to purchase in bulk. Note that unlike paying less for a given UPC, this channel of lowering prices requires some storage costs. As such, it is not entirely clear if we should think of bulk-purchases as a pure reduction of per-unit costs or as a costly way of lowering the price of the consumption bundle. When constructing my measure of bulk savings, I proceed as follows: Following Griffith et al. (2009), I break out each product module into five quintiles of the size distribution. I then define UPCs in the top two quintiles (or the Top 40% of the within-product-module size distribution) as bulk items.[8] I then define bulk savings as follows:

$$\begin{aligned} \textit{Bulk Effects}_{i,t} = & \underbrace{\ln\left(\sum_{m\times b}\left(\frac{\sum_{k\in m\times b}\sum_j P_{j,t,k}Q_{j,t,k}}{\sum_{k\in m\times b}\sum_j Q_{j,t,k}}\cdot\sum_{k\in m\times b}Q_{i,t,k}\right)\right)}_{\text{Cost of Bundle at Actual Bulk Share}} \\ & - \underbrace{\ln\left(\sum_{m}\left(\frac{\sum_k\sum_j P_{j,t,k}Q_{j,t,k}}{\sum_k\sum_j Q_{j,t,k}}\cdot\sum_{k}Q_{i,t,k}\right)\right)}_{\text{Cost of Bundle at Average Bulk Share}}, \end{aligned}$$

where i and j denote households, m indexes product modules, b indexes bulk and non-bulk, t indexes time and k denotes UPC-level items. Intuitively, bulk savings are the difference in the cost of the purchased bundle at actual bulk shares (for each product module separately) and the cost of the bundle had the household purchased it at "average" bulk shares, both evaluated at average per-unit prices for a given product-module-by-bulk-level combination. In that sense, the bulk savings measure abstracts away from actual prices—all savings arise solely from allocation a higher budget share to bulk items (assuming bulk items are cheaper). I normalize this variable to be centered at zero every period by subtracting its period-mean:

$$\text{Bulk Effects}^*_{i,t} = \text{Bulk Effects}_{i,t} - \frac{1}{N}\sum_j^N \text{Bulk Effects}_{j,t} \tag{2.8}$$

In Column 2 of Table 2.3, we can see that the potential savings from bulk purchases are quite meaningful: The standard deviation is 6.1% and the interquartile range is 7.4%. This is particularly true in the tails where moving by just 5 percentiles (i.e., from p95 to p90 or from p10 to p5) results in approximately 2.5% lower prices for a given bundle.

Going back to the framework in Section 2.3, it will be convenient to think of the bulk and effort margins as the overall wedge between observed expenditure and consumption expenditure. I also define this aggregate savings measure as follows:

$$\text{Expenditure Wedge}_{i,t} = \text{Price Effects}^*_{i,t} + \text{Bulk Effects}^*_{i,t} \tag{2.9}$$

[8]Note that I use *total quantities* to define item size. For example, a twelve-pack of cans of Coca-Cola will be defined as equaling 144 fluid ounces (12 times 12 fluid ounces), roughly similar to two-liter bottles of Coca-Cola (135.2 fluid ounces).

Since the expenditure wedge combines the two measures of adjustment that keep the quality of the bundle fixed, looking at the distribution of the expenditure wedge gives us a sense of the maximum possible savings through paying lower prices for a given UPC and purchasing a higher fraction of bulk items. The standard deviation of the expenditure wedge is 9.9% and the interquartile range is 11.9%.

Quality Effects: Purchasing Goods with Lower Average Prices

My definition of quality adjustments leverages the richness of the Nielsen Panel. Within a given product module, households can move down the quality ladder and purchase UPCs that have lower prices *on average*. This measure of quality abstracts away from the realized prices any given household pays and compares the average per-unit price of the items the household actually purchased to the average per-unit price of that product-module-by-bulk level. To make this explicit, quality adjustments are given by

$$\text{Quality Effects}_{i,t} = \underbrace{\ln\left(\sum_m \sum_k \frac{\sum_j P_{j,t,k} Q_{j,t,k}}{\sum_j Q_{j,t,k}} \cdot Q_{i,t,k}\right)}_{\text{Cost of Actual Bundle at Mean Prices}} - \underbrace{\ln\left(\sum_{m\times b}\left(\frac{\sum_{k\in m\times b}\sum_j P_{j,t,k} Q_{j,t,k}}{\sum_{k\in m\times b}\sum_j Q_{j,t,k}} \cdot \sum_{k\in m\times b} Q_{i,t,k}\right)\right)}_{\text{Cost of Average Bundle at Mean Prices}_{i,t}},$$

where m denotes product modules, b denotes bulk vs. non-bulk, and k indexes UPCs within a product m times bulk level b. As above j indexes households other than i, t indexes time at monthly, quarterly, or annual frequency, and P and Q denote prices and quantities, respectively. Just like with shopping effort, I normalize the quality measure so that it is centered at 0 in every period:

$$\text{Quality Effects}^*_{i,t} = \text{Quality Effects}_{i,t} - \frac{1}{N}\sum_j^N \text{Quality Effects}_{j,t} \tag{2.10}$$

To illustrate the logic of this measure of quality, consider the example in Figure 2.A6. A every period, the household can move along the quality ladder by purchasing store-brand eggs, name-brand cage-free eggs, store-brand organic eggs, or name-brand pasture-raised eggs. The quality effect metric will compare the cost of eggs at the average price of the eggs actually purchased to the cost of the same number of eggs at the average price of eggs in the same period. As I show in Table 2.3, the quality effects are more variable than even the combined price and bulk effects: The standard deviation of the quality metric is 13.3% and its interquartile range is 16.4%.

In principle, a household could pay very low prices for items that tend to be very expensive on average (e.g., through buying in bulk or waiting for deals, by going to different stores, or by using coupons). Similarly, a household may pay very high prices for eggs that tend to be cheap on average, for example by purchasing eggs at a corner store or by buying eggs that are frequently discounted at full price. In practice, paying higher prices (a high effort margin) and purchasing higher quality items are positively correlated (a correlation of 0.226 at annual frequency). Similarly, purchasing higher quality bundles is correlated with purchasing smaller quantities (0.121 at annual frequency) while buying in bulk and lowering UPC-level prices are not correlated at all.

Validating the Quality Measure

Given the distinction between prices and quality, it is important to validate that my quality metric really does represent differences in the quality of the consumption bundle. For example, one might be concerned that my measure of quality effects overstates differences in product quality and really picks up differences in *prices*.[9] In this section, I will provide evidence that my quality measures really do pick up *quality* differences, not just price effects.

A natural question to ask is whether the my quality metric varies by permanent income. For this purpose, I assign each household an earnings rank based on their average earnings within a given cohort.[10] As we can see in Panel (a) of Figure 2.1, bundle quality is monotonically increasing in lifetime income rank, despite the fact that the quality metric is constructed without levering any information on income (or even on the total amount of expenditure). In addition, the slope of the quality measure that is constructed locally is almost identical to the slope of a national comparison. This suggests that my quality measure picks up true variation in purchasing behavior rather than a positive correlation between local incomes and prices.

Prices paid, on the other hand, are not very strongly related to lifetime income. The slope of the prices and lifetime income profile follows a rough u-shape: Prices paid relative to average are *falling* from the 1st to the 20th percentile, then they are relatively flat from the 15th to 40th percentile, slowly increasing from the 40th to 80th percentile, and then the are rapidly increasing for the highest lifetime income percentiles. For local prices, this pattern in generally true as well. However, the increase in prices for the highest income percentiles is much starker for national than local prices. This suggests that national price comparisons do pick up some of the covariance between local incomes and local prices. Regardless of this fact, the range of prices paid is much more narrow than that of quality and lifetime income

[9] It should be noted that the opposite might be true as well. Using mean UPC-prices may treat quality differences as differences in prices. For example, the same (in a UPC-code sense) gallon of milk could be purchased at Whole Foods or a large discount store. To the extent that the shopping experience at Whole Foods is more pleasant than at the discount store, this may reflect *quality* and not price differences. A similar argument can be applied to commute time.

[10] This way, my earnings ranks are only about household earnings ranks within their age group instead of picking up the age profile of earnings.

ranks explain more than an order of magnitude more of the variation in quality than prices paid.[11]

As a second test show the relationship between my quality and price measures and county-level unemployment rates. Controlling for household characteristics, time and market fixed effects local unemployment rates are highly correlated with the quality of consumption. Going from a county-level unemployment rate of 5 percent to a rate of 10 percent, the quality of the average consumption bundle falls by 5 percentage points. Effective prices, on the other hand, are not strongly correlated with local economic conditions. If anything, it seems like effective prices are *increasing* in the local unemployment rate for unemployment rates above 5 percent. One potential explanation for this is that households in areas with very high unemployment have less access to cars that make cheaper stores more easily accessible.

Finally, we can check to what extent the measures of consumption are correlated with household level measures of socio-economic status. In Table 2.4, I perform six such tests. In the first column, I regress the quality measure on log household income and a set of household level controls. We can see that a one percent increase in current income increases the quality of consumption by 0.06%. To test to what extent current or lifetime income drives this effect, I regress quality on lifetime income rank and current income in column (2). We can see that the coefficient on current income falls by about two-thirds while a one percentile increase in permanent income increases the quality metric by 0.14 percent. From this, we can infer that the quality of the consumption bundle is driven both by permanent and current income. To test whether wealth predicts the quality of consumption, I regress the quality measure on the ZIP code-level house price from Zillow for the subsample of households living in a single family home, a condo, or a co-op.[12] We can see that higher house prices go along with substantially higher quality of consumption, even conditional on income. In the fourth column, I leverage a matched data set between the HomeScan data and the Survey of Consumer Finances (details on the matching procedure can be found in Appendix Section 2.9). Based on their SCF matches, I group households into within-cohort net-worth ventiles and regress the quality of consumption on the imputed rank in the wealth distribution. Higher imputed net worth ranks result in substantially higher consumption quality according to each measure. Conditional on income, going from the 5th to the 15th ventile (so the 25th to the 75th percentile) of the within-cohort wealth distribution results in consuming module bundles of 1.5% higher quality. In the fifth column, I regress the quality measure on five educational attainment indicators: less than high-school, a high-school degree, some college, a college degree, or a post-graduate degree. The quality of consumption is monotonically increasing in educational attainment, even conditional on current income. This is consistent with the fact that more educated households are likely

[11]Regressing the quality and price measures on lifetime income ranks (either linearly or on 100 income rank indicators) yields an R^2 of 0.0629 and 0.0643 for the quality metric but only an R^2 of 0.0021 and 0.0054 for the effort metric.

[12]Unfortunately, the Nielsen Panel data does not elicit home-ownership directly but only asks household whether they live in a single family home, a two-party home, a multi-family home, or a mobile home. Then the survey separately elicits whether households live in a condo or co-op.

to have higher lifetime incomes (even conditional on current income). In the final column, I include all of these variables at the same time. All coefficients are statistically significant and economically meaningful. Current income and current wealth (as measured by house prices and imputed wealth based on the SCF) both predict the quality of consumption; and so do lifetime income ranks and household-level educational attainment. Given that household-level education is highly predictive of the quality of consumption even conditional on current income, lifetime income, and imputed wealth suggests that my quality metric also captures the *preferences* of more educated households to some extent. For example more highly educated households may prefer organic pasture-raised eggs to battery-cage eggs at any level of income. Throughout, I will include a household fixed effect, so none of my results are driven by this potential for preference heterogeneity.

2.5 Decomposing the Post-Retirement Expenditure Drop

In my first set of results, I will focus on dynamic adjustments around retirement. These allow me to follow the same households as it enters retirement and trace out their expenditure profile, cost savings arising from paying lower prices and buying bulk as well as the quality of consumption. My general empirical strategy for the event studies is the following framework:

$$\text{Outcome}_{i,t} = \theta_i + \sum_{\tau=-A}^{B} \mathbb{1}\{T_{i,t} = \tau\} \cdot \delta_\tau + \boldsymbol{X}'_{i,t}\gamma\epsilon_{i,t}, \tag{2.11}$$

where θ_i = is a household fixed effect, δ_τ = are leads and lags for last known date of employment (before retirement or unemployment), and $\boldsymbol{X}$ is a vector of household composition controls (household size, the relationship of the adult heads of household, presence of children). The coefficient of interest are the δ_τ which tells us how the outcome of interest evolves relative to the last year before retirement.[13]

Income, Expenditures, and Shopping Behavior

Throughout, I will present dynamic event study coefficients where coefficients are plotted relative to the last year in which the household member was still employed so that we

[13]Note that I am not including time fixed effects in my baseline results. The retirement-consumption puzzle is about falling expenditures for a given household, not falling expenditures relative to a growth trend. However, I also estimate a two-way fixed effect version of my results in the appendix. Recent work (Borusyak et al. 2022, Callaway and Sant'Anna 2021, de Chaisemartin and D'Haultfœuille 2020, Sun and Abraham 2021) has shown that estimating models like the one above with ordinary least squares will not generally yield unbiased results, particularly in the presence of treatment effect heterogeneity. Therefore, I will actually estimate this model using the estimator proposed by de Chaisemartin and D'Haultfœuille (2020) in my baseline specification.

know that the retirement occurred at some point in period 0.[14] As we can see in Figure 2.2, household income is flat or trending down before retirement and then falls abruptly. An important note is that the income drop is about twice as big in the HRS as it is in the Nielsen Panel. Potential explanations for this are the topcoding of high incomes in the HomeScan data, the fact that higher-educated households are over-represented in the Nielsen Panel, or measurement error in the income measure discussed in Section 2.2. In Figure 2.3, we can see that HomeScan-covered expenditures and the broader set of HRS-covered expenditures follow broadly similar trajectories, particularly after retirement. Six years into retirement, scanner-covered expenditure has fallen by about 8 percent while total non-durable expenditure in the HRS has fallen by around 15 percent. In both cases, this suggest some smoothing by households as incomes fall by substantially more than expenditures. Nonetheless, the marginal propensity out of the post-retirement income drop is large and the declines expenditure are economically meaningful.

Turning the question how households adjust their shopping behavior in response to retirement, Panels (a) and (b) of Figure 2.4 show event study results for savings arising from paying lower UPC-level prices and purchasing larger quantities (i.e., bulk savings). We see that there is very little evidence of pre-trends for prices paid relative to the average while bulk savings fall smoothly through retirement. Four years into retirement (at the trough of prices paid and one year after the trough for bulk savings), households pay around 0.2 percent lower prices for a fixed UPC than they did just prior to retirement while they save an additional 0.2 percent on per-unit prices by changing their bulk allocation. The quality of the consumption bundle, on the other hand, changes significantly. As we can see in Panel (c), households rapidly substitute towards cheaper, lower-quality items within narrowly defined categories of consumption. Four years into retirement, quality of the purchased bundle relative to the average bundle has fallen by about 1.2 percentage points.

These results have stark implications. In Panel (a) of Figure 2.5, I decompose the change in the household's deviation from the cost of the "average" bundle (that is, a bundle consisting of the same consumption categories, but at average prices, average bulk shares, and average quality). We can see that bundle quality, rather than price and bulk effects, is the main margin of shopping adjustment. Quality adjustments are three to four times as large as the combined adjustment from paying lower prices for a given good or increasing bulk purchases. This suggests that the main shopping adjustment is one that is costly in utility terms rather than one that yields similar levels of utility at lower expenditures through lower prices. In Panel (b), I present the path of expenditures, the path of *price and bulk-corrected* expenditures, and the path of quantities only (so declines in expenditure that arise from purchasing fewer items or substitutions across product modules). Consumption expenditures fall by almost as much as uncorrected expenditures. There also is very evidence of households "learning" how to save money: the wedge between corrected and uncorrected

[14] Since the HRS is a bi-annual survey, this means that we can only group the leads and lags into two year bins. Therefore period -4 refers to four to three years prior to retirement, $t = -2$ refers to two to one years prior to retirement, $t = 0$ refers to zero to one years since retirement, and so on.

expenditures *falls* later into retirement and has fully disappeared six years into retirement. Even accounting for downward adjustments in item quality, quantities fall by almost 7%. Taken together, these two pieces of evidence imply that the post-retirement expenditure drop is a real drop in consumption expenditures, with utility-constant shopping effects explaining *at most* 5 to 10 percent of the expenditure drop.

Home Production and Expenditure Timing

An important caveat is that the estimates in Figure 2.5 still only capture *market* consumption. To the extent that households cut durable goods and work-related expenses and adjust their consumption bundles towards more home production, *expenditures* will overstate the consumption drop. In order to test these channels, I perform four separate tests: First, I make use of the detailed nature of the Scanner data and group goods into three categories: Goods that cannot be substituted with home production (e.g., shampoo, trash bags), goods that are substitutes for consumption away from home but are not inputs for home production (e.g., ready-to-eat foods), and home production inputs (e.g., unprepared produce and meat, flour, fresh eggs). I then estimate the expenditure drop for each category separately. As we can see in Panel (a), spending on non-substitutable goods falls by more than spending on home-consumption goods and spending on home production inputs do. However, spending on home consumption and home production inputs also falls. Given that expenditure on these categories also falls, at-home consumption and home production cannot offset declines in consumption elsewhere. In Panel (b), I conduct a similar exercise and leverage the fine-grained information in the CAMS module of the HRS to construct expenditure variables for home production goods (food at home, cleaning products, gardening products) and their immediate market good substitutes (food away from home, cleaning services, gardening services). Importantly, CAMS also includes information on time spent in each of these activities (time spent cleaning, doing gardening work, preparing meals and cleaning up and time spent shopping or running errands). In Panel (b) of Figure 2.6, we can see that expenditures on home production inputs fall almost as much as expenditures on the equivalent market goods. Therefore, it seems unlikely that home production could offset the decline in spending on market production. As a complement to this, we can consider households' time use: As we can see in Panel (c), the total time engaged in market work per week decreases by about 22 hours at retirement. However, the time spent preparing meals, shopping, cleaning, and gardening increases by only 4.5 hours for the first two years after retirement and then slowly declines thereafter.

Last, I investigate to what extent the large decline in total expenditures in the HRS could be the result of expenditure timing (for example, households purchasing new durables right before retirement, thereby elevating pre-retirement expenditures). To test this, I make use of the RAND HRS CAMS Data File that constructs a measure of *consumption* that accounts for principal repayment of mortgages and the fact that the consumption of durables occurs over time rather than all at once.[15] As we can see in Panel (d) of Figure 2.6, consumption spending does fall by less than total spending, but the declines in consumption expenditure are still substantial and there is no evidence of a spike in expenditures *before* retirement. The declines in total spending and are substantially larger than declines in non-durable spending alone (with consumption spending falling in between the two).

Mechanisms and Heterogeneity

An important question is how much the adjustments at retirement vary across households. Hurst (2008) argues that many studies of the post-retirement expenditure drop actually find *zero* adjustments at the median so that most of the drop is driven by a relatively small set of households who might be myopic and did not sufficiently plan for retirement. On the other hand, Bernheim et al. (2001) find that even households with relatively low declines in post-retirement income or high wealth cut their consumption after retirement. In their setting, households with low assets and larger post-retirement income drops do experience the largest declines in expenditures, but the expenditure drops are ubiquitous across the income and wealth distributions (the only group not seeing an expenditure decline is the set of high-post-retirement-income high-asset households).

To tackle the question of the heterogeneity of the consumption drop, I perform several tests in the spirit of Bernheim et al. (2001). First, I group households into terciles according to their post-retirement income drop. In order to not pick up any one-time fluctuations, I define the income drop as the log difference in the average income over the three years preceding and immediately succeeding retirement. I then estimate the expenditure drop, the drop in consumption expenditure, as well as the expenditure wedge and quality adjustments for each tercile of the income drop. Similar to the finding of Bernheim et al. (2001), expenditures fall for every tercile of the income drop. Nonetheless, the expenditure drop for households with the smallest income drop is only about half as large as that for the two other terciles, suggesting an important role for current income. The wedge between expenditure and consumption relevant expenditure falls significantly only for households with the largest income drop, but even these declines are very small relative to the overall decline in expenditure. The quality of the bundle, too, declines the most for households with the largest income drop, with quality adjustments among the first tercile about three to four times the size of the adjustment for the top tercile.

As a second test, I match the Nielsen Consumer Panel to the Survey of Consumer Finances to impute household wealth (see Section 2.9 for details) and group households by their pre-retirement imputed wealth. Here, the findings are also resembling those of Bernheim et al. (2001): Expenditures fall for each tercile of imputed wealth, but they fall the most for households with the lowest pre-retirement wealth. Savings arising from paying lower purchases and more bulk purchases are largest for the middle tercile with very little

[15] In particular, the RAND HRS CAMS Data File uses information from the Consumer Expenditure Survey to separate out interest expenses and principal payments on mortgages. For durables, it applies a per-period flow usage transformation following the approach of Hurd and Rohwedder (2006).

adjustment for either the top or the bottom tercile of pre-retirement wealth. Bundle quality, on the other hand, falls the least for the top tercile of wealth, with the drops for the bottom two terciles looking quite similar.

In results not reported here, I further investigate heterogeneity by educational attainment and pre-retirement income. Results are similar in the sense that it is households with lower incomes and lower education attainment who see the largest declines in expenditures, corrected expenditures, and bundle quality. This suggest that price and bulk effects at retirement are not only small on average, they are also small across many socioeconomic observables. On the other hand, households with lower current income, households with lower permanent income, and households with lower levels of wealth are all more affected by falling expenditures at retirement. These results are also borne out in the CAMS data which has joint information on assets, income, and expenditures. Total expenditure falls by much more for households with low wealth-to-income ratios or large post-retirement income drops. In terms of non-durables, point estimates are generally negative, but in most cases, I cannot reject no declines in expenditures for households with high wealth or a small post-retirement income drop (see Appendix Figure 2.9). From a social welfare perspective, these groups are likely to have relatively high marginal utilities, suggesting that the post-retirement expenditure drop also has important consequences for aggregate welfare.

Together, these results suggest that households really are differentially insured against the prospect of retirement. Some households see much smaller declines in current income that also go along with much smaller drops in expenditure. Wealthier households are also much better insured against large declines in current income after retirement, suggesting that differences in wealth accumulation are an important factor in explaining the heterogeneity in the magnitude of the post-retirement expenditure drop. These results are similar to the findings of Ganong and Noel (2019) and Ganong et al. (2020a) in the context of unemployment and general fluctuations in labor earnings. In their setting, households with low liquidity are much more sensitive to unemployment or other labor-demand driven fluctuations in earnings. Considering more extreme cases, Ganong and Noel (2022) show that 70% of mortgage defaults are driven by adverse life events (shocks to current and future income) rather than negative equity or the interaction between negative equity and adverse life events. My results indicate that liquidity is an important driver of consumption behavior not just in light of these unanticipated shocks but even in light of anticipated shocks to current income. This points towards an important role for present bias or mental-accounting consumption behavior that puts weight not just on permanent income but *current* income as well.

Robustness

One potential concern with my estimates is that my main measures of shopping adjustment are all defined at the national, rather than local level. To investigate to what extent this has an effect on my results, I re-estimate the event studies for price effects, bulk effects, and quality adjustments making only local comparisons when constructing my shopping measures. For this, I leverage the fact that Nielsen divides its panel in to markets. These are 76 areas, which sometimes align with metro areas but need not be contiguous (e.g., rural counties surrounding a metro area might be defined as one market while the urban core could be defined as another). As we can see in Figure 2.A9, the price and bulk effects are *smaller* when using local comparisons while the quality effects are very similar to the effect estimated using national comparisons.

I further explore whether my event study results are robust to my choice of estimator. In order to do so, I re-estimate my event studies using the estimator proposed by de Chaisemartin and D'Haultfœuille (2020). As we can see in Figures 2.A11 and 2.A12, the income and expenditure drops are a bit smaller. However, expenditures still discretely drop after retirement. The estimated shopping effects are very similar to those of my baseline specification with price effects being a bit larger and bulk effects being a bit smaller (see Figure 2.A13). Overall bundle savings are almost identical to my baseline specification as are quality adjustments (see 2.A14). Taken together, shopping effects explain at most 10% of the post-retirement expenditure drop so the result that the post-retirement expenditure drop is not driven by shopping effects is robust to the choice of estimator.

2.6 When Do Shopping Effects Matter?

Given the results of the previous section, it is worth revisiting the question when shopping effects matter for household level prices and the quality of the consumption bundle. In order to do so, I report the effect of household head labor market attachment in the spirit of Kaplan and Menzio (2015). An important deviation from their estimates is that I break out labor market attachment for male and female household heads separately. In Panel A of Table 2.6, I report estimates for the set of households with two household heads and an "average" household age between 25 and 54. In Column (1), we can see that households with a *female* head not working full time are paying around 0.9% lower prices while the male heads' labor market participation has no effect on household-level prices. Controlling for household income, the effect for female household heads gets a bit smaller, although households with female household heads working part time or not working at all are still paying around 0.75% lower prices. Including a household fixed effect, the coefficient on the female heads' labor market attachment gets cut in half for non-employment and falls by a factor of four for female heads working part time.

With respect to the quality of the consumption bundle, we can see that much of the variation explained by household heads' labor market attachment is the result of the effect on earnings. While households with fewer employed household heads consume much cheaper bundles, conditional on income, this effect disappears or even reverses. Including a household fixed effect, the effect of labor market attachment on the quality of the consumption bundle is still statistically significant, but of much smaller economic magnitude than for the specification without household fixed effects or income controls.

In Panel B, we can see that a similar story applies for household ages 55 to 74. Unlike for prime-age households, the male head's labor market attachment now is predictive of household-level prices. Households with a female head not working full time pay between 0.7% and 1% lower prices, while households with a male head not working full time pay between 0.4% to 0.5% lower prices. Controlling for household income, all coefficients shrink a bit, but most are still statistically significant. Including a household fixed effect, male heads' labor market participation is no longer predictive of household level prices. Households with female heads not working full time pay around 0.2% lower prices, a similar magnitude to the maximum effect of retirement on household level prices estimated in Section 2.5. For bundle quality, patterns are quite similar: Households with heads not working full time are purchasing significantly cheaper items, an effect almost entirely driven by the effect of work on household income. Including a household fixed effect, not working full time reduces the quality of the consumption bundle by 0.6 to 0.75% with no effect of part time work for male household heads.

Together, these results suggest that the division of labor and household-level labor supply choices play an important role in explaining household level prices. If household-level prices were only a measure of household's opportunity cost of time, we would expect male household heads' employment status to also be an important determinant of household-level prices. However, most coefficients on male heads' employment status are not significantly different from zero and reasonably precisely estimated. More generally, a lot of the dispersion of household level prices can be explained by household fixed effects rather than within-household across time variation.

The Case of Unemployment

Given the apparently large role of fixed household characteristics in terms of explaining household-level prices, it is instructive to investigate the adjustment around a shock different from retirement: unemployment. In some ways retirement and unemployment are quite similar: Both retired and unemployed households have more free time to engage in money-saving shopping activities or home production.[16]

However, there are also very important differences for these two shocks. Unemployment represents a (mostly) unanticipated shock to income. Retirement, on the other hand, is (mostly) anticipated. Therefore, ex-ante it seems likely that there are more pre-cautionary shopping effects for households entering retirement rather than households entering an unemployment spell. In addition, unemployment spells decrease lifetime income (a result going back to Jacobson et al. 1993), while retirement should leave lifetime income unchanged.

[16] An important note is that households facing unemployment will also spend some of their time looking for a job. Work by Krueger and Mueller (2010) suggests that the time spent looking for a job would still leave ample time to exert shopping effort to lower household-level prices. Households facing adverse labor market shocks also have other adjustment margins. For example, Koustas (2018) finds that ride-share drivers a substantial fractions of lost earnings in primary jobs with ride-share earnings, suggesting that flexible labor supply adjustments may be a very important aspect of households' self-insurance behavior against adverse labor market shocks.

Therefore, we would expect larger adjustments for unemployment than retirement. A factor pushing in the opposite direction is that unemployment is (usually) temporary while retirement is an absorbing state. To the extent that households consumption and shopping behavior are driven by *current* income, this would result in larger adjustments for households entering retirement.

In Figure 2.10, I present event studies for households undergoing non-employment spells. Since demographics are only elicited once a year, I cannot directly observe when households become unemployed. Rather, I center all coefficients at the last quarter before the unemployment spell so that $t = -1$ corresponds to the last known quarter of employment and $t = 4$ corresponds to the quarter in which a household member was non-employed. To make sure that most of the non-employment I am picking up is unemployment (rather than people voluntarily withdrawing from the labor force), I further restrict attention to unemployment spells that end after no more than two years. Estimating the effects on a balanced window, we see that during these non-employment spells, the brunt of the shopping adjustment falls on quality changes. Between the last quarter of known employment and the first quarter of known non-employment, bundle quality falls by about 1.1%. Prices paid fall by around 0.3% over the same time horizon while savings from bulk allocations are unresponsive to the non-employment spell. A notable feature of the adjustment is that effects on household-level prices and bundle quality revert back over time, though only prices recover fully. One potential explanation for the persistence of the effect on bundle quality is that unemployment tends to go along with declines in life-time income.

Implications for Using Expenditure as a Proxy of Consumption

Together, my results suggest that an important determinant of household-level shopping behavior is a household or type fixed effect. Households with female heads who are not working full time pay substantially lower prices, but this appears to be largely a result of differences in the household production function between different households. Within households, household heads' employment status does not affect household-level prices much. In the case of male heads of household, their employment status rarely matters for household-level prices irrespective of the choice of estimation procedure. This either implies that men and women have substantially different opportunity costs of time or that shopping is part of a larger intra-household bargaining problem for which opportunity costs of time are just one of many considerations. Another implication of these patterns it that shopping effects matter a lot for cross-sectional comparisons, but much less for within-household-across-time comparisons. That means that analyses relying on expenditures as a proxy for consumption are likely to be good approximations of true consumption adjustments as long as the underlying variation is within household.

2.7 Conclusion

Is the post-retirement expenditure drop a true drop in consumption? Decomposing expenditure into shopping behavior that keeps the consumption basket constant and quality and quantity adjustments that are costly in utility terms, I find that at least 90% of the expenditure decline after retirement represent true declines in quantities or quality adjustments. I further investigate whether substitution towards home production or expenditure timing can explain the expenditure drop. Expenditures declines are ubiquitous across all sub-components of consumption and even expenditures on home production inputs fall in retirement.

These patterns have important policy implications: The high sensitivity of households' consumption to current income implies that the payout scheme is a crucial policy lever when designing pension systems (irrespective of their funding mechanism). The main source of a constant stream of income for households in retirement is Social Security which only replaces 40% of pre-retirement earnings on average. One potential remedy for the large post-retirement expenditure drop is to increase replacement rates of Social Security, particularly at the lower end. More generally, my findings suggests that retirement schemes that guarantee constant cash flows may have substantial welfare benefits relative to systems without recurring payouts. This question of optimal payout schemes is particularly important as future retirees will be ever more likely to be drawing from defined-benefit plans. Given the failure of many households to annuitize their wealth, policies to increase annuity take-up or other ways to derive stable income flows from private retirement accounts are likely to have substantial welfare benefits. This is particularly important as the US transitions from a private retirement system mostly composed of defined benefit plans with guaranteed income flows to one dominated by defined contribution plans without any pre-set withdrawal strategies.

My results suggest multiple avenues for future research. On the empirical side, it would be interesting to use richer information on assets and income to explore in more detail which households are more insulated from the post-retirement expenditure drop. This work could then explore heterogeneity in the mechanism underlying the expenditure drop. For some households, it may be the result of low-wealth and a large decline in income, forcing the large adjustment. For wealthier households, important avenue for future research will be to disentangle to what extent the low rates of dis-saving are driven by large bequest motives, a failure to annuitize wealth holdings, or high sensitivity to current income even among the wealthy. On the theoretical side, the most common explanation for a high sensitivity of consumption to current income are high returns on illiquid assets or a combination of liquidity constraints and present focus. However, both classes of models are hard to reconcile with drops in expenditures in retirement even for relatively wealthy households. Models of consumption that can rationalize these behaviors would be an important contribution to our understanding of households' consumption behavior.

2.8 Figures and Tables

Figure 2.1: Validating the Quality Measure

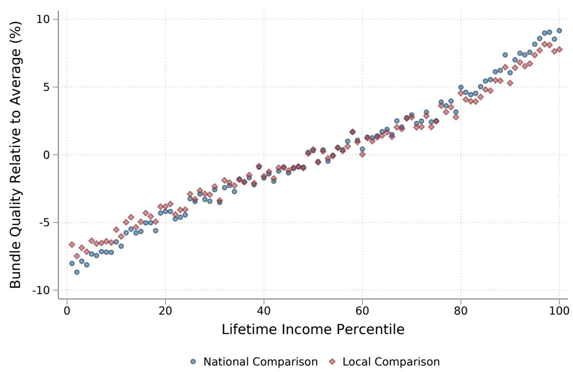

(a) Bundle Quality and Lifetime Income

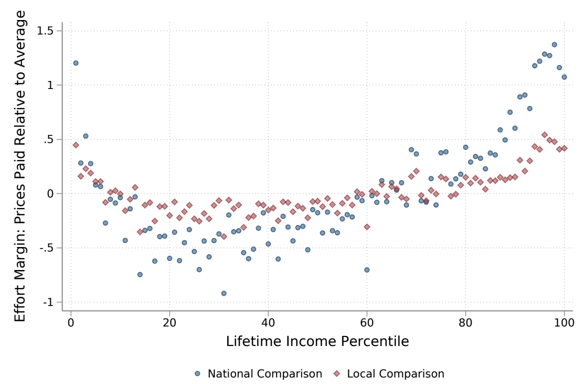

(b) Prices Paid and Lifetime Income

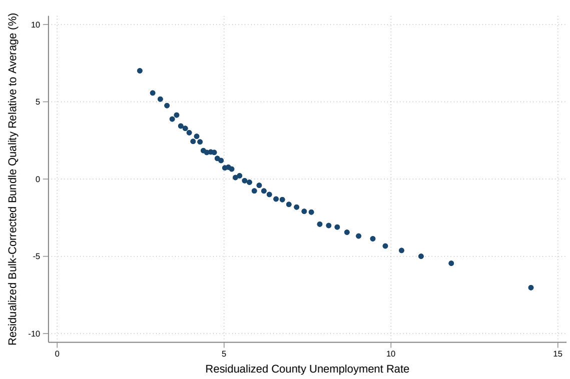

(c) Bundle Quality and Local Unemployment

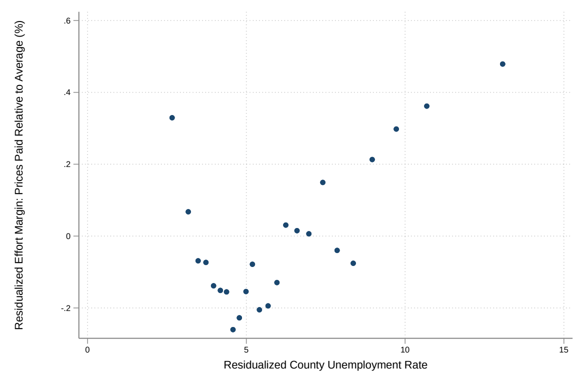

(d) Prices Paid and Local Unemployment

Notes: This figure presents two validation exercises for disentangling price, bulk, and quality effects. The effort and quality margins are as defined in Equations 2.7 and 2.10. The underlying data for all figures is the baseline sample of Table 2.1.

Figure 2.2: The Path of Income Around Retirement

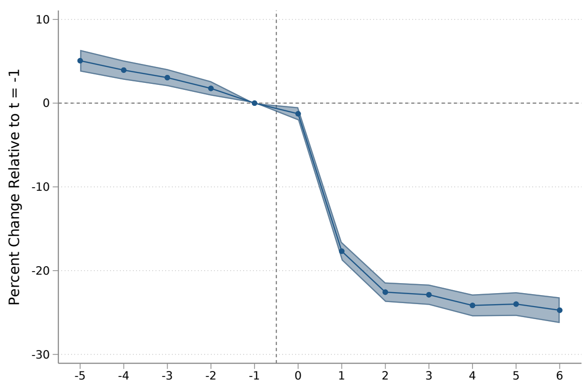

(a) Scanner Data

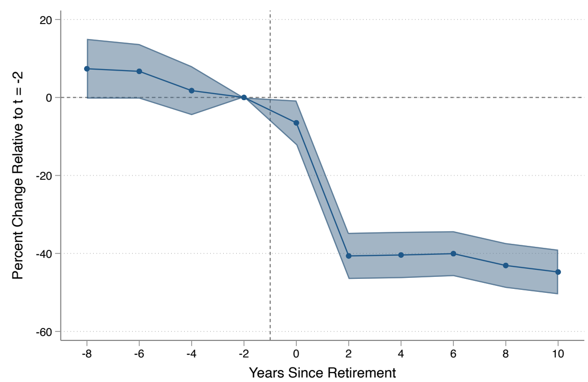

(b) Consumption Survey

Notes: This figure presents the path of income around retirement in the Kilts Nielsen Consumer Panel and the Health and Retirement Survey estimated according to Equation 2.11. The underlying data for Panel (a) is the retirement sample of Table 2.1, for Panel (b) is the retirement sample of Table 2.2. Standard errors are clustered at the household level.

Figure 2.3: The Path of Expenditure Around Retirement

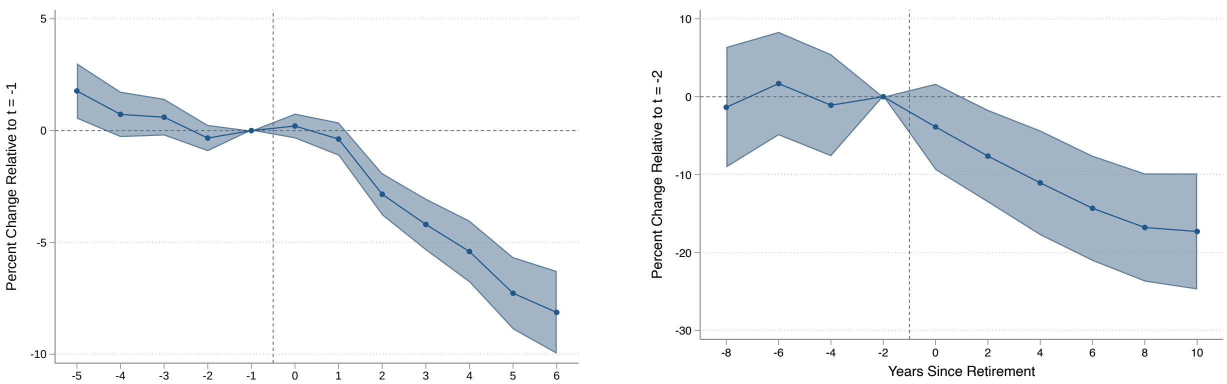

(a) Scanner Data (b) Consumption Survey

Notes: This figure presents the path of expenditure around retirement in the Kilts Nielsen Consumer Panel and in the Health and Retirement Survey estimated according to Equation 2.11. Expenditure in the KNCP is defined as spending on non-magnet data product modules covered in all panel years. Expenditure in the HRS is total reported expenditure on non-durables (excluding housing). The underlying data for Panel (a) is the retirement sample of Table 2.1, for Panel (b) is the retirement sample of Table 2.2. Standard errors are clustered at the household level.

Figure 2.4: Shopping Adjustments Around Retirement

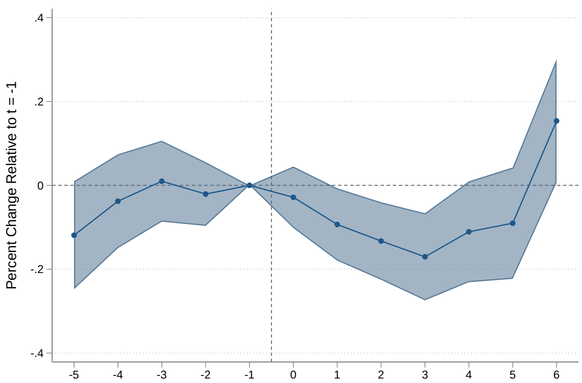

(a) Prices Paid

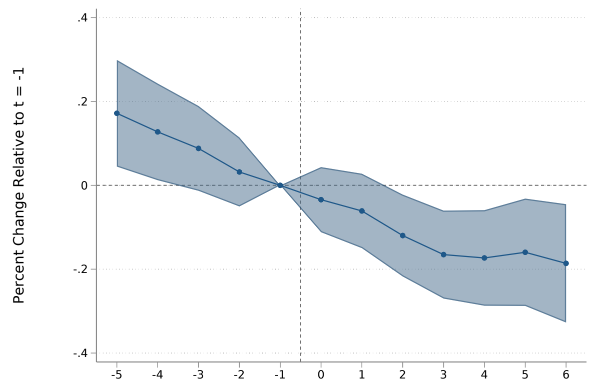

(b) Cost Relative to Average Bulk Allocation

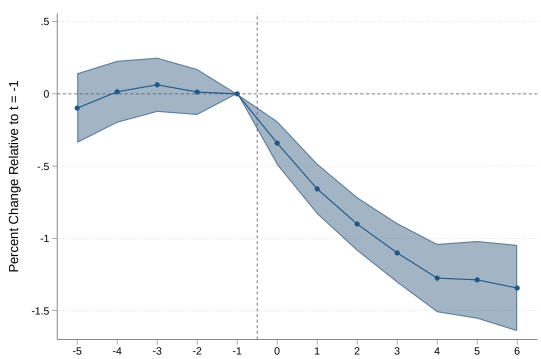

(c) Bundle Quality

Notes: This figure presents the evolution of price effects, bulk effects, and quality adjustments estimated according to Equation 2.11. The underlying data for each panel is the retirement sample of Table 2.1. Standard errors are clustered at the household level.

Figure 2.5: Decomposing the Post-Retirement Expenditure Drop

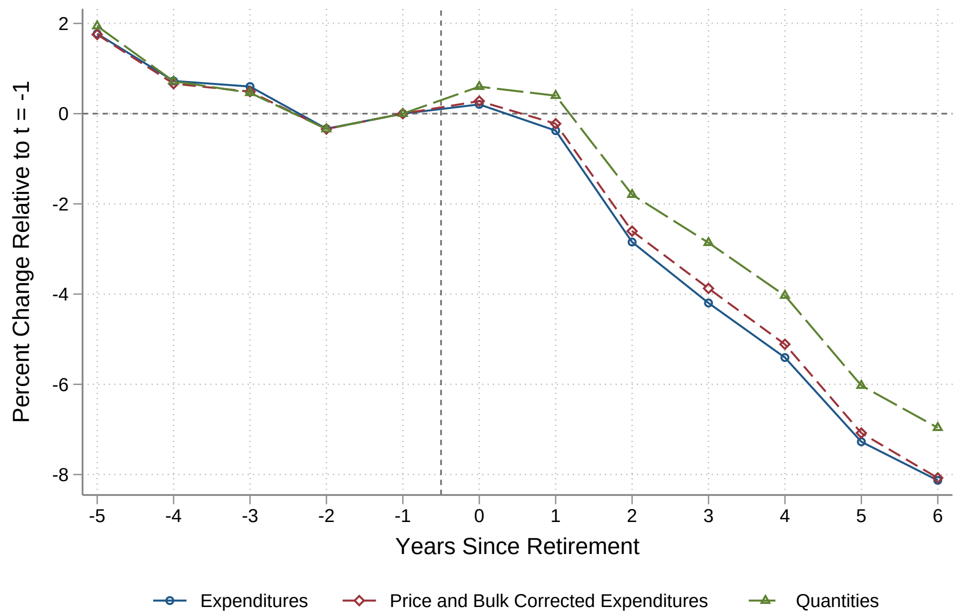

(a) Expenditures and Corrected Expenditures

Notes: This figure presents the evolution of price effects, bulk effects, and quality adjustments estimated according to Equation 2.11 as well as uncorrected expenditure, expenditure at average prices and the average bulk allocation, and expenditure at average prices, the average bulk allocation, and average within-product module quality. The underlying data for each panel is the retirement sample of Table 2.1.

Figure 2.6: Home Production, Home Consumption and Expenditure Timing in Retirement

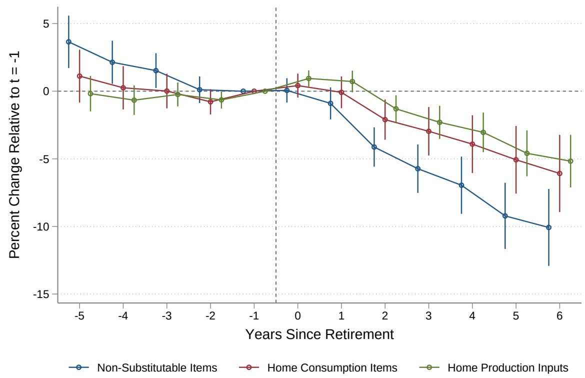

(a) Subsets of Spending by Substitutability

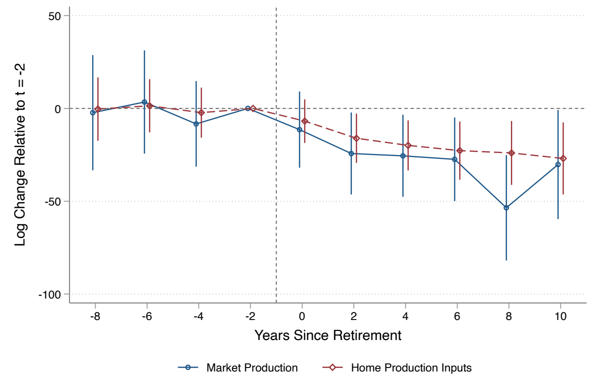

(b) Market and Home Production Expenditure

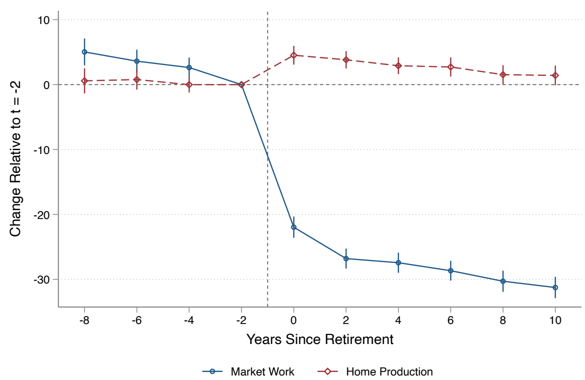

(c) Time Working vs. in Home Production

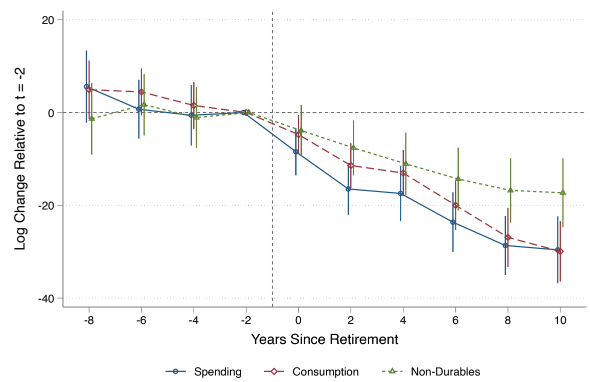

(d) Total vs. Consumption Expenditure

Notes: This figure presents the path of expenditure categories around retirement in the Kilts Nielsen Consumer Panel as well as expenditure categories and time use in the Health and Retirement Survey estimated according to Equation 2.11. Non-substitutable items are non-food grocery items, health and beauty items, and general merchandise. Home consumption items are prepared foods and deli items, food items that do not require preparation, and alcohol. Home production inputs are dry grocery items, packaged meat, fresh produce, frozen meat and vegetables, and dairy products. Home production items in the HRS are food at home, cleaning products, gardening products expenditures. Market production items are their immediate market good substitutes (food away from home, cleaning services, and gardening services). Home production time is time spent cleaning, doing gardening work, preparing meals and cleaning up and time spent shopping or running errands. The underlying data for Panel (a) is the retirement sample of Table 2.1; the data for Panels (b)-(d) is the retirement sample of Table 2.2. Standard errors are clustered at the household level.

Figure 2.7: Heterogeneity by Size of Income Drop

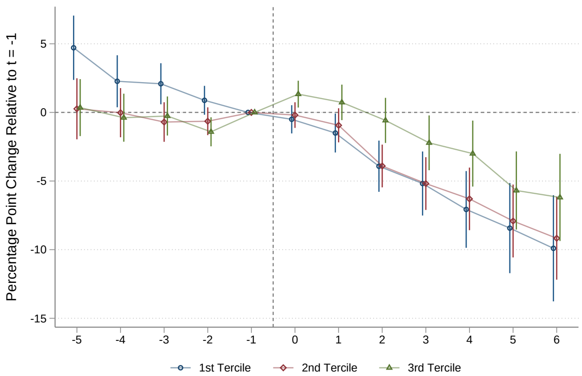

(a) Expenditures

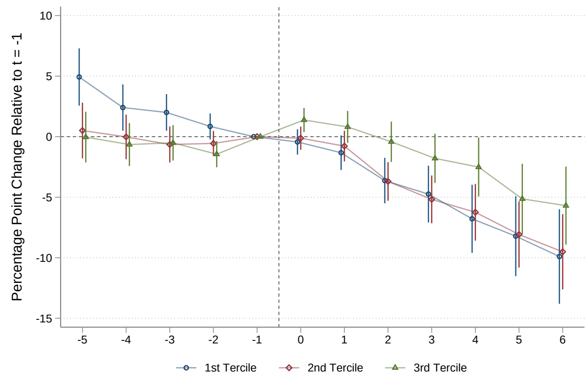

(b) Price & Bulk Corrected Expenditures

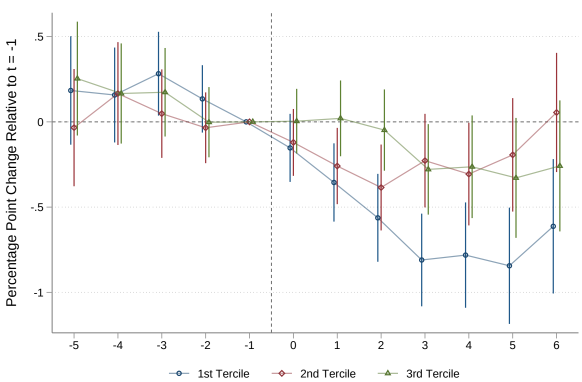

(c) Price and Bulk Effects

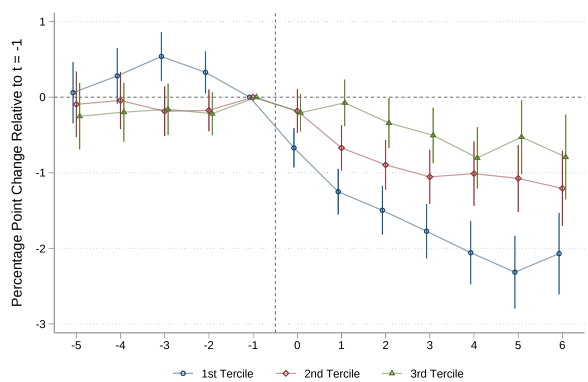

(d) Quality Adjustments

Notes: This figure presents the evolution of expenditures, price and bulk corrected expenditures, price and bulk effects, and quality adjustments for each tercile of the post-retirement income drop. The income drop is defined as the log difference in the average income over the three years preceding and immediately succeeding retirement. The underlying data for each panel is the retirement sample of Table 2.1. Standard errors are clustered at the household level.

Figure 2.8: Heterogeneity by Pre-Retirement Wealth

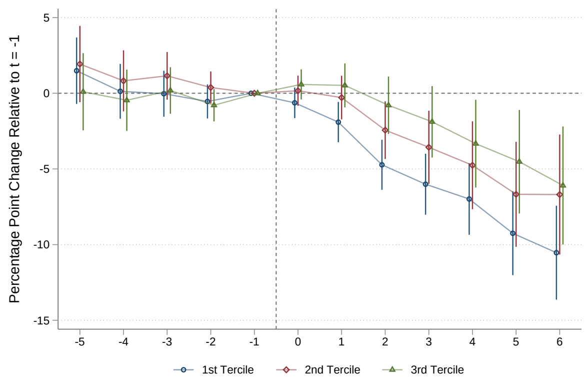

(a) Expenditures

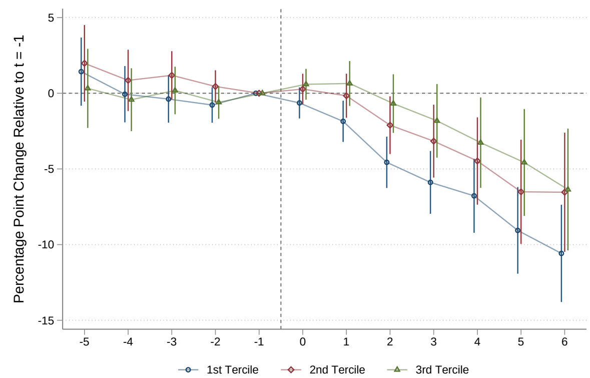

(b) Price & Bulk Corrected Expenditures

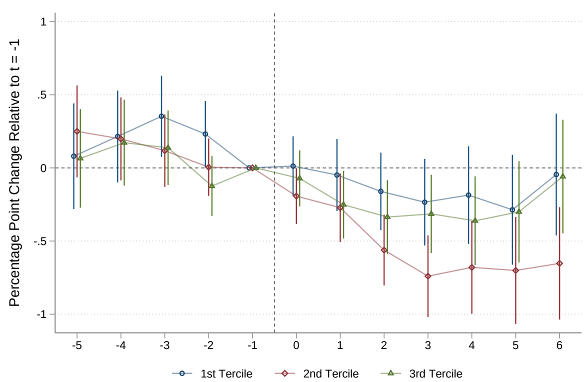

(c) Price and Bulk Effects

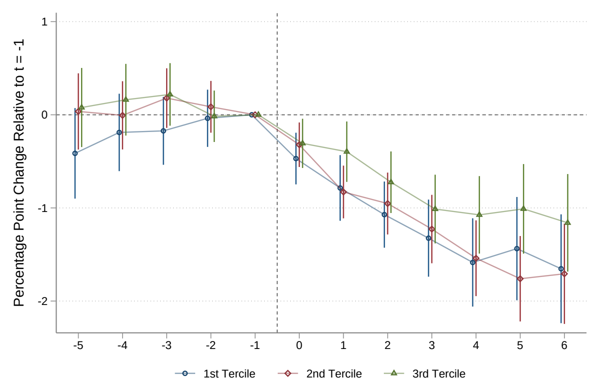

(d) Quality Adjustments

Notes: This figure presents the evolution of expenditures, price and bulk corrected expenditures, price and bulk effects, and quality adjustments for each tercile of pre-retirement household wealth, where household wealth is a cell based match to the Survey of Consumer Finances. The underlying data for each panel is the retirement sample of Table 2.1. Standard errors are clustered at the household level.

Figure 2.9: Heterogeneity of Broader Expenditures by Income and Wealth

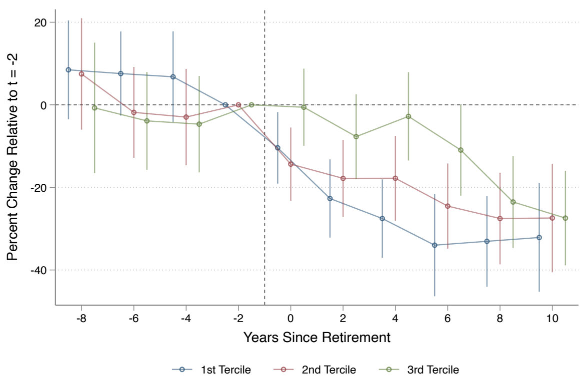

(a) By Income Drop: Expenditure

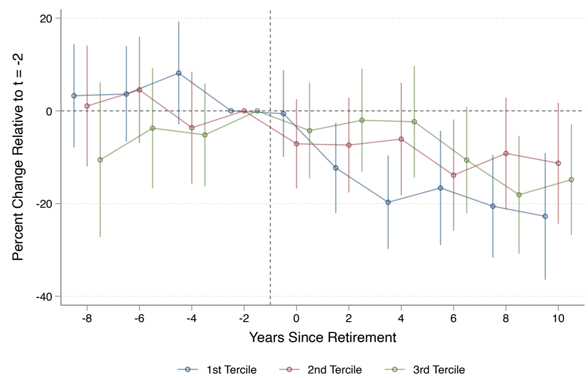

(b) By Income Drop: Non-Durables

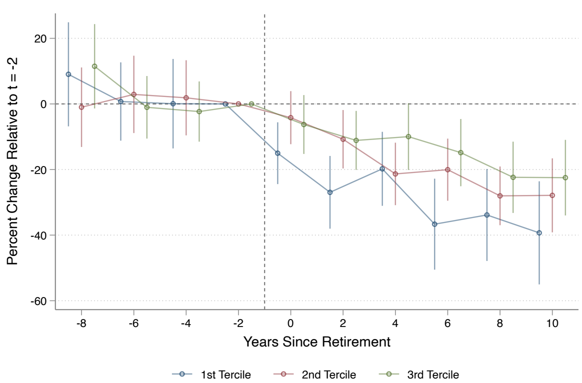

(c) By Wealth-to-Income: Expenditure

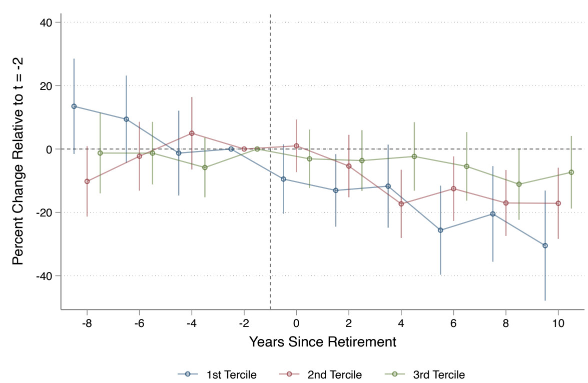

(d) By Wealth-to-Income: Non-Durables

Notes: This figure presents the evolution of total expenditure and non-durable expenditure by the size of the post-retirement income drop and the pre-retirement wealth-to-income ratios. The underlying data for each panel is the retirement sample of Table 2.2.

Figure 2.10: Shopping Adjustments Around Nonemployment

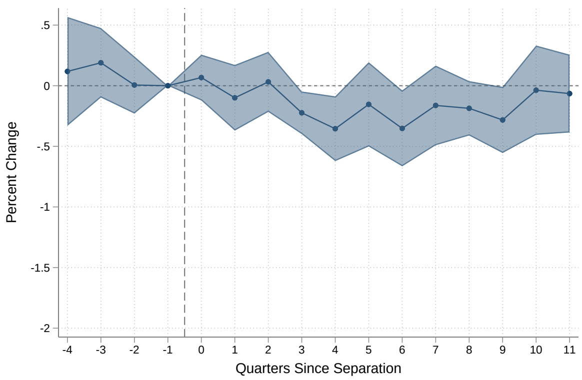

(a) Prices Paid

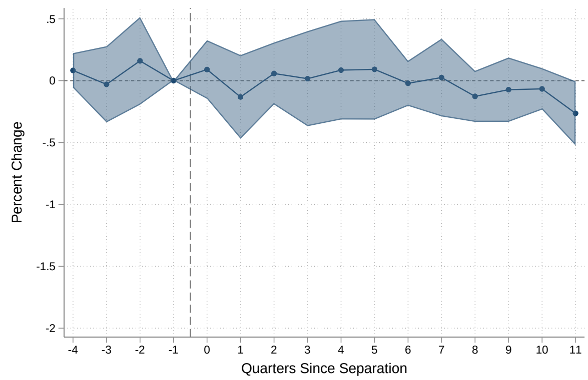

(b) Cost Relative to Average Bulk Allocation

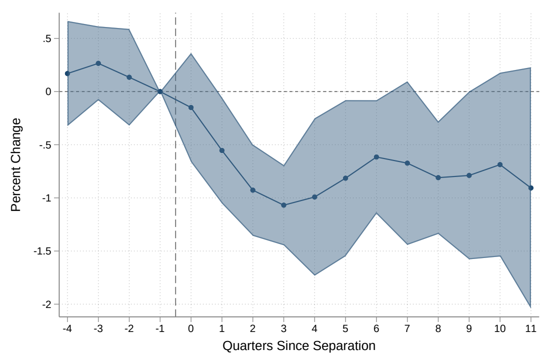

(c) Bundle Quality

Notes: This figure presents the evolution of price effects, bulk effects, and quality adjustments estimated according to Equation 2.11 for households undergoing a non-employment spell lasting at most two years. Households heads are reporting to be employed at $t = -1$, report non-employment at $t = 3$, and are reporting being employed again in either $t = 7$ or $t = 11$. Standard errors are clustered at the household level.

Table 2.1: Kilts Nielsen Consumer Panel Summary Statistics

	(1) Baseline Sample			(2) Retirement Sample		
	Mean	Median	SD	Mean	Median	SD
Male Employment	0.73	1	0.45	0.46	0	0.50
Male Full-Time	0.68	1	0.47	0.37	0	0.48
Female Employment	0.60	1	0.49	0.47	0	0.50
Female Full-Time	0.48	0	0.50	0.33	0	0.47
Two Heads	0.51	1	0.50	0.72	1	0.45
No Male Head	0.28	0	0.45	0.21	0	0.41
No Female Head	0.20	0	0.40	0.074	0	0.26
Married	0.48	0	0.50	0.69	1	0.46
Household Size	2.61	2	1.46	2.03	2	0.87
Any Children Present	0.34	0	0.47	0.047	0	0.21
Male Age	49.8	50	12.7	64.6	65	6.48
Female Age	48.4	49	12.9	62.6	63	6.40
Non-Hispanic White	0.70	1	0.46	0.84	1	0.37
Black	0.12	0	0.33	0.087	0	0.28
Hispanic	0.12	0	0.33	0.037	0	0.19
Male Education	13.9	14	2.36	14.3	14	2.37
Female Education	13.9	14	2.17	14.3	14	2.14
Total Expenditure	4766.4	4106.6	3003.4	5283.4	4712.5	3035.5
Binned Household Income	73555.8	57543.8	55324.2	68912.3	57543.8	47256.8
Years in Panel	7.62	7	4.73	10.4	11	3.71
Any Retirement	0.090	0	0.29	1	1	0
Total Retirements	0.095	0	0.32	1.02	1	0.38
Any Unemployment Spell	0.097	0	0.30	0.029	0	0.17
Total Unemployment Spells	0.11	0	0.36	0.030	0	0.19
Observations	814938			107877		

Notes: This table presents summary statistics for the three main analysis samples. Male always refers to the male household head, female refers to the female household head. Education is expressed in years. Household income is deflated by the PCE, total expenditures are deflated by the CPI Food at Home for Urban Consumer; both are expressed in 2012 dollars. Binned household income in the Nielsen Consumer Panel is translated into dollars based on the mean income in the same income bin the IRS Statistics of Income database. For the baseline sample, all values are weighted using projection factors to make sample representative of US population. For the retirement sample, all values are weighted by the inverse of years in the sample, so that each household has the same weight. Total observations correspond to household years.

Table 2.2: Summary Statistics Health and Retirement Survey Data

	Mean	Median	SD
Panel A: Not Retired			
Household Size	2.5668	2	1.3633
Married	0.6293	1	0.4830
Age	58.201	57	8.1996
White	0.7047	1	0.4562
Black	0.1759	0	0.3808
Hispanic	0.1572	0	0.3640
Years of Education	13.092	13	3.2349
Log Household Income (2012 USD)	10.815	11.098	1.8031
Log Total Expenditure (2012 USD)	10.597	10.621	0.7381
Log Nondurable Expenditure (2012 USD)	9.9373	9.9635	0.7508
Weekly Hours Work for Pay	30.687	40	20.492
Weekly Hours Homeproduction	25.485	20	23.740
N	19067		
Panel B: Retired			
Household Size	2.0645	2	1.0727
Married	0.5460	1	0.4979
Age	71.207	71	9.7238
White	0.7873	1	0.4092
Black	0.1562	0	0.3631
Hispanic	0.08608	0	0.2805
Years of Education	12.564	12	3.0480
Log Household Income (2012 USD)	10.367	10.443	1.3565
Log Total Expenditure (2012 USD)	10.337	10.354	0.7762
Log Nondurable Expenditure (2012 USD)	9.7941	9.8257	0.8073
Weekly Hours Work for Pay	3.5184	0	10.181
Weekly Hours Home Production	26.457	21	24.009
N	42892		
Panel C: Full Sample			
Household Size	2.2191	2	1.1926
Married	0.5717	1	0.4948
Age	67.204	66	11.053
White	0.7618	1	0.4260
Black	0.1623	0	0.3688
Hispanic	0.1079	0	0.3103
Years of Education	12.726	12	3.1161
Log Household Income (2012 USD)	10.505	10.626	1.5222
Log Total Expenditure (2012 USD)	10.411	10.430	0.7744
Log Nondurable Expenditure (2012 USD)	9.8350	9.8646	0.7942
Weekly Hours Work for Pay	11.655	0	18.794
Weekly Hours Home Production	26.166	20	23.933
N	62664		

Notes: This table presents summary statistics for the HRS CAMS sample, broken out by retired and not-yet-retired households.

Table 2.3: Distribution of the Effort and Quality Distributions

Mean	Price Effects	Bulk Effects	Price + Bulk Effects	Quality Adjustments
SD	7.97	5.76	9.93	11.93
p5	-13.14	-9.51	-15.17	-19.96
p10	-8.66	-7.09	-11.56	-15.01
p25	-3.45	-3.44	-6.03	-7.36
p50	0.38	0.1	-0.28	0.35
p75	4.15	3.47	5.79	7.62
p90	8.49	6.73	11.95	14.27
p95	11.48	8.97	16.12	18.63

Notes: This table presents moments for my measure of price effects, bulk effects, combined price and bulk effects, as well as the quality adjustment metric. These statistics correspond to the baseline sample of the data presented in Table 2.1.

Table 2.4: Validating the Quality Measures

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: National Quality						
Log Household Income	0.0643	0.0231	0.0523	0.0609	0.0564	0.0192
	(0.0005)	(0.0007)	(0.0006)	(0.0006)	(0.0006)	(0.0009)
Lifetime Income Rank		0.1375				0.1010
		(0.0021)				(0.0026)
Log Home Price			0.0463			0.0404
			(0.0007)			(0.0008)
Imputed Wealth Ventile				0.1657		0.0245
				(0.0060)		(0.0061)
Highschool					1.3839	1.9243
					(0.3287)	(0.4070)
Some College					3.0698	3.0269
					(0.3272)	(0.4060)
College Grad					4.6408	3.8905
					(0.3292)	(0.4087)
Postgrad					6.1632	4.7137
					(0.3394)	(0.4205)
Panel B: Local Quality						
Log Household Income	0.0572	0.0220	0.0537	0.0544	0.0502	0.0188
	(0.0006)	(0.0007)	(0.0007)	(0.0006)	(0.0006)	(0.0009)
Lifetime Income Rank		0.1177				0.1043
		(0.0022)				(0.0028)
Log Home Price			0.0133			0.0069
			(0.0008)			(0.0009)
Imputed Wealth Ventile				0.1507		0.0402
				(0.0062)		(0.0067)
Highschool					1.1226	1.7712
					(0.3293)	(0.4123)
Some College					2.5226	2.8361
					(0.3274)	(0.4109)
College Grad					3.8805	3.7851
					(0.3299)	(0.4147)
Postgrad					5.3974	4.8033
					(0.3414)	(0.4291)
Observations	790983	770597	615243	598557	790983	453658

Notes: This table reports regressions of the national and local quality measure on six separate measures of socioeconomic status: Current income (in 2012 USD), lifetime income percentile, log real house prices from Zillow (for households living in single-family homes or condos/coops), imputed household wealth from the Survey of Consumer Finances, and indicators for educational attainment with no high school degree as the reference category. All regressions control for household size, household composition (two household heads, married, otherwise related household heads, a single male or female head), and the age and presence of children. The estimation sample for each regression is the baseline sample from Table 2.1. Clustered standard errors in parentheses.

Table 2.5: Reported MPCs at Retirement

	20% Income Increase			20% Income Decline		
	MPC (pp.)	ln MPC	MPC $\neq 0$	MPC (pp.)	ln MPC	MPC $\neq 0$
DID Estimate	7.512	0.258	0.085	-1.440	-0.096	0.012
	(1.882)	(0.090)	(0.026)	(2.140)	(0.045)	(0.015)
N	2850	1234	2850	3113	2646	3113

Notes: This table reports difference-in-difference coefficients from a regression of reported MPCs out of a hypothetical income increase (decline) just before and just after retirement. Estimates are estimated following de Chaisemartin and D'Haultfœuille (2020) and are robust to treatment effect heterogeneity. Bootstrap standard errors clustered at the household level in parentheses.

Table 2.6: Shopping Behavior and Employment

	(1) Prices	(2) Prices	(3) Prices	(4) Quality	(5) Quality	(6) Quality
Panel A: Households Ages 25-54						
Male Part Time	0.190	0.234	-0.0482	-2.069***	1.830***	-0.315
	(0.158)	(0.158)	(0.131)	(0.335)	(0.315)	(0.229)
Male Not Employed	0.0533	0.152	0.0354	-5.225***	-0.211	-0.500**
	(0.109)	(0.109)	(0.117)	(0.198)	(0.194)	(0.176)
Female Part Time	-0.913***	-0.775***	-0.193*	-1.751***	0.472**	-0.346*
	(0.0900)	(0.0899)	(0.0788)	(0.162)	(0.151)	(0.137)
Female Not Employed	-0.922***	-0.762***	-0.478***	-3.212***	-0.0605	-0.630***
	(0.0733)	(0.0759)	(0.0823)	(0.137)	(0.129)	(0.141)
Income Controls		X			X	
Household FE			X			X
N	286923	286923	286923	286919	286919	286919
Panel B: Households Ages 55-74						
Male Part Time	-0.390**	-0.216	-0.0505	-3.056***	-0.421	-0.224
	(0.142)	(0.142)	(0.0818)	(0.351)	(0.294)	(0.176)
Male Not Employed	-0.520***	-0.309**	-0.0769	-3.887***	-0.536**	-0.645***
	(0.106)	(0.105)	(0.0698)	(0.197)	(0.186)	(0.121)
Female Part Time	-1.019***	-0.904***	-0.207**	-2.125***	-0.297	-0.607***
	(0.139)	(0.137)	(0.0786)	(0.263)	(0.246)	(0.147)
Female Not Employed	-0.698***	-0.571***	-0.191*	-2.790***	-0.423*	-0.751***
	(0.0970)	(0.0967)	(0.0795)	(0.198)	(0.185)	(0.144)
Income Controls		X			X	
Household FE			X			X
N	247646	247646	247646	247646	247646	247646

Notes: This paper reports the effect of work status of male and female household heads on household level prices and household bundle quality. The underlying sample is the baseline sample in Table 2.1, restricted to households with two household heads. Household age is defined as the average age of the two household heads, rounded to the nearest integer. The omitted category in both cases is a household head working full time (at least 30 hours a week). All regressions control for household size, the relationship between the household heads, indicators for the age and presence of children and year fixed effects. Regressions are weighted to be representative of the US population. Clustered standard errors in parentheses.

2.9 Appendix

Decomposing Expenditure

Since we are interested in decomposing quantities, prices, bulk effects, and the quality of the consumption bundle, we can proceed as follows. Fixing a "bundle" at the product-module level, let p denote prices at the item level, b denote bulk allocations, and μ denote quality, we have

$$e_i = \text{Cost of Bundle at } p_i,\ b_i,\ \mu_i$$

which we can express as

$$= \text{Cost of of Bundle at } \bar{p}_i,\ \bar{b}_i,\ \bar{\mu}_i \cdot \frac{\text{Cost of Bundle at } p_i,\ b_i,\ \mu_i}{\text{Cost of of Bundle at } \bar{p}_i,\ \bar{b}_i,\ \bar{\mu}_i}$$

This fixes the bundle at the product-module level. A household with a given bundle can pay more or less than the average by purchasing varieties of higher (or lower) quality, by paying more (or less) for a given variety k or by buying more (or less) in bulk. Taking logs, we have:

$$\begin{aligned}\ln e_i &= \ln\big(\text{Cost of of Bundle at } \bar{p}_i,\ \bar{b}_i,\ \bar{\mu}_i\big)\\ &\quad + \ln\big(\text{Cost of Bundle at } p_i,\ b_i,\ \mu_i\big) - \ln\big(\text{Cost of of Bundle at } \bar{p}_i,\ \bar{b}_i,\ \bar{\mu}_i\big)\end{aligned}$$

Below, I show that the deviation of the cost of the bundle deviation from the cost of the bundle at average barcode-level prices, the average bulk allocation, and average quality is additively separable in each of the sub-components. Therefore, we have:

$$\begin{aligned}\ln e_i &= \ln\big(\text{Cost of of Bundle at } \bar{p}_i,\ \bar{b}_i,\ \bar{\mu}_i\big)\\ &\quad + \ln\big(\text{Cost of Bundle at } p_i\big) - \ln\big(\text{Cost of Bundle at } \bar{p}_i\big)\\ &\quad + \ln\big(\text{Cost of Bundle at } b_i\big) - \ln\big(\text{Cost of Bundle at } \bar{b}_i\big)\\ &\quad + \ln\big(\text{Cost of Bundle at } \mu_i\big) - \ln\big(\text{Cost of Bundle at } \bar{\mu}_i\big)\\ &\quad = \text{Quantity Effects}_i + \text{Price Effects}_i + \text{Bulk Effects}_i + \text{Quality Effects}_i\end{aligned}$$

We can decompose expenditure into quantity, quality, price, and bulk effects as follows. Starting with expenditure, we have:

$$\text{Expenditure}_{i,t} = \sum_k P_{i,t,k} Q_{i,t,k}$$

Denoting the within-m average price by $\bar{P}_{m,t}$, we can decompose this into the actual bundle at product-module average prices and the ratio of the actual cost of the bundle and the cost of the bundle at module-average prices:

$$= \left(\sum_m \sum_{k\in m} \bar{P}_{t,m} Q_{i,t,k}\right) \cdot \frac{\sum_k P_{i,t,k} Q_{i,t,k}}{\sum_m \sum_{k\in m} \bar{P}_{t,m} Q_{i,t,k}}$$

Taking logs, we have

$$\ln \text{Expenditure}_{i,t} = \ln\left(\sum_m \sum_{k\in m} \bar{P}_{i,t,m} Q_{i,t,k}\right) + \ln\left(\frac{\sum_k P_{i,t,k} Q_{i,t,k}}{\sum_m \sum_{k\in m} \bar{P}_{i,t,m} Q_{i,t,k}}\right)$$

Importantly, the first term captures changes in quantities only (either increasing $Q_{i,t,k}$ for fixed k or switching the bundle between k and l for $k \neq l$). We can think of the last term as an aggregate shopping margin that is captures the log difference between the actual cost of the bundle household i decided to buy and cost of the same product-module-level quantities at product-module-level mean prices. We can decompose this adjustment margin as follows:

$$\textit{Aggregate Margin}_{i,t} = \ln\left(\sum_k P_{i,t,k} \cdot Q_{i,t,k}\right) - \ln\left(\sum_m \left(\frac{\sum_{k\in m}\sum_j P_{j,t,k} Q_{j,t,k}}{\sum_{k\in m}\sum_j Q_{j,t,k}} \cdot \sum_{k\in m} Q_{i,t,k}\right)\right)$$

Summing over the products of actual quantities and average prices within a product module first and then aggregating over product modules will yield the same total as directly summing over the products of actual quantities and average prices over all UPCs. Therefore, we can add and substract the actual bundle at average prices and get

$$\begin{aligned}
&= \ln\left(\sum_k P_{i,t,k} \cdot Q_{i,t,k}\right) - \ln\left(\sum_m \left(\frac{\sum_{k\in m}\sum_j P_{j,t,k} Q_{j,t,k}}{\sum_{k\in m}\sum_j Q_{j,t,k}} \cdot \sum_{k\in m} Q_{i,t,k}\right)\right) \\
&+ \ln\left(\sum_m \sum_{k\in m} \frac{\sum_j P_{j,t,k} Q_{j,t,k}}{\sum_j Q_{j,t,k}} \cdot Q_{i,t,k}\right) - \ln\left(\sum_k \frac{\sum_j P_{j,t,k} Q_{j,t,k}}{\sum_j Q_{j,t,k}} \cdot Q_{i,t,k}\right)
\end{aligned}$$

We can re-arrange this to yield

$$\begin{aligned}
&= \ln\left(\sum_k P_{i,t,k} \cdot Q_{i,t,k}\right) - \ln\left(\sum_k \frac{\sum_j P_{j,t,k} Q_{j,t,k}}{\sum_j Q_{j,t,k}} \cdot Q_{i,t,k}\right) \\
&+ \ln\left(\sum_m \sum_{k\in m} \frac{\sum_j P_{j,t,k} Q_{j,t,k}}{\sum_j Q_{j,t,k}} \cdot Q_{i,t,k}\right) \\
&- \ln\left(\sum_m \left(\frac{\sum_{k\in m}\sum_j P_{j,t,k} Q_{j,t,k}}{\sum_{k\in m}\sum_j Q_{j,t,k}} \cdot \sum_{k\in m} Q_{i,t,k}\right)\right) \\
&= \textit{Effort Margin}_{i,t} + \textit{Uncorrected Quality Margin}_{i,t}
\end{aligned}$$

Turning to the uncorrected quality margin, we can follow Griffith et al. (2009) and we break out each product module into five quintiles of the size distribution. Then, we can define UPCs in the top two quintiles (or the Top 40% of the within-product-module size distribution) as bulk items.[17] We can now add and subtract the cost of the quantities the household actually bought for a given product-module-by-bulk level but at the product-module-by-bulk mean price:

$$
\begin{aligned}
&= \ln\left(\sum_m \sum_k \frac{\sum_j P_{j,t,k} Q_{j,t,k}}{\sum_j Q_{j,t,k}} \cdot Q_{i,t,k}\right) - \ln\left(\sum_m \left(\frac{\sum_k \sum_j P_{j,t,k} Q_{j,t,k}}{\sum_k \sum_j Q_{j,t,k}} \cdot \sum_k Q_{i,t,k}\right)\right) \\
&+ \ln\left(\sum_{m\times b} \left(\frac{\sum_{k\in m\times b} \sum_j P_{j,t,k} Q_{j,t,k}}{\sum_{k\in m\times b} \sum_j Q_{j,t,k}} \cdot \sum_{k\in m\times b} Q_{i,t,k}\right)\right) \\
&- \ln\left(\sum_{m\times b} \left(\frac{\sum_{k\in m\times b} \sum_j P_{j,t,k} Q_{j,t,k}}{\sum_{k\in m\times b} \sum_j Q_{j,t,k}} \cdot \sum_{k\in m\times b} Q_{i,t,k}\right)\right)
\end{aligned}
$$

As before, we can rearrange this to get:

$$
\begin{aligned}
&= \ln\left(\sum_m \sum_k \frac{\sum_j P_{j,t,k} Q_{j,t,k}}{\sum_j Q_{j,t,k}} \cdot Q_{i,t,k}\right) - \ln\left(\sum_{m\times b} \left(\frac{\sum_{k\in m\times b} \sum_j P_{j,t,k} Q_{j,t,k}}{\sum_{k\in m\times b} \sum_j Q_{j,t,k}} \cdot \sum_{k\in m\times b} Q_{i,t,k}\right)\right) \\
&+ \ln\left(\sum_{m\times b} \left(\frac{\sum_{k\in m\times b} \sum_j P_{j,t,k} Q_{j,t,k}}{\sum_{k\in m\times b} \sum_j Q_{j,t,k}} \cdot \sum_{k\in m\times b} Q_{i,t,k}\right)\right) \\
&- \ln\left(\sum_m \left(\frac{\sum_k \sum_j P_{j,t,k} Q_{j,t,k}}{\sum_k \sum_j Q_{j,t,k}} \cdot \sum_k Q_{i,t,k}\right)\right) \\
&= \textit{Quality Margin}_{i,t} + \textit{Bulk Margin}_{i,t}
\end{aligned}
$$

Finally, we can put all of these pieces together and get:

$$
\begin{aligned}
\ln \text{Expenditure}_{i,t} &= \ln\left(\sum_m \sum_{k\in m} \bar{P}_{i,t,m} Q_{i,t,k}\right) - \textit{Quality Margin}_{i,t} - \textit{Effort Margin}_{i,t} - \textit{Bulk Margin}_{i,t} \\
&= \textit{Quantity Effects}_{i,t} - \textit{Quality Margin}_{i,t} - \textit{Effort Margin}_{i,t} - \textit{Bulk Margin}_{i,t}
\end{aligned}
$$

Intuitively, households can bring down the price of their product-module level bundle by paying less for the same UPC, by buying more in bulk, or by substituting towards lower quality items within a product module. In practice, not all UPCs in the KNCP data have interpretable quality information. For example, some households collect "magnet data" which refers to items that are sold by weight and do not have an associated UPC (e.g., fresh meat at the counter or by-weight produce). Households record these purchases in counts (instead of weight or volume), which makes interpreting the quantity for these items

[17] Note that I use *total quantities* to define item size. For example, a twelve-pack of cans of Coca-Cola will be defined as equaling 144 fluid ounces (12 times 12 fluid ounces), roughly similar to two-liter bottles of Coca-Cola (135.2 fluid ounces).

impossible. Therefore, I drop magnet data from my analysis. Another issue are durables and semi-durables recorded in “General Merchandise”. While most goods in this category have their own UPC (allowing me to compare prices paid for any specific item), it is unclear to what extent price differences within a product module will reflect quality differences or differences in some other attributes. For example, one such product module is calendars. While I can observe exactly how much any one household paid for a given calendar (at the UPC level), it is unclear to what extent all price differences truly reflect quality differences in the same narrow sense as this is true in the other levels of consumption I observe. A calendar might be bigger or smaller (e.g., have an individual page for each day or a page for every workweek) so that this isn't the same narrow comparison I make elsewhere. Therefore, I keep these purchases when estimating my main shopping effort measure but drop them for my quality comparisons. For robustness, I also compute the shopping effort measure after excluding these purchases.

Data Definitions and Details on Sample Construction

Matching the Nielsen Panel to the SCF

When matching information on asset holdings from the Survey of Consumer Finances to the Nielsen Panel, I follow the following procedure: For each iteration of the SCF from 2004 to 2019, I restrict the Nielsen Panel to the years surrounding the year the SCF was conducted (so I use KNCP information elicited in 2003, 2004, and 2005 to match to the 2004 SCF). I then match exactly on household structure (a cohabitating couple, a single female head, a single male head), household race (using the household head's race in the SCF and the reported “household race” from the Nielsen Panel), and on the employment status of the household head (employed or not employed). I the employ the coarsened exact matching algorithm by Iacus et al. (2012) to match on educational attainment of the household head (less than high school, high school, some college, a college degree, a post-graduate degree), income quintile (using the distribution of income in the Nielsen Panel), the number of children (no children, one child, two or more children), and the age of the household head.

Appendix Figures

Figure 2.A1: Product Hierarchy in Nielsen

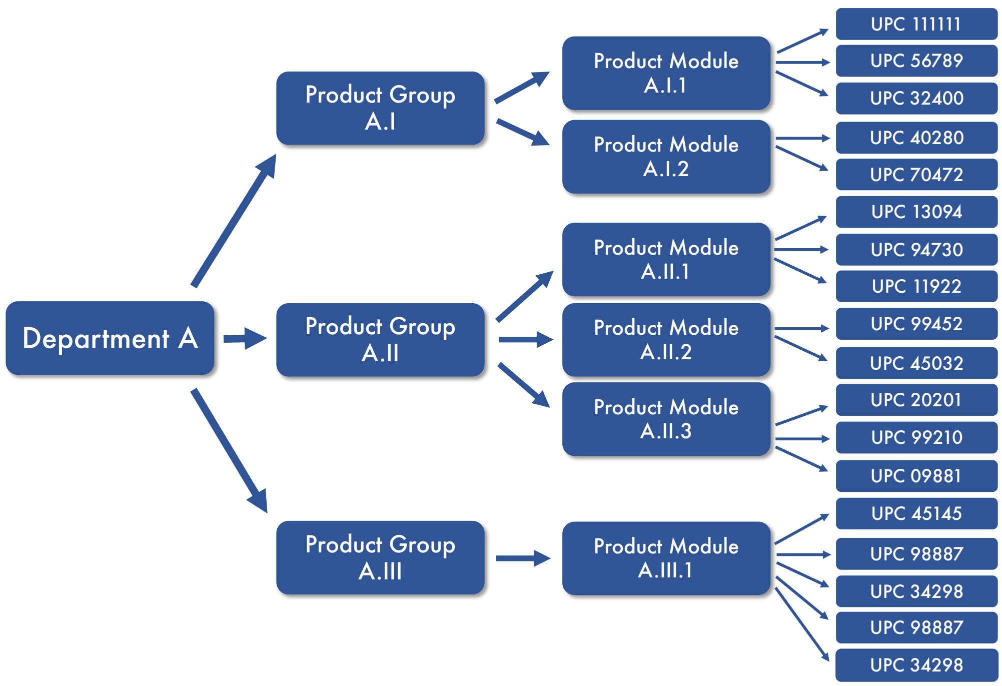

Notes: An illustration of the product hierarchy in Nielsen. Each of nine departments is organized into various product groups (a total of 110), which in turn are disaggregated into product modules (a total of 1,298). Each product group may nest one or more product modules and each product module will correspond to a large number of different UPCs.

Figure 2.A2: Distribution of Age at Retirement in Nielsen

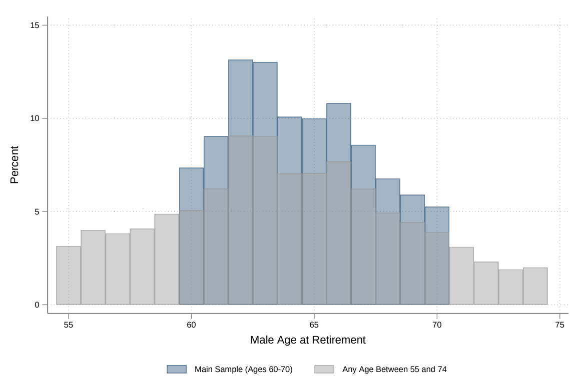

(a) Men

(b) Women

Figure 2.A3: An Illustration of UPCs capturing different items

Notes: An illustration of three different sets of UPCs that amount to an exactly identical quantity of an identical good.

Figure 2.A4: Illustration of the Effort Margin

(a) Store-Brand Eggs: $2.77

(b) Store-Brand Eggs: $3.80

(c) Name-Brand Pasture-Raised Eggs: $5.99

(d) Name-Brand Pasture-Raised Eggs: $7.49

Notes: An illustration of the effort margin. This figure shows two different products with different prices at different points in time or at different retailers. Shopping effort refers to buying a *given* UPC at a lower price than the average for that UPC. Prices are actual prices at specific stores in Alameda County and *not* taken from the Kilts Nielsen Nielsen Panel.

Figure 2.A5: Illustration of the Bulk Margin

(a) 6-Count Eggs: \$0.332 each

(b) 18-Count Eggs: \$0.229 each

(c) 36-Count Eggs: \$0.226 each

(d) 60-Count Eggs: \$0.202 each

Notes: An illustration of the bulk margin. This figure shows the prices of four distinct UPCs within the product module "Eggs" with different degrees of "bulkiness". Bulk savings means purchasing eggs in larger quantities in order to reduce per-unit prices (abstracting away from realized prices). The presented prices are actual prices at specific stores in Alameda County and *not* taken from the Kilts Nielsen Nielsen Panel.

Figure 2.A6: Illustration of the Quality Margin

(a) Store-Brand Eggs: $2.77

(b) Name-Brand Cage-Free Eggs: $3.69

(c) Store-Brand Organic Eggs: $3.99

(d) Name-Brand Pasture-Raised Eggs: $7.49

Notes: An illustration of the quality margin. This figure shows the prices of four different UPCs within the same product module at a single point in time. Quality adjustment means purchasing eggs that are cheaper or more expensive on a per-unit basis. Prices are actual prices at specific stores in Alameda County and *not* taken from the Kilts Nielsen Nielsen Panel.

Figure 2.A7: The Relationship Between Bundle Quality and Prices Paid

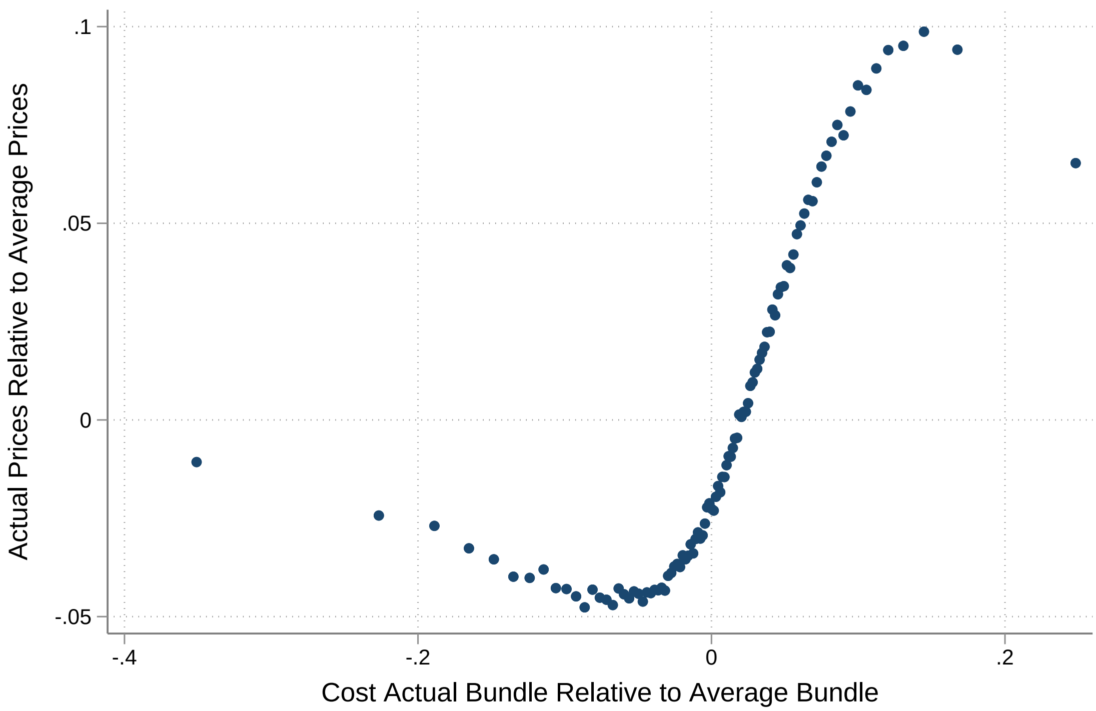

Notes: This figure presents a non-parametric binned scatter plot as proposed by Cattaneo et al. (2021). The effort and quality margins are as defined in Equations 2.7 and 2.10. The underlying data for this binned scatter plot is the baseline sample in Table 2.1.

Figure 2.A8: Decomposing the Post-Retirement Expenditure Drop

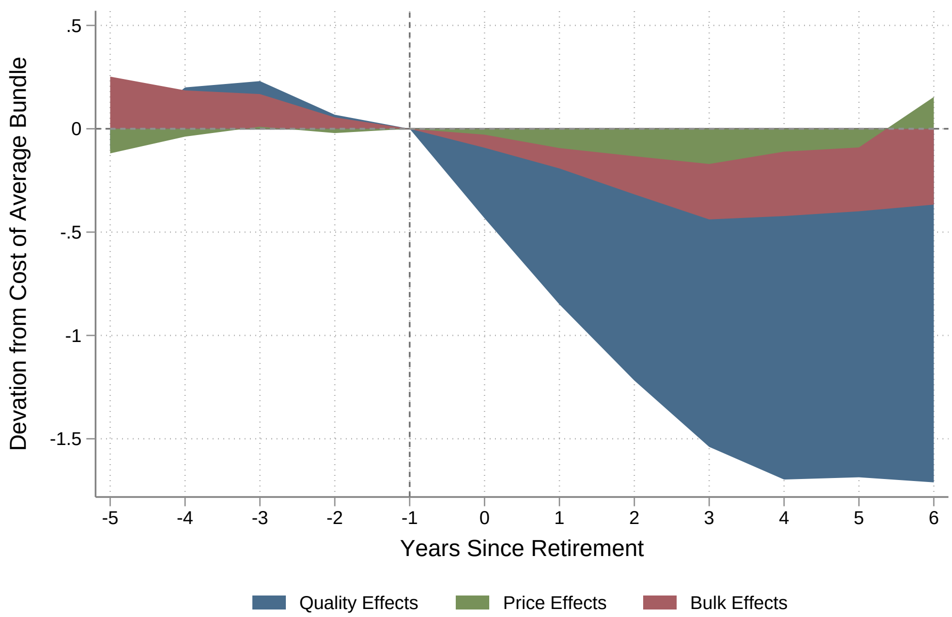

(a) Shopping Margins of Adjustment at Retirement

Notes: This figure presents the evolution of price effects, bulk effects, and quality adjustments estimated according to Equation 2.11 as well as uncorrected expenditure, expenditure at average prices and the average bulk allocation, and expenditure at average prices, the average bulk allocation, and average within-product module quality. The underlying data for each panel is the retirement sample of Table 2.1.

Figure 2.A9: Local Shopping Adjustments Around Retirement

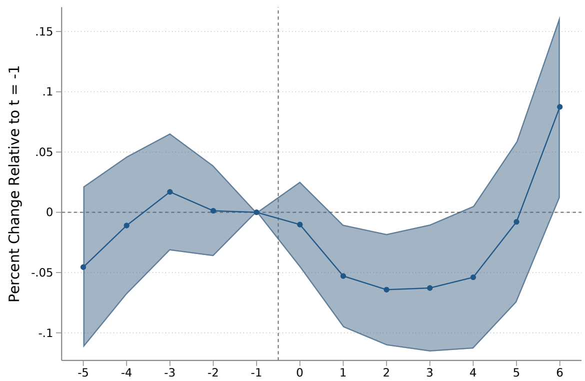

(a) Prices Paid

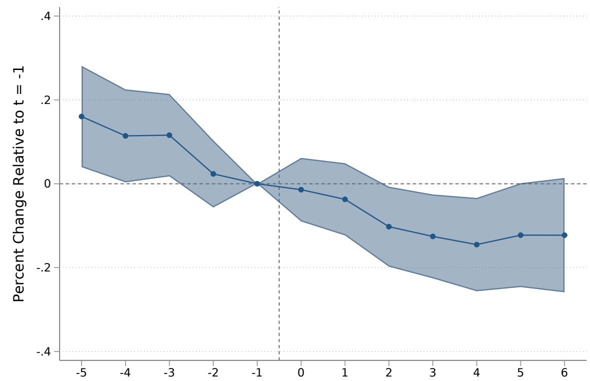

(b) Cost Relative to Average Bulk Allocation

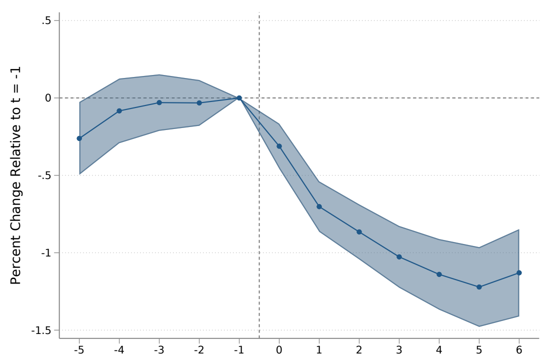

(c) Bundle Quality

Notes: This figure presents the evolution of price effects, bulk effects, and quality adjustments estimated according to Equation 2.11. All shopping margins of adjustment are defined in terms of deviations from local market averages. The underlying data for each panel is the retirement sample of Table 2.1. Standard errors are clustered at the household level.

Figure 2.A10: Further Validations of Quality Measures

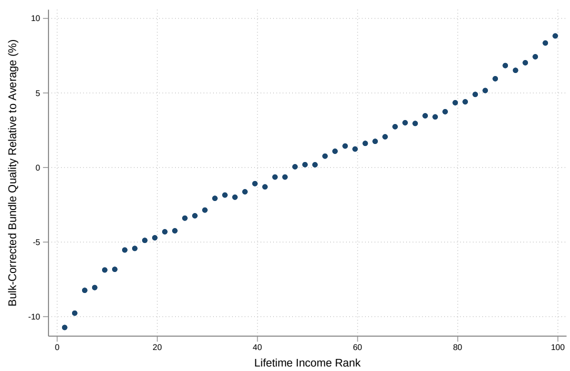

(a) Bundle Quality and Lifetime Income

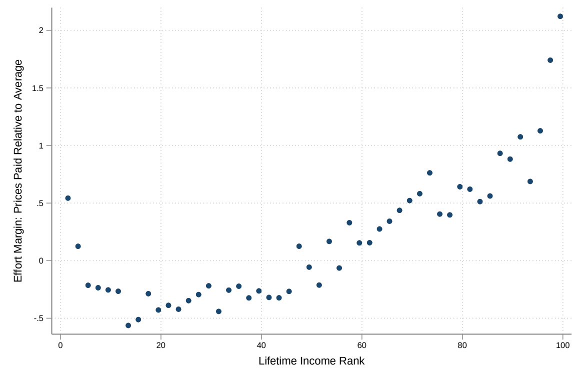

(b) Prices Paid and Lifetime Income

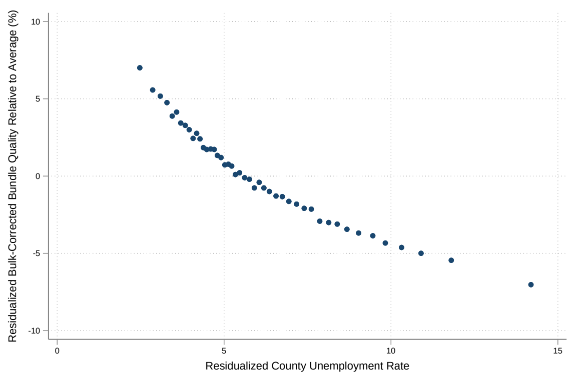

(c) Bundle Quality and Local Unemployment

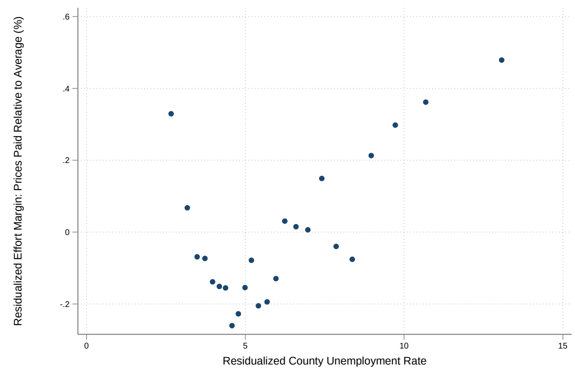

(d) Prices Paid and Local Unemployment

Notes: All figures present non-parametric binned scatter plots as proposed by Cattaneo et al. (2021) and control for household composition, household size, the age and presence of children, household age, a year fixed effect, and a scantrack market fixed effect. The effort and quality margins are as defined in Equations 2.7 and 2.10. The underlying data for Panels (a) and (b) is the set of households in the baseline sample who are in the data for at least three years; the data for Panels (c) through (d) is the baseline sample of Table 2.1.

Figure 2.A11: de Chaisemartin and D'Haultfœuille (2020) Estimates: Path of Income Around Retirement

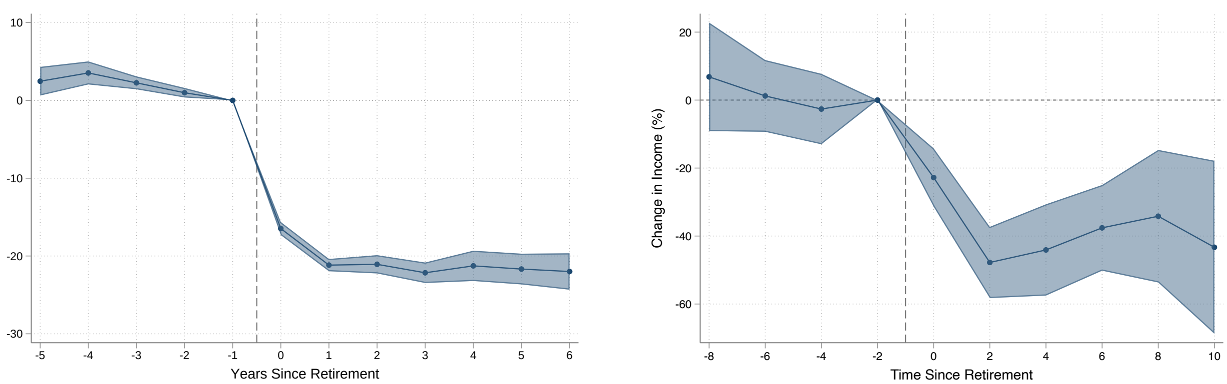

(a) Scanner Data (b) Consumption Survey

Notes: This figure presents the evolution of income around retirement in the KNCP and HRS estimated with the event study estimator proposed by de Chaisemartin and D'Haultfœuille (2020). The underlying data for Panel (a) is the retirement sample of Table 2.1; the underlying data for Panel (b) is the retirement sample of Table 2.2. Bootstrap standard errors are clustered at the household level.

Figure 2.A12: de Chaisemartin and D'Haultfœuille (2020) Estimates: Expenditure Around Retirement

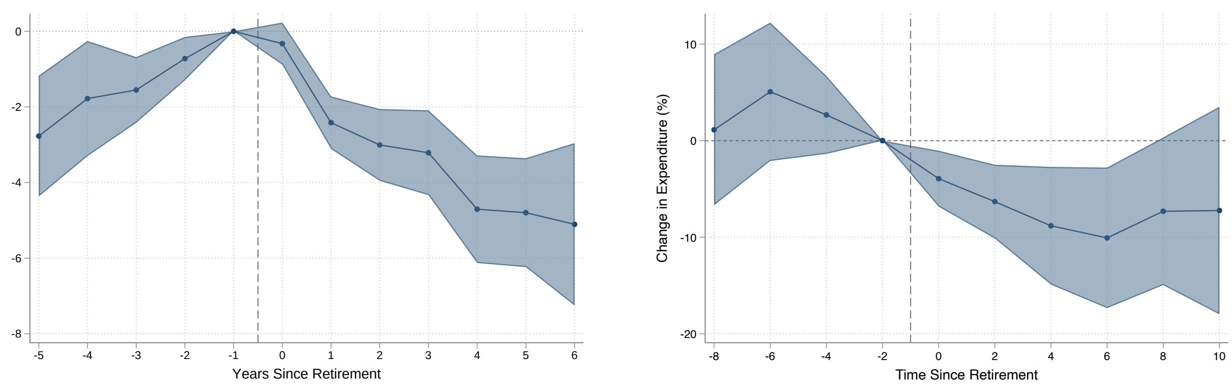

(a) Scanner Data (b) Consumption Survey

Notes: This figure presents the evolution of expenditures around retirement in the KNCP and HRS estimated with the event study estimator proposed by de Chaisemartin and D'Haultfœuille (2020). The underlying data for Panel (a) is the retirement sample of Table 2.1; the underlying data for Panel (b) is the retirement sample of Table 2.2. Bootstrap standard errors are clustered at the household level.

Figure 2.A13: de Chaisemartin and D'Haultfœuille (2020) Estimates: Price and Bulk Savings Around Retirement

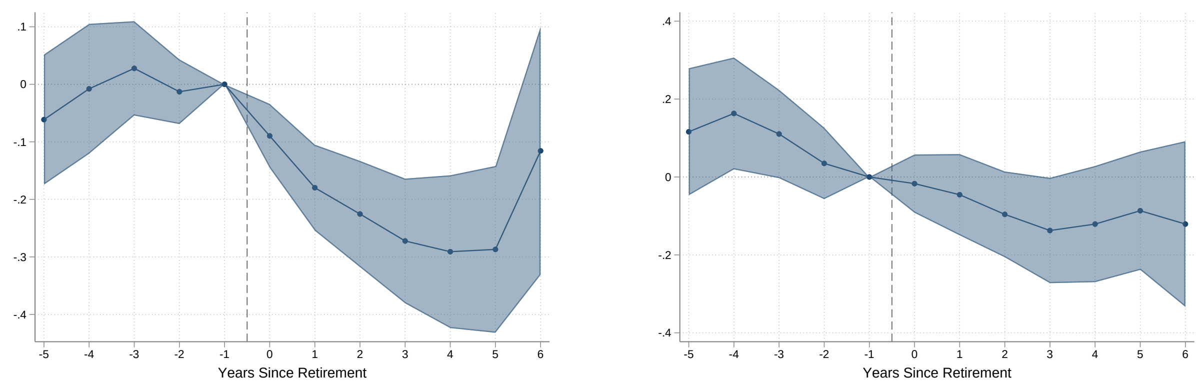

(a) Prices Paid (b) Cost Relative to Average Bulk Allocation

Notes: This figure presents the evolution of price and bulk effects around retirement estimated with the event study estimator proposed by de Chaisemartin and D'Haultfœuille (2020). The underlying data for both panels is the retirement sample of Table 2.1. Bootstrap standard errors are clustered at the household level.

Figure 2.A14: de Chaisemartin and D'Haultfœuille (2020) Estimates: Consumption Quality Around Retirement

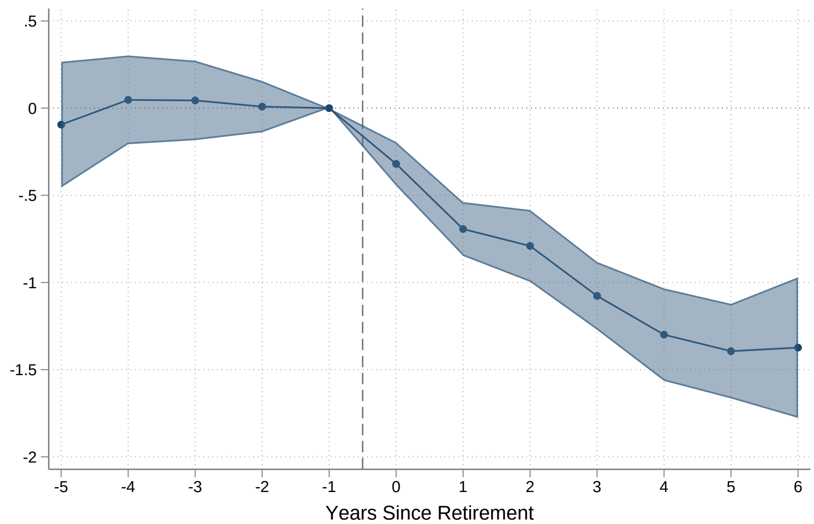

(a) Bundle Quality

Notes: This figure presents the evolution of quality adjustments around retirement estimated with the event study estimator proposed by de Chaisemartin and D'Haultfœuille (2020). The underlying data for both panels is the retirement sample of Table 2.1. Bootstrap standard errors are clustered at the household level.

Figure 2.A15: de Chaisemartin and D'Haultfœuille (2020) Estimates: Home Production, Home Consumption and Expenditure Timing in Retirement

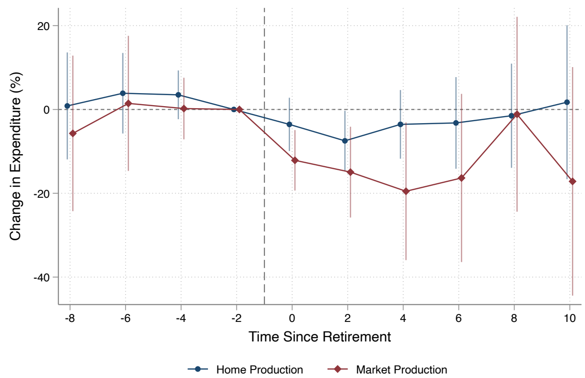

(a) Market and Home Production Expenditure

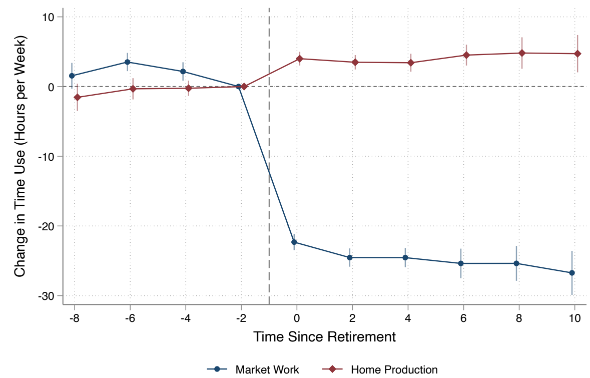

(b) Time Working vs. Home Production

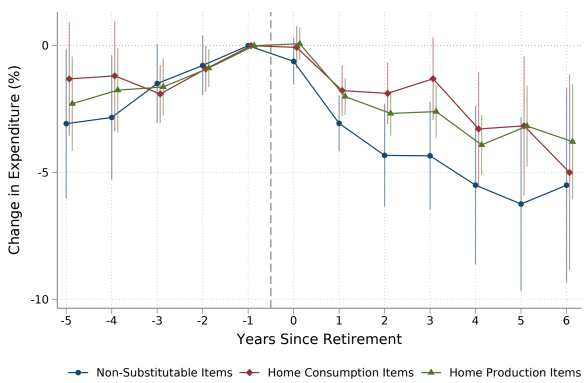

(c) Subsets of Spending by Substitutability

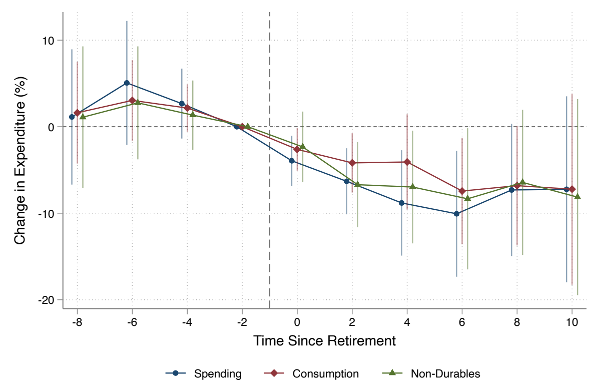

(d) Total vs. Consumption Expenditure

Notes: This figure presents the path of expenditure categories around retirement in the Kilts Nielsen Consumer Panel as well as expenditure categories and time use in the Health and Retirement Survey estimated according to Equation 2.11. Non-substitutable items are non-food grocery items, health and beauty items, and general merchandise. Home consumption items are prepared foods and deli items, food items that do not require preparation, and alcohol. Home production inputs are dry grocery items, packaged meat, fresh produce, frozen meat and vegetables, and dairy products. Home production items in the HRS are food at home, cleaning products, gardening products expenditures. Market production items are their immediate market good substitutes (food away from home, cleaning services, and gardening services). Home production time is time spent cleaning, doing gardening work, preparing meals and cleaning up and time spent shopping or running errands. The underlying data for Panels (a), (b), and (d) is the retirement sample of Table 2.2; the data for Panel (c) is the retirement sample of Table 2.1. Bootstrap standard errors are clustered at the household level

Figure 2.A16: Flows Into Non-employment in Nielsen and Aggregate Initial UI Claims

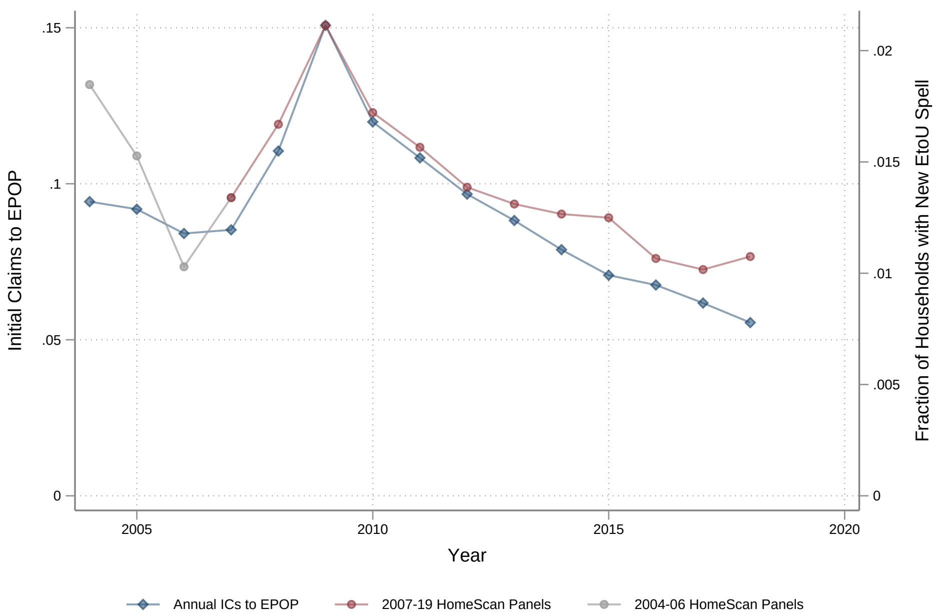

Notes: This figure compares flows into non-employment in the Kilts Nielsen Consumer Panel to the aggregate ratio of initial unemployment insurance claims to the employed population. In 2007, the Nielsen Nielsen Panel was expanded to include 50% more households and be a more accurate representation of national demographics.

Chapter 3

Unemployment Insurance as a Financial Stabilizer: Evidence from Large Benefit Expansions

Disclaimer: This project uses confidential credit record micro data provided through a California Policy Lab agreement with Experian, which has reviewed all results in this paper for inadvertent disclosure. All findings and opinions are those of the authors alone and do not necessarily represent the opinions of the California Policy Lab or Experian.

This chapter is based on joint work with Sreeraahul Kancherla.

One of the core rationales for government provision of unemployment insurance is a market failure: Households cannot fully insure against unemployment risk and they also use credit to sufficiently smooth consumption over the course of unemployment spells. This appears to be borne out in the data: Faced with unemployment, household borrowing is limited (e.g., Ganong and Noel 2019, Olafsson and Pagel 2018a). In addition, more generous unemployment insurance regimes appear to reduce delinquencies (Di Maggio and Kermani 2016) and foreclosures (Hsu et al. 2018). However, a lack of borrowing could also be explained by the effect of unemployment on permanent income. Going back to Jacobson et al. (1993), a famous result of the labor economics literature is that job loss induces *permanently* lower incomes. If this decline in permanent income is the dominant factor in driving household consumption, then declines in consumption and a lack of borrowing are not the result of credit failures but rather permanent earnings risk households cannot insure against. In this Chapter, I can provide some indirect evidence on this mechanism. As part of the US policy response to Covid-19, the generosity of the UI system increased substantially with maximum benefit durations reaching almost two years and replacement rates often exceeding 100% (Ganong et al. 2020b). If the effect of unemployment on consumption drops and delinquencies runs through permanent income, then this expansion of UI benefits should have at most modest effects. However, as we will see below, the increases in UI generosity substantially reduced the sensivity of local financial conditions to local unemployment rates. This suggests that liquidity is an important mechanism in explaining delinquency and default responses to unemployment shocks.

3.1 Introduction

Job loss induces substantial financial stress: Households experiencing temporary unemployment spells are more likely to default on their loans (Braxton et al., 2020; Hurd and Rohwedder, 2010) or to file for bankruptcy (Keys, 2018). Liquidity—as opposed to wealth—seems to be a crucial determinant of consumption smoothing behavior, with liquidity-constrained households appearing much more sensitive to adverse shocks (Gerardi et al., 2017; Ganong and Noel, 2020; Ganong et al., 2020a). An important policy question is the extent to which targeted liquidity provision from unemployment insurance (UI) benefits insulate households from these adverse financial effects of job loss. Empirical evidence in this area has primarily focused on the micro-level impacts of the UI system. Using survey data, Hsu et al. (2018) leverage heterogeneity in UI generosity across states and over time to show that workers' job loss translates into less financial distress during more generous benefit regimes.

At the macroeconomic level, Kekre (2021) shows that UI can stabilize aggregate economic conditions if unemployed households have higher marginal propensities to consume (MPCs) than employed households or if UI alleviates precautionary savings motives. In a similar vein, McKay and Reis (2021) show that unemployment insurance can insulate households from uninsurable income shocks and unemployment, potentially making counter-cyclical increases in UI generosity optimal. Finally, Landais et al. (2018) show that counter-cyclical UI benefit expansions can be welfare enhancing in matching models by increasing labor market tightness in slumps. However, *empirical* evidence assessing the magnitudes of these potential mechanisms is relatively limited. Hsu et al. (2018) use their micro results to argue that UI expansions during the Great Recession prevented 1.3 million foreclosures, but this result relies on a partial equilibrium analysis extrapolation that simply multiplies the micro elasticities with the benefit extensions during the Great Recession. The best empirical evidence of the role of UI in smoothing *aggregate* economic conditions comes from Di Maggio and Kermani (2016). Leveraging heterogeneity in local benefit generosity and estimating the effects of Bartik shocks on local economies, Di Maggio and Kermani (2016) find that more generous UI regimes attenuate the effect of adverse shocks on employment and earnings growth. The key mechanism in their analysis is the financial accelerator channel: Delinquencies rise by much less in more generous UI regimes, preventing banks from tightening lending standards in economic downturns.

We shed new light on this role of UI as a stabilizer of aggregate financial conditions, leveraging the enormous increases in unemployment rates during the Covid-19 pandemic. As part of its policy response, the United States engaged in an unprecedented expansion of the unemployment insurance system: Maximum benefit durations were increased from lows of 12 weeks to highs of up to 76 weeks and supplemental payments from $300 to $600 increased replacement rates to close to or substantially over 100% (Ganong et al., 2020b). In addition, eligibility requirements were loosened so that virtually all unemployed workers were eligible for UI, even independent contractors or those with inadequate pre-displacement earnings. Importantly, many of these changes mimic previous policy recommendations intended to make the unemployment insurance system a better tool for macroeconomic stabilization (Chodorow-Reich and Coglianese, 2019), making this period an ideal setting to empirically investigate the macroeconomic stabilization potential of the unemployment system.

To illustrate the Covid policy response's unparalleled magnitude, in Figure 3.1 we map unemployment and delinquency rates at the county level between 2019 and 2021. This exercise reveals a stark geographic pattern. Looking first at Panels (a) and (c), we see stark increases in county-level unemployment rates between 2019 and 2020. However, Panels (b) and (d) reveal that county-level delinquencies actually *decreased* over the same time period, reflecting a decoupling of unemployment shocks and delinquencies. This phenomenon stands in stark contrast to the Great Recession, in which both measures simultaneously spiked. Turning now to Panels (e) and (f), we see unemployment rates falling back down between 2020 and 2021 accompanied by a small continued decline in delinquencies. To what extent does this new disconnect between local unemployment and delinquency rates throughout the pandemic reflect the efficacy of different Covid policy choices in local financial stabilization?

In this paper, we show that a majority of this decoupling between delinquency and unemployment is attributable to Covid-era expansions in the UI system rather than other contemporaneous policy responses. We leverage a nationally representative sample of administrative credit records from Experian, aggregated to the county-month level. As a benchmark, we estimate the delinquency-unemployment sensitivity over time, using state-by-month fixed effects to absorb contemporaneous policy changes (which were generally set or carried out at the state level). We show that prior to the pandemic, local financial conditions were highly sensitive to local economic conditions with county-level unemployment rates being highly predictive of county-level delinquency rates. We then show that the sensitivity of local financial conditions to local unemployment rates collapsed during the time expanded UI was in effect. The sensitivity of auto loan delinquencies with respect to the local unemployment rate fell by 50%. For consumer credit cards, this decline in the delinquency-unemployment rate sensitivity was 66%.

In order to isolate the effect of UI expansion from other Covid-era policy responses that would have simultaneously changed unemployment sensitivity, we employ a staggered event study design around the withdrawal from the Federal UI program in late 2021. We argue that withdrawal was plausibly driven by political and ideological concerns about the generosity of the UI system rather than a response to state-level economic conditions. The delinquency-unemployment rate sensitivity increases substantially after withdrawal—by about the same magnitude as prior estimates of sensitivity drops—suggesting that the Covid-era drop in the delinquency-unemployment rate sensitivity was due to expanded UI benefits. These findings are qualitatively robust to alternate specification choices, and we also present placebo treatments that show that we do not see increases in delinquencies on loan types that are likely to be unaffected by UI policies. Using a back-of-the-envelope calculation that keeps the delinquency-unemployment rate sensitivity fixed at pre-UI expansion levels, we estimate that UI expansions prevented about 59% of counterfactual delinquency-months. This aggregate financial stabilization effect was in addition to other widely acknowledged benefits of UI benefit provision, such as sustaining aggregate household consumption.

This paper contributes to a growing literature on the effects of pandemic UI policies. Ganong et al. (2022) use bank account data to show that expanded UI played a significant role in explaining aggregate consumption dynamics but had very limited labor market effects. Replacement rates were so high that households receiving UI built up substantial savings buffers despite one-month MPCs out of UI ranging from 0.26 to 0.43. Similarly, Coombs et al. (2022) show that workers affected by the abrupt withdrawal from federal UI had relatively small job finding responses while the MPC out of the benefit cut was 0.52. Similar to our results on the effects of aggregate financial conditions, these high MPCs combined with the sheer magnitude of the UI policy response suggest substantial aggregate effects of the Covid-era UI benefit expansions. Unlike these more micro-level papers, our paper explicitly focuses on the aggregate effects of UI expansions. Rather than estimating whether any one household is more insulated from adverse shocks under a more generous UI regime, our paper therefore answers the question whether *macroeconomic* conditions can be stabilized with UI as a policy instrument.[1]

We also add to a nascent *empirical* macro literature on the benefits of UI provision. A large empirical micro tradition has attempted to separately estimate both the consumption-smoothing benefits (Gruber, 1997b; Ganong and Noel, 2019) and the job search disincentive costs of UI (Katz and Meyer, 1990; Card et al., 2007) towards calibrating models of optimal benefit provision (Baily, 1978; Chetty, 2006, 2008; Schmieder and von Wachter, 2017). On the other hand, past empirical research in macroeconomics has mainly focused on estimating aggregate labor market disincentive effects of the UI system (Hagedorn et al., 2013; Johnston and Mas, 2018; Chodorow-Reich et al., 2019; Boone et al., 2021). Our paper fills this research gap by providing direct evidence that counter-cyclical increases in UI generosity can have substantial benefits in improving aggregate *financial* conditions. As such, our paper is one of the first to provide empirical support for the theoretical work on the financial macro stabilization effects of UI.

Section 3.2 details the numerous federal stabilization policies enacted during the pandemic, with a focus on the UI system. Section 3.3 describes our credit bureau microdata and aggregation procedure in detail. Turning to empirics, Section 3.4 explains our first esti-

[1] An important caveat to our results is that we cannot separately identify whether our results are due to liquidity provision to the unemployed (potentially stimulating aggregate demand through high MPCs among the unemployed) versus reductions in pre-cautionary savings motives or spillovers to employed households. We offer suggestive county-level evidence in favor of the former view by estimating heterogeneous treatment effects of UI benefit withdrawal on counties with low and high unemployment rates (see Section 3.5 for more details).

mation strategy and estimates aggregate financial sensitivity to unemployment shocks over time. Section 3.5 then explains how we use the 2021 staggered federal UI withdrawal to calibrate UI-specific effects in attenuating financial sensitivity during the pandemic. Section 3.6 uses these estimates to provide an aggregate estimate for delinquencies prevented by pandemic UI policies. Finally, Section 3.7 concludes.

3.2 The Pandemic Policy Environment

Unemployment Insurance Policies

In this paper, we analyze the aggregate financial effects of introduction and withdrawal from pandemic unemployment programs enacted during the Covid-19 pandemic. These programs, which were first created as part of the 2020 CARES Act, had three major components. The Pandemic Emergency Unemployment Compensation (PEUC) first extended the maximum duration of unemployment benefits by 13 weeks. This largely mirrored prior ad-hoc federal benefit extensions during the Great Recession (see Figure 3.3, which plots the evolution of federal UI benefit duration extensions over time). Together with existing state-level policies—both existing statutory durations and automated cyclical UI extension triggers—total eligible benefit durations totalled up to a maximum of 99 weeks. The other two components of pandemic unemployment policy were novel and reflected a broad desire to provide rapid liquidity to affected workers. The Federal Pandemic Unemployment Compensation (FPUC) program introduced a $600 supplement to existing weekly benefit amounts, which increased replacement rates above 100% for low to medium wage workers.[2] The Pandemic Unemployment Assistance (PUA) program additionally extended benefits to otherwise ineligible workers, such as those who had otherwise exhausted benefit eligibility, independent contractors, or those with an insufficient working history.[3]

While PEUC and PUA were authorized through December 31st, the CARES Act originally set FPUC supplements to expire on July 26th. The program was not reauthorized despite congressional efforts[4] and was partially replaced by the Lost Wages Assistance Program (LWA), which instead provided a temporary six-week $300 UI benefit supplement until September 6th.[5] On December 27th, all three CARES programs—FPUC, PUA, and

[2]Ganong et al. (2020b) show that statutory replacement rates exceeded 145%. In sum, this component of the program paid out over $263 billion in benefits, totaling 7% of total personal income over this period.

[3]More information on unemployment agencies' implementations of these policies can be found [here]. The California UI system also provides an excellent and accessible breakdown of the various UI programs and their resulting changes to pandemic benefits [here].

[4]The 2020 HEROES Act, which would have extended benefit supplements, passed the House but was not taken up in the Senate [link].

[5]The program was made possible by presidential order, as LWA program funding came from redirected FEMA disaster relief funds originally earmarked in the 2020 CARES Act [link]. Using Chase bank account data, Ganong et al. (2022) find that receipt of LWA supplements was inconsistent and depended highly on state agencies; while most benefits were paid in September, some Wisconsin and New Jersey recipients

PEUC—were extended until March 13th as part of the new Continued Assistance Act.[6] PEUC and PUA duration extensions were renewed for a further 11 weeks, with FPUC reauthorized for a smaller \$300 supplement. These policies were extended for a final time on March 11th as part of the American Rescue Plan, which reauthorized i) the \$300 FPUC supplement as well as ii) new 29 week benefit duration extensions for PUA and PEUC claimants. As part of the bill, each pandemic unemployment insurance program was designed to expire September 4, 2021.

Following a weak jobs report in May 2021, however, some state governors expressed concern that UI benefit availability had suppressed workers' job search and was impeding economic recovery. As we highlight in Section 3.5 when discussing our empirical strategy, this belief was arguably driven by ideological, rather than financial, concerns.[7] 26 states consequently terminated access to FPUC benefits ahead of the scheduled September expiration (22 states in June, three in July, and one in August), generating relatively sharp state-level variation in both access and generosity of UI benefits. Figure **??**, which plots UI continued claims throughout 2021, highlights the stark nature of benefit expiry: almost 4.5 percent of the labor force (nearly 5 million people) lost access to UI in September, with another 1.5 percent of the labor force losing UI access during the early phase-out from June through August.[8] Other authors have leveraged this variation across states as a shock to UI benefit access, finding relatively small increases in job-finding rates but large MPCs out of UI benefits for benefit losers (Coombs et al., 2022).

In sum, these pandemic programs made the UI program substantially more generous even compared to past recessions. Workers eligible under regular claims had a potential benefit duration up to 99 weeks in some states, equalling benefit durations at the height of the Great Recession. Moreover, FPUC supplements greatly increased benefit levels; replacement rates for some workers almost tripled compared to normal program levels. The introduction of PUA also dramatically expanded access to benefits for otherwise ineligible workers. Figure 3.4, which plots the insured (IUR, red line) and regular unemployment rate (UR, in blue) over time from January 2000 to December 2021, depicts an immense increase in aggregate insured for even regular workers. Indeed, the IUR-UR ratio nearly doubled in March 2020 compared to the Great Recession, from about 50% to almost 100%. Taken together with special programs like PUA[9] (green line), the IUR-UR ratio was around 150% until federal

received benefits well into October. Given the haphazard nature of LWA payments, we are unable to cleanly assign receipt for different counties over time and do not include the program variation in our analysis.

[6] The CAA also authorized Mixed Earner Unemployment Compensation (MEUC), which provided \$100 supplements for self-employed workers receiving benefits.

[7] As an illustrative example, South Carolina governor Henry McMaster claimed in early May that "[the] labor shortage is being created in large part by the supplemental unemployment payments that the federal government provides claimants on top of their state unemployment benefits . . . [it] has turned into a dangerous federal entitlement, incentivizing and paying workers to stay at home rather than encouraging them to return to the workplace" [link].

[8] In addition, workers in early phase-out states who continued to receive UI benefits through the regular UI program lost access to the \$300 supplement.

[9] The regular IUR, taken from the BLS, does not include federal programs like PUA. We construct the

program expiry in late 2021.[10]

Many of these Covid-era changes to the UI system correspond to existing policy proposals improving the macroeconomic stabilization component of UI. Writing before the pandemic, Chodorow-Reich and Coglianese (2019) point out that UI had historically played a minor role in macroeconomic stabilization. Duration expansions are usually implemented with lags and only affect a small subset of workers (since relatively few workers become long-term unemployed). Baseline take-up rates of UI are also quite low at 30-50% (Blank and Card, 1991), implying limited scope for UI to stabilize aggregate economic conditions. In order to make UI into a macro stabilization tool, they made five recommendations: (i) increased eligibility and take-up of regular UI, (ii) full federal financing of the expanded benefit (EB) program, (iii) removing look-back provisions for EB, (iv) automatic extensions of benefit durations in times of very high unemployment, and (v) automatic increases in UI generosity during recessions. During Covid, UI eligibility and access increased massively through relaxations of earnings tests as well as the reduction of administrative hurdles. The EUC and PUA programs both extended benefit durations and provided supplemental payments that significantly increased UI replacement rates. In effect, we can think of Covid-era UI as a temporary implementation of (i), (iv), and (v). Therefore, these UI expansions provide an excellent framework to test whether these changes actually help in stabilizing macroeconomic conditions.

Other Pandemic Policies

It is important to note that the Covid legislative response included many non-UI policies that may have affected credit outcomes. For example, the CARES Act also instituted a mortgage forbearance program that allowed borrowers with federally backed mortgages to defer payments for up to 18 months.[11] Since forbearance immediately affects payment status on mortgage loans by deferring payments, we assume all changes in mortgage delinquencies are driven by either forbearance or the general Covid policy response outside of expanded UI. Another policy immediately affecting financial conditions is the ongoing moratorium on federal student loan payments. Payments were paused effective March 20th and the Office of Federal Student Aid also stopped collections on defaulted loans and set the interest rate on Department of Education-backed loans to 0%. Given that the vast majority of student

insured unemployment rate including pandemic programs by 1) computing the ratio of all-programs and continuing claims weeks (which include both regular claims and special federal programs), 2) multiplying by the regular IUR.

[10] Note that the ratio can surpass 100% since the two statistics' underlying populations do not exactly line up.

[11] Initially, the program allowed for 180 days of forbearance with a borrower-side option to extend forbearance for another 180 days. Borrowers with mortgages backed by Freddie Mac or Fannie Mae could extend forbearance for up to 180 more days, provided their account when into forbearance before February 28, 2021. Households with mortgages backed by the Department for Housing and Urban Development, the Department of Veterans Affairs, and the Rural Housing Service could request an additional 180 days of forbearance provided they first entered forbearance before June 30, 2020.

loans are federal loans, this policy meant that most student loans were reported to creditors as "current" starting on March 20, 2020.

Beyond credit market policies, the Covid policy response included many actions aimed at providing liquidity and insulating households from the economic fallout of the pandemic. The federal government provided three rounds of Economic Impact Payments ("stimulus checks") ranging from $500 to $1,400 per household member. The American Rescue Plan provided an expanded and fully refundable Child Tax Credit of $3,600 per child under the age of six (and $3,000 per child between the ages of six and seventeen). The CARES Act also instituted the Paycheck Protection Program, a policy designed to keep existing labor market matches intact by providing forgivable loans to employers provided they were mostly used to make payroll payments. For a more detailed description and analysis of many of these policies, see Chetty et al. (2020). Our empirical setup is designed to control for many of these federal policy changes, as well as other policies at the state level.[12]

Our general approach for isolating the effect of changes to the unemployment insurance system on local financial conditions is as follows. First, we restrict primary attention to credit cards and auto loans, which were both unaffected by explicit policy responses. Second, our baseline analysis focuses on county within state-by-month dynamics. As such, all common variation in delinquencies driven by the general policy response will be absorbed by the state-by-month fixed effect. Third, we focus on the effects of the local unemployment rate on local financial conditions. Many of the other policy responses were not directly targeted at the unemployed (and some of them were explicitly attempting to keep labor market matches intact): To the extent that we see declines in the sensitivity of local financial conditions to local unemployment, it is very likely that this "dampening" is driven by increases in the generosity of the UI system. We return to these points while discussing our design in Section 3.4.

3.3 Data

Our analysis principally makes use of aggregated credit bureau microdata matched to county-level Local Area Unemployment Statistics (LAUS) data.

Credit Data

Our credit bureau microdata comes from the University of California's Consumer Credit Panel (UC-CCP), which covers a nationally representative random 2% sample of households (together with associated borrowers and household members) with their associated credit histories for each quarter from 2004 to 2021. The data, which originates from Experian and is made possible through a data use agreement with the California Policy Lab, contains detailed information about credit holders over time: person-level variables include geographic

[12]Several states enacted their own policies to expand or mimic federal reforms: for example, California provided two rounds of stimulus checks for state residents.

identifiers, demographic information, credit scores, bankruptcies, and new inquiries for credit. A novel aspect of our data relative to other credit bureau data sets is that we additionally also see raw tradeline-level information about each loan, such as monthly payment history, credit limits and balances, loan type (e.g., credit card vs auto loan), delinquency status, and deferments.[13]

Our principal goal is to construct detailed measures of aggregate financial distress over the Covid-19 pandemic. We therefore begin by extracting person-level records between the first quarter of 2017 and the first quarter of 2021. By leveraging the loan-level payment history information, we then reconstruct a monthly panel of loan-level delinquencies, linked to each consumer and their county of residence over time.[14] We aggregate these person-level records to the county-month level, separated by loan type (e.g., credit cards, auto loans, mortgages), to form our main analysis data set. To better understand potential aggregate demand-side credit responses to unemployment shocks, we similarly construct and merge on county-month counts of new loans, new loan balances, and new loan inquiries. Table 3.1 describes our final aggregate credit data set, which covers a balanced panel of 3,107 counties, 5.7 million unique consumers, and over 30 million unique loans between January 2017 and March 2022.

Employment Data

We obtain county-level monthly employment and unemployment rates from the Local Area Unemployment Statistics (LAUS), as published by the Bureau of Labor Statistics.[15] County-level LAUS data are not seasonally adjusted and are available for virtually all counties and county-equivalents in the US. Our main measure for local economic conditions is the the county-level unemployment rate. One potential issue with using not seasonally adjusted data is that the unemployment rate is highly cyclical (see again Figure 3.4). However, it seems reasonable to think that unemployment insurance insures against both business cycle and seasonal fluctuations. Ex-ante, it is not clear to what extent households behave differently in response to seasonal vs. cyclical unemployment. That being said, our results are robust to seasonally adjusting both the unemployment and the delinquency rates as most of the seasonal variation of unemployment will get absorbed by a state-by-month fixed effect in our baseline specification.

[13] By comparison, other credit panels (such as the New York Fed's Consumer Credit Panel) are often "rolled-up" to the person-level and may not include associated borrowers or household members. Further background information on this data and comparisons to other credit panels can be found on the UC-CCP's website here.

[14] In particular, we utilize the fact that for each loan Experian also reports the last 64 months of payment history. We extract and reshape these payment histories to form a monthly dataset. In Figure 3.A1, we benchmark our constructed panel against public CFPB mortgage delinquency data, finding that a very similar percentage of mortgages are delinquent over time in both data sets. Indeed, the principal differences for 90+ day delinquencies stems from the fact that the CFPB's publicly available data is rounded to the nearest tenth.

[15] The LAUS data can be downloaded here. Our data was downloaded as of May 3, 2021.

Recent work by Boone et al. (2021) has argued that LAUS data may not be the best measure when estimating the aggregate labor market effects of UI policies, as LAUS relies on state-level information to impute county-level unemployment rates. We are interested in the interaction between local economic conditions and local financial distress, however, rather than conditions themselves. In addition, our main specification will include a state-by-month fixed effect that should purge local unemployment rates of the common state component of unemployment for all counties in a given state. A related but separate concern in using LAUS data is that small county populations may generate large sampling error in the unemployment rate or delinquency rates. To mitigate this concern, 1) we drop small counties with credit data on fewer than 10 people and 2) weight by population size in all specifications. Our results are also qualitatively robust to increasing the county size restriction.[16]

3.4 How Elastic are County Delinquencies to Local Unemployment?

We start by qualitatively examining how county-level delinquencies respond to increases in local unemployment rates. The intuition for our analysis here is straightforward: since UI benefits provide needed liquidity to otherwise constrained households (Ganong and Noel, 2019), benefit expansions should attenuate aggregate delinquency responses to unemployment (Di Maggio and Kermani, 2016). Descriptively, we would therefore expect a reduced effect of unemployment rate increases on local delinquency rates after Covid UI policies are enacted. To begin, Figure 3.5 plots a binned scatter plot of overall county delinquencies against local unemployment rates separately both during Covid (March 2020 to August 2021) and pre-Covid (January 2017 to February 2020). We find the stabilization prediction bears out in the data: we see a much larger pre-pandemic (in red) slope compared to during the pandemic (in blue).

To better understand these patterns, we extend this setup to a regression framework with explicit dynamics and disaggregation by credit type (e.g., credit cards or auto loans).[17] First, we define the county-level delinquency rate for credit type k and county c in state $s(c)$ and month t as $y_{s(c),k,t}$. We regress $y_{s(c),k,t}$ on a state-by-month fixed effect and the local unemployment rate interacted by time dummies in the following estimating question:

$$\text{Delinquency Rate}_{s(c),k,t} = \delta_{s(c),t} + \sum_{\tau}^{T} \beta_t \, UR_{c,t} \cdot \mathbb{1}\{t = \tau\} + \varepsilon_{c,t} \tag{3.1}$$

Our main objects of interest, $\{\beta_t\}_{\tau}^{T}$, summarize the impact of a 1 percentage point increase in county-level unemployment on the county's aggregate delinquency rate for each month

[16] One reason for this is that we have relatively few small counties: see Figure 3.A2, which shows the distribution of average observed county populations in the credit data.

[17] Disaggregation by credit type is particularly important during the Covid pandemic period. As discussed in Section 3.2, policies such as the student loan payment moratorium or mortgage forbearance affected discrete credit groups, so estimating any-delinquency outcomes masks substantial heterogeneity across types.

from January 2017 to March 2022. To reduce expositional clutter in what follows, we refer more concisely to these treatment effects β_t as the *delinquency-unemployment sensitivity* in each period. Our choice of estimating equation is motivated by three considerations. First, UI expansion was not the only policy response during this time period: other policies, both directly within the credit market (mortgage forbearance, the student loan payment moratorium) and in providing stimulus (e.g., economic impact payments and the expanded Child Tax Credit) could have also affected this sensitivity over time. Since these reforms were largely invariant to the local unemployment rate, direct effects should be captured in the time component of the fixed effect $\delta_{s(c),t}$. Many other policies that may be affected by unemployment shocks, such as Extended Benefit triggers, are set at the state level and so separating out the state-time component using our fixed effect is also important for dealing with these contemporaneous confounders. Third, since the regular unemployment insurance system is a state-run program, treatment variation is at the state level.

Turning now to results, Figure 3.6 plots the coefficients $\{\beta_t\}_{\tau}^{T}$ separately for credit cards, auto loans, mortgages, and student loans for each month between January 2017 and March 2022. We prefer to report disaggegated estimates this way to ensure comparability across our sample time period, since mortgage and student loan payment obligations changed during the pandemic. In all plots, the shaded area shows periods when pandemic UI policies including UI supplements were in effect: darker grey implies full pandemic UI, while light grey starting in June 2021 denotes the beginning of UI phase-outs.[18]

We focus on auto loans and credit cards, which were not subject to any Covid policies or credit reporting changes. A first striking feature of these graphs is their cyclicality: the delinquency-unemployment sensitivity tends to rise in the fall months and fall each spring during the pre-pandemic period. After the start of the pandemic, this pattern changes: credit cards, for example, exhibit a mostly flat estimated sensitivity while pandemic UI is in effect. A second notable result is the substantial drop in estimated sensitivity during the pandemic period. Looking first at credit cards throughout the pandemic period in Panel (b), a 1 percentage point increase in the county unemployment rate is associated with roughly a 0.075 percentage point increase in county-level delinquencies (about 0.05 in the first shaded UI period, 0.1 in the second period). This average represents a 66% drop in the average sensitivity compared to the pre-pandemic period value of 0.225, indicating that Covid policies were associated with to substantial reductions in county-level credit card delinquency risk. Panel (a) highlights a similarly large effect for auto loans, at least early in the pandemic: between March and August of 2020, the average delinquency-unemployment sensitivity was about 50% lower than its pre-pandemic average (0.4 to 0.2).

[18] While major Covid economic stabilization policies began at the end of March 2020 with the CARES Act, March 2020 can be regarded as potentially treated due to lenders preemptively waiving delinquency reporting in expectation of federal legislation. In contrast to later federal legislation only covering student loans and mortgages, many prominent *credit card* issuers (including Goldman Sachs, US Bank, Truist, and Discover) also announced temporary forgiveness programs for March 2020. This MarketWatch article provides an illustrative sample of popular news coverage on preemptive supply-side credit policies at the time [link]. Anecdotally, lenders ceased idiosyncratic delinquency waivers after the introduction of the CARES Act.

A third takeaway from these plots is that these drops are largely coincident with the shaded areas when UI policies are in effect, and estimates appear to increase in reaction to policy withdrawals. For credit cards, for example, this picture is especially stark: the only increases in the delinquency-unemployment sensitivity are during the unshaded non-pandemic UI periods. We interpret this timing as suggestive potential evidence that our results flow through a UI-liquidity channel, as most other policies were unaffected by contemporaneous UI expirations. This pattern is intuitively quite plausible given the notably more generous pandemic UI policy environment. Recall that UI benefit replacement rates often exceeded 100% (Ganong et al., 2020b), so unemployed workers were receiving more income than before during employment. Looking at this in bank account microdata, Ganong et al. (2022) find that both income and aggregate checking account balances for the unemployed were about 20% and 50% higher respectively than *employed* workers (matched on pre-displacement characteristics).[19] Given this context, and assuming roughly similar debt spend-down out of UI and earned income, the additional benefits appear a strong candidate explanation for these sensitivity drops. We return to this point in Section 3.5, where we utilize the staggered expiration of UI benefits to directly estimate the proportion of the delinquency-unemployment sensitivity drop that is attributable to UI.

One concern is that we may be measuring reductions in *reported* financial distress instead of actual financial distress: creditors may have simply not reported delinquencies during the pandemic. As a data validation check, panels (c) and (d) of Figure 3.6 reestimates Eq. 3.1 for student loans and mortgages, where we know delinquencies were not reported. Looking first at student loans in Panel (c), while we see similar (though more muted) pre-pandemic cyclicality to credit cards and auto loans, we see consistently near zero sensitivity during the pandemic. We interpret this as a useful check on our data: due to the student loan payment moratorium, we should indeed see no reported delinquencies. Panel (d), covering mortgages, also provides a similar validation check as a federal mortgage forbearance policies were in effect between March 2020 and August 2021. In this case, however, our estimates are relatively small rather than zero. This finding reflects two factors. First, not all mortgages were necessarily subject to forbearance policies; the CARES Act policy only applied to federally-backed mortgages, such as those through Fannie Mae, Freddie Mac, Veterans Affairs, or the Federal Housing Administration. While some private mortgage servicers may have followed the federal policy, in the data we see some servicers reporting delinquencies during the pandemic. As of 2018, federally-backed mortgages reflected about 70% of all mortgages (Housing Finance Policy Center, 2020); we therefore interpret this percentage as a lower bound on the number of mortgages potentially affected by forbearance. Secondly, forbearance policies were enacted upon request rather than automatically through servicers: while we do not observe forbearance enactment for individual mortgages, incomplete take-up of this option may further explain nonzero estimated sensitivity. Even despite these two factors, however,

[19] Using a constructed series of redistributed national accounts data, Blanchet et al. (2022) additionally find that UI distributions constituted about a third of monthly income for bottom 50% households (see Figure 8, in particular, of their paper).

we see a large and relatively consistent drop in the mortgage delinquency-unemployment sensitivity during the pandemic.[20]

Next, we probe our estimates for robustness. For our results so far, we follow the typical definition and define loans as delinquent if they are over 30 days past due. A reasonable question for the financial stabilization interpretation is whether our results largely reflect continued nonpayment on older loans or new nonpayments. To examine this point, we re-estimate Equation 3.1 by instead using the shorter term 30-89 day delinquency rate to better capture short-run nonpayments. The results are qualitatively very similar; auto loan sensitivity seems largely driven by short-term delinquencies, while credit cards are more evenly split between short and longer-term delinquencies (our estimated levels and drop are about half of the previous all-delinquency estimate). One other concern with our outcome variable construction is that our results could be mechanically driven by demand-side responses for additional credit during the pandemic: if consumers take out additional loans, then the aggregate delinquency rate (delinquencies as a fraction of all loans) would mechanically decrease. Figure 3.8 thus re-estimates Equation 3.1 but replaces the delinquency rate with new loans per capita, disaggregating into auto loans and credit cards. We find little evidence of compensating loan count increases that would drive our results: while some point estimates for credit cards are statistically significant, they are largely precisely estimated near zero. Indeed, the largest estimates for credit cards imply a 0.005 increase in loans per capita for each percentage point increase in the unemployment rate.

3.5 Effects of the Pandemic UI Phase-Out

To what extent do these reductions in the delinquency-unemployment sensitivity reflect UI versus other contemporaneous Covid policy changes? In this section, we disentangle these effects by exploiting the aforementioned staggered loss of benefits for UI claimants across states between July and September 2021.[21] These withdrawals happened relatively quickly: looking across the 22 states that withdrew from federal UI programs in June, public announcements typically gave a month or less of forewarning for the policy change. These withdrawals were unlikely to have been driven by local government budgetary conditions: the federal program would have expired in September regardless, and all spending on UI benefits was covered by federal funds.

A common public interpretation was that the withdrawals were motivated by political considerations rather than labor market conditions, consistent with other research that highlights the role of political polarization as impetus for recent state-level policy changes

[20] Our estimates for mortgages increase substantially towards the end of our sample period, possibly reflecting the fact that the mortgage forbearance program ended in August 2021.

[21] Coombs et al. (2022) use the same variation in related work to examine employment and earnings responses in payroll-linked banking data, finding relatively small increases in job-finding rates and aggregate earnings increases of $900 million for benefit-losing workers in early withdrawal states. These workers also lost access to about $7.6 billion total in UI transfers, however, constituting a substantial aggregate net loss in income for affected households.

(DellaVigna and Kim, 2022). Indeed, an illustrative public announcement from Gov. Brad Little of Idaho signalled broader ideological opposition to continued UI benefits, saying in mid-May that his *"decision [was] based on a fundamental conservative principle – we do not want people on unemployment"* [link]. Reflecting this consideration, 21 of the 22 early withdrawal states were led by Republican governors; the sole Democratic governor, John Bel Edwards of Louisiana, led a largely Republican-leaning state (58.5%-39.9% Republican-Democrat vote shares during the 2020 presidential election).

We exploit the sharp timing of these changes in an event study framework to examine how UI withdrawal affected the delinquency-unemployment sensitivity. The key variation is across different states' month of exit from federal UI policies: given that these withdrawals were politically motivated, we see these events as plausibly uncorrelated with local credit market conditions. Following our previous specification, we estimate a dynamic event study variation of Equation 3.1 that also includes state-by-month fixed effects:

$$\text{Delinquency Rate}_{c,t,m} = \delta_{s(c),t} + \sum_{\tau}^{T} \beta_t \, UR_{c,t} \cdot D_{s(c),\tau} + \varepsilon_{c,t} \tag{3.2}$$

where now $D_{s(c),\tau}$ is an indicator that equals one if county c in state s withdrew from federal pandemic UI programs in month τ. In all regressions, we use a balanced sample of counties and plot estimates for 6 months before and 5 months after the policy change to allow for visual inspection of pre-trends. As before, we again disaggregate delinquency rates by loan type to ensure comparability over time and to the previous set of results.

We present our estimates for auto loans and credit cards in Figure 3.9. As before, we begin by discussing results for the first two categories. We see little evidence of pre-trends for auto loans or credit cards: point estimates before state-level UI withdrawals are near zero and statistically insignificant. Moreover, both credit types show a sharp effect of withdrawal on the delinquency-unemployment sensitivity: after about 4 months, the estimated sensitivity increases by about 0.2 percentage points (or 68%) for auto loans and 0.13 percentage points (144%) for credit cards. These treatment effects are qualitatively quite large, constituting 68% and 144% increases respectively compared to the month before withdrawal.

We now compare these treatment effects to the total sensitivity drops in Panels (a) and (b) of Figure 3.6. There, the sensitivity change after the introduction of Covid policies is about -0.15 for credit cards and -0.2 for auto loans. If we assume UI withdrawal had similar or symmetric effects on local financial stabilization to pandemic UI introduction, our phase-out estimates imply that the UI channel represents the vast majority of the total stabilization arising from Covid pandemic policies: almost all of the auto loans sensitivity drop, and 86% of the credit card sensitivity drop. Given the substantial amount of relief policies passed during the pandemic, both directly within the credit market (mortgage forbearance, the student loan payment moratorium) and in providing stimulus (e.g., economic impact payments and the expanded Child Tax Credit), we interpret this as strong evidence for substantial aggregate financial stabilization provided by the unemployment insurance system.

We conclude this section by considering three potential extensions and robustness checks for our estimates. In Figure 3.10, we re-estimate regressions for auto loans and credit cards using a 30-89 day delinquency measure to assess the extent to which our estimates may reflect newer or older nonpayments. Our results largely mirror the previous discussion of Figure 3.7: while the short-term auto loans estimates are about 2/3 of the total sensitivity increase (about 0.14 of the previous 0.2 increase after 4 months), about half of our credit card estimate appears to be driven by shorter-term delinquencies. We also again test whether our estimated sensitivity changes could be driven by demand-side changes in the number of loans taken out by consumers. Figure 3.11 estimates the effect of the phase-outs on the per-capita number of loans in each county. As before, our estimates are economically and generally statistically insignificant: the largest estimates, for auto loans after 4 months, imply a 0.002 change in per capita loans after a 1 percentage point change in the unemployment rate.

3.6 How Did Pandemic UI Affect Aggregate Delinquencies?

We conclude our analysis by providing a back-of-the-envelope calculation for the amount of aggregate delinquencies prevented by federal UI policies during the Covid pandemic. Our framework is motivated by our previous intuition for macro effects of UI: since benefit expansions provide increased liquidity to harder-hit counties, they effectively attenuate aggregate delinquency responses to unemployment shocks. To construct a macro counterfactual, we should thus reset the aggregate delinquency-unemployment sensitivity to empirical pre-pandemic levels, and calculate the difference between observed and otherwise predicted delinquencies over time. We illustrate these ideas, first in a simplified way and using Figure 3.12 as a visual aid. Panel (a) starts with a stylized reproduction of Figure 3.5, the empirical delinquency-unemployment relationship before and during the pandemic. As represented in panels (b) and (c), under a simplified attenuation framework UI policies can only impact delinquencies through a change in the curves' slope. Differences in intercepts thus reflect other existing Covid policies, such as stimulus checks or CTC expansion. Panel (d) illustrates our proposed calculation for aggregate delinquency effects: after removing intercept differences, the distance between the pre-Covid and during-Covid curves represent the prevented delinquencies at each value of the unemployment rate. We can thus sum across unemployment rates to yield the total number of delinquencies prevented.

We extend these base ideas to a fully dynamic framework, just as before in Section 3.4. One complication is that delinquencies are not an absorbing outcome, so a delinquency prevented in a given month does not imply that the delinquency cannot occur later on. We thus compute delinquency-months as our preferred measure of prevented financial distress. Our implementation proceeds in several steps. First, we re-estimate an augmented form of Equation 3.1:

$$\text{Delinquency Rate}_{s(c),k,t} = \delta_{s(c),t} + \alpha \mathbb{1}(t \in [\underline{\tau}, \bar{\tau}]) + \beta_t \, UR_{c,t} + \tilde{\beta} \mathbb{1}(t \in [\underline{\tau}, \bar{\tau}]) \, UR_{c,t} + \varepsilon_{c,t} \quad (3.3)$$

where $[\underline{\tau}, \bar{\tau}]$ is a shorthand for the Covid UI period, between March 2020 and August 2021. The first new coefficient in our estimation, α, provides for a level shift in delinquencies after the introduction of pandemic UI. The second term, $\tilde{\beta}$, separately estimates a direct shift in the delinquency-unemployment sensitivity in the same period. In essence, we will "turn off" these pandemic policy effects to construct our counterfactual delinquency series. Note that this way of constructing counterfactuals is quite conservative: We assume that increased UI generosity does not have any effect after August 2021, an assumption which undercounts prevented delinquencies if expanded UI benefits allowed households to build up precautionary savings. We use these estimates to construct two new monthly series for our counterfactual calculations, as seen in Figure 3.13. We begin with Panel (a), which proceeds for auto loans. The blue line plots fitted values from Equation 3.3, representing the estimated evolution of the delinquency rate. As a reassuring check on our estimation, this series roughly matches the dynamics of actual observed delinquencies over time (grey line). The red line, however, instead plots fitted values where the α and $\tilde{\beta}$ effects are removed from the blue line between March 2020 and August 2021. This second series thus represents a designed counterfactual where we have removed the effect of federal Covid policies. We can then calculate the number of monthly prevented delinquency-months as the difference between our estimated counterfactual (red) and estimated status quo (blue) series for each month, multiplied by the number of loans for that credit type in our data. To arrive at a total sum for delinquency-months prevented, we simply sum this measure over the Covid UI period, between March 2020 and August 2021.

This back-of-the-envelope calculation delivers stark results. For credit cards, we estimate that UI prevented about 59.3% of all potential delinquency-months in this time frame; for credit cards, we estimate a slightly larger net effect of about 59.6% of potential delinquency-months. While these effects are quite large, this came at a price: total federal pandemic UI program spending was about $674 billion[22], implying a cost of about $8,864 per delinquency-month prevented across the two credit types. Note that this estimate computes the direct cost; our results cannot identify the effects on other types of credit, overall credit smoothing, or aggregate spending-side responses that would all mitigate the final cost figure.

As a last step, we briefly review the robustness of our results to estimation design. One potential consideration is that comparisons across counties within a state-month are problematic due to county-level heterogeneity in responsiveness over time, and so within-county variation is better suited to our design. To address this, we re-estimate our results by replacing our state-month fixed effect with separate county and month fixed effects and reproduce our previous results in Appendix 3.9 as Figures 3.A3-3.A8. Our estimates are qualitatively quite similar: we again find reduced seasonality and a large drop in the delinquency-unemployment sensitivity during the pandemic, though the drops here are larger in percentage terms (Figure 3.A3). We also find a clear effect of the phase-out on the delinquency-unemployment sensitivity (Figure 3.A6), though now somewhat smaller than our previous state-month fixed effect estimates. Altogether, these differences lead to a qualitatively similar conclusion that

[22]Taken from Department of Labor official calculations of federal pandemic UI spending, available here.

UI policies instead explain about 60% of the total delinquency-unemployment sensitivity drop during the pandemic (Figure 3.A8). Though we prefer our prior estimates as better absorbing confounding state-level policies, we view this replication as broadly similar and reassuring evidence that our estimates are indeed quite robust.

3.7 Conclusion

In this paper, we use administrative credit bureau data to investigate the local financial effects of UI benefit expansions during the Covid-19 pandemic. At the micro level, if UI provides targeted liquidity to financially constrained households, then expansions should attenuate delinquency responses to unemployment. At the macro level, expanded UI represents large injections of liquidity into areas hit with adverse economic shocks and can be thought of as rapid counter-cyclical fiscal policy at the local level, targeted towards populations with potentially high marginal propensities to consume. Therefore, any micro stabilization might actually *understate* the effect of UI on aggregate economic conditions. We overcome this problem by directly estimating whether increasing the generosity of the UI system insulates aggregate financial conditions from economic shocks. We have three main findings. First, we estimate 50-66% reductions in the county-level delinquency-unemployment sensitivity after the introduction of Covid policies, driven both by changes in new delinquencies and continued nonpayment on existing delinquencies. Furthermore, this finding is qualitatively robust to placebo tests on unaffected credit types and demand-side responses.

At the same time, our first design cannot disentangle the effects of UI policies from other contemporaneous policies that would have also mitigated unemployment shocks. We thus next leverage the late 2021 staggered phase-out of federal UI to isolate the UI-specific component of the pandemic sensitivity drop. We estimate large sensitivity increases after UI withdrawal using a dynamic event study design, finding a 68-144% increase in sensitivity after 4 months (compared to the month before withdrawal). We find no evidence of pre-trends, supporting a casual interpretation of our results. As before, we again find that this result is robust to placebo tests and demand-side changes. Assuming that changes in the delinquency-unemployment sensitivity are symmetric with respect to UI expansions, our estimates imply that over 86% of our prior estimated Covid-era sensitivity drop is attributable to UI policies.

How should we think about these results in terms of delinquencies rather than sensitivities? In a last step, we assess the aggregate implications of our results and calculate the total delinquency-months prevented by UI policies. Using a simple framework to construct counterfactual delinquencies over the pandemic, we estimate that UI expansions prevented about 59% of potential delinquency-months for a cost of $8,864 per month. While this suggests that preventing any one delinquency month was very costly, these financial stabilization effects are in addition to the effects on micro-level household welfare and the aggregate *spending* effects. Ganong et al. (2022) estimate that the $600 and $300 dollar supplements boosted aggregate spending by 2.9% and 1.3%, respectively. Our estimates show that beyond the immediate effect of UI on aggregate demand, Covid-era expansions of UI also substantially stabilized aggregate financial conditions.

The Covid-19 pandemic saw unprecedented and prolonged increases in unemployment. Our results imply that UI policies were enormously successful in attenuating corresponding delinquency increases at the aggregate level. Ganong et al. (2022) show that the adverse labor market effects of UI expansion were small while the aggregate spending effects were large, a result mostly driven by the fact that substantial fraction of UI recipients seem to be high-MPC *types* rather than households with temporarily high MPCs because of liquidity constraints. This is consistent with our result that financial conditions became more sensitive to unemployment rates as soon as the UI expansions expired. Ganong et al. (2022) argue that their results suggest that front-loading of expanded benefits might be optimal policy in terms of trading off stimulating demand and increasing disincentives to work. Our results can be read as cautionary evidence that such front-loading may come at the cost of under-stabilizing financial conditions compared to smoother payout paths of UI supplements, presumably at levels that do not lead to median replacement rates substantially above 100%. An analysis of how to optimally trade off these two effects is a promising avenue for future research.

3.8 Figures and Tables

Figure 3.1: The Geography of Delinquency, Before and During the Pandemic

(a) Unemployment Rates, 2019

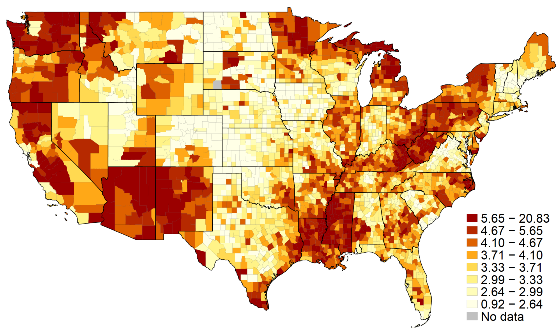

(b) All Delinquencies, 2019

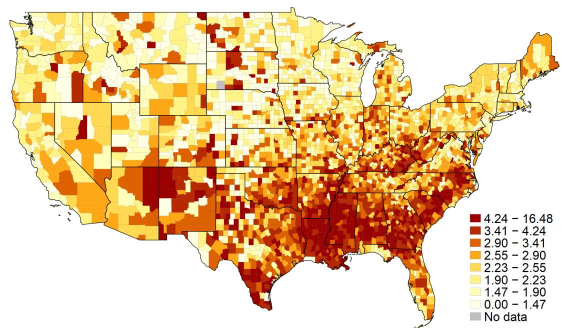

(c) Unemployment Rates, 2020

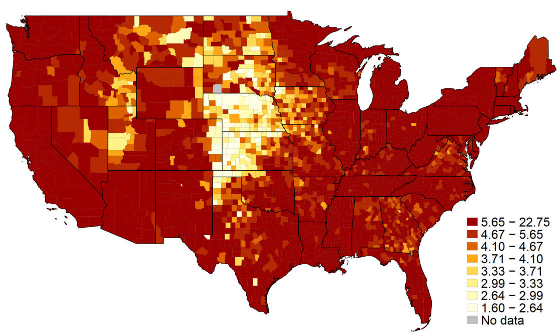

(d) All Delinquencies, 2020

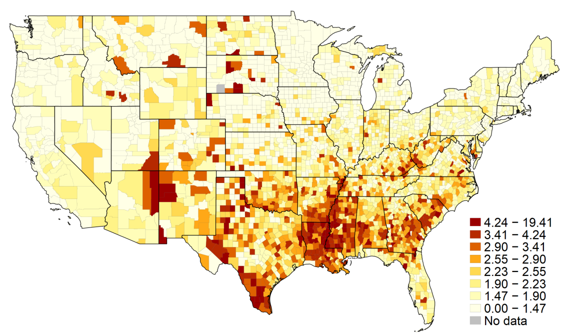

(e) Unemployment Rates, 2021

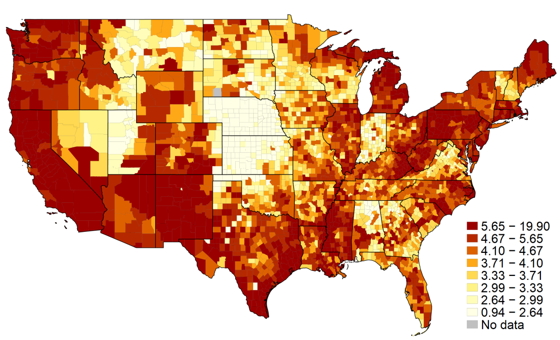

(f) All Delinquencies, 2021

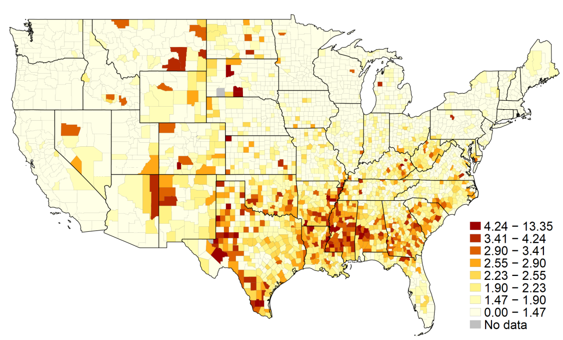

Notes: This figure graphs mean county-level delinquency and unemployment rates between 2019 and 2021. Shading for each measure represents 8 equally-spaced bins for 2019 values. Delinquency rates are constructed using our county-month aggregation of credit bureau microdata; more details on data construction can be found in Section 3.3.

Figure 3.2: The 2021 UI Phase-Out

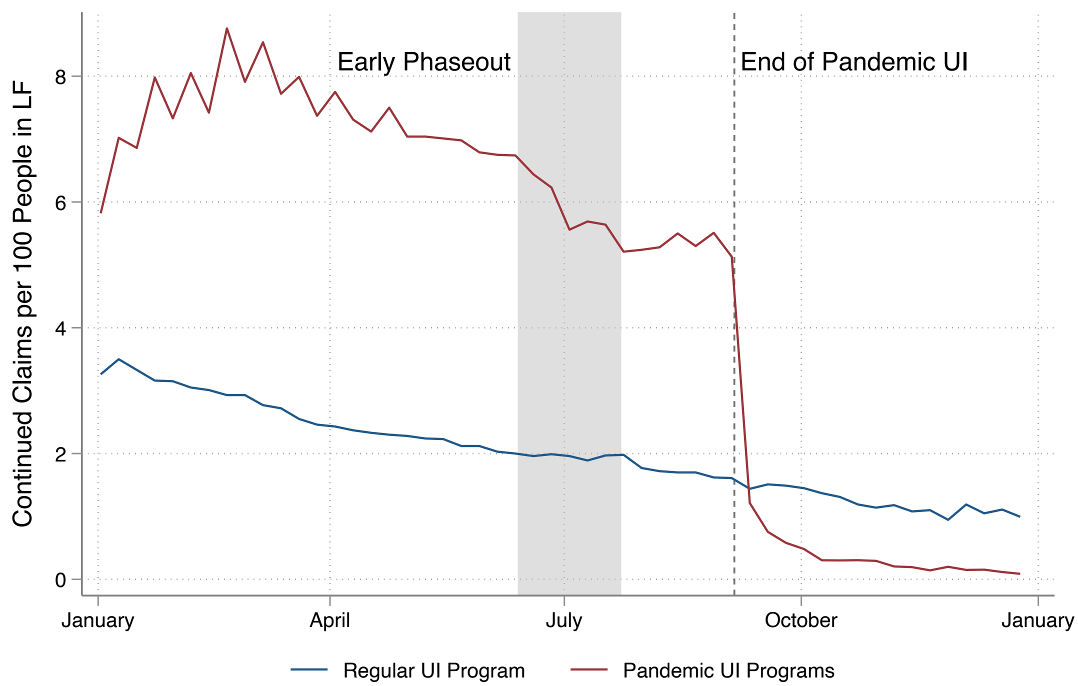

Notes: This figure plots continuing UI claims over time, for each week between 1/1/2021 and 1/1/2022, to highlight the stark drop in UI claimants after state-level withdrawals from federal pandemic UI programs. See Sections 3.2 and 3.5 for more details on the underlying policy variation. Our calculations are based on the Department of Labor's ETA 539 Weekly Claims data.

Figure 3.3: Federal UI Duration Extensions, 2000-2021

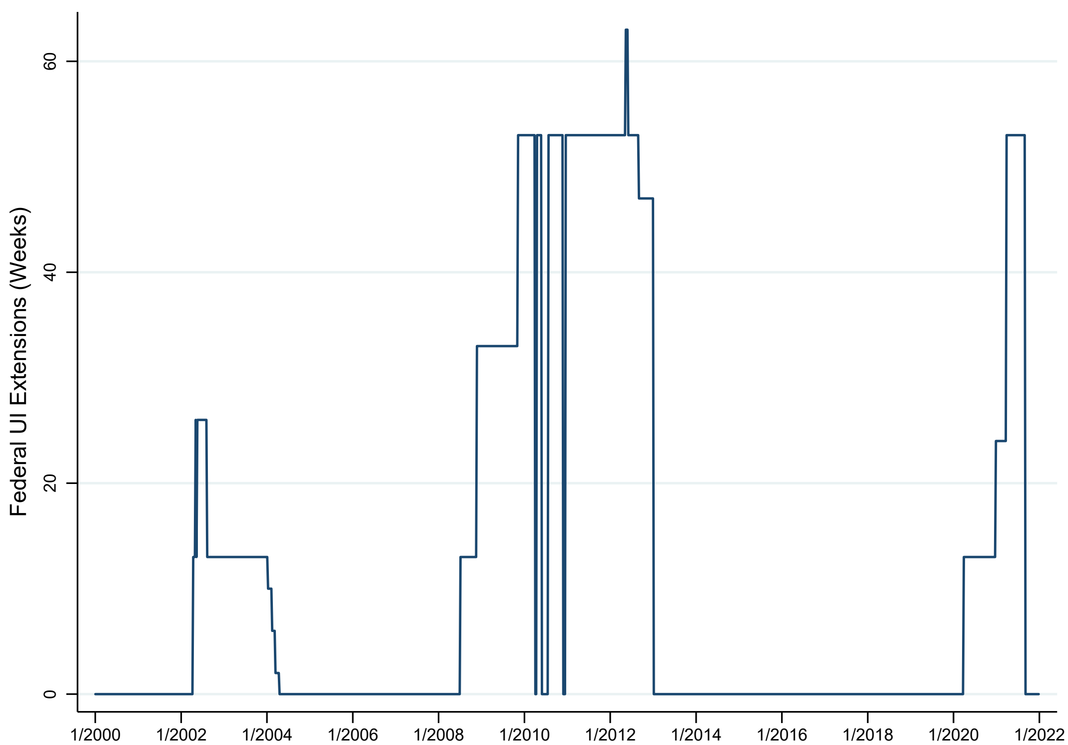

Notes: This figure plots the number of maximum total federal UI weeks available to new initial claimants for each week between January 2020 and December 2021. Importantly, this figure plots only federal weeks available to claimants: UI recipients could also access upto 48 total additional weeks from state-specific UI programs, depending on whether UI trigger policies were in effect.

Figure 3.4: Insured and Total Unemployment Rates, 2000-2021

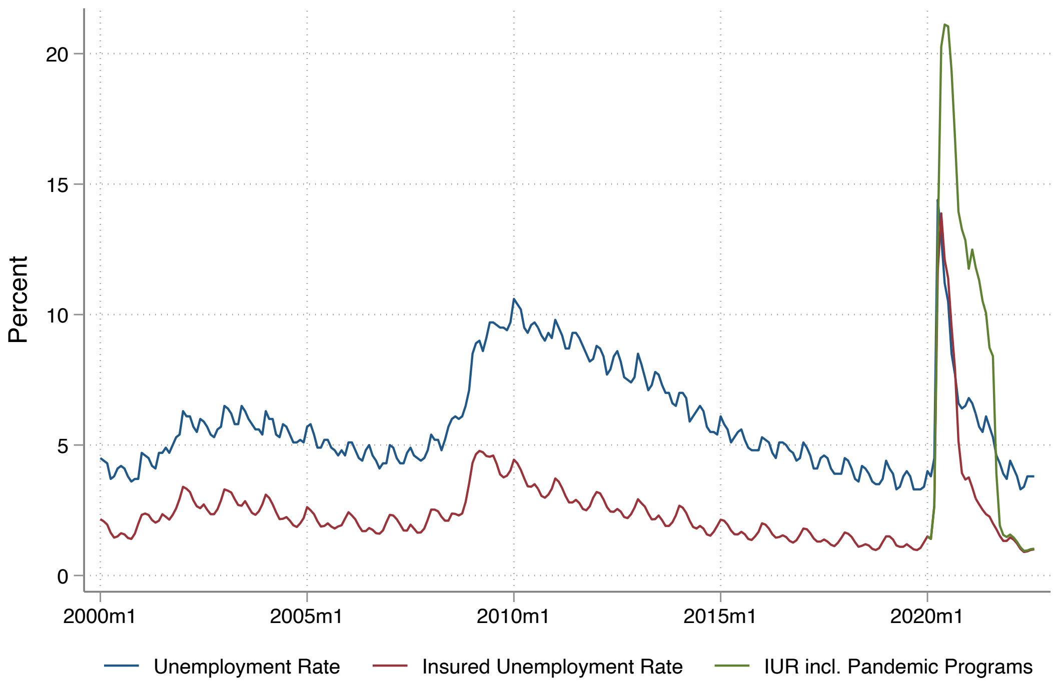

Notes: This figure plots the unemployment rate (UR), insured unemployment rate (IUR), and our constructed insured unemployment rate including pandemic programs. Our constructed series adjusts for the large expansions of UI eligibility during the pandemic through the federal PUA program. The first two series are from the Bureau of Labor Statistics. We construct the all-programs insured unemployment rate here by 1) computing the ratio of all-programs (regular UI, PEUC, PUA) and continuing claims weeks (which include both regular claims and special federal programs), 2) multiplying by the regular IUR.

Figure 3.5: Delinquencies vs Unemployment, Before and During Covid

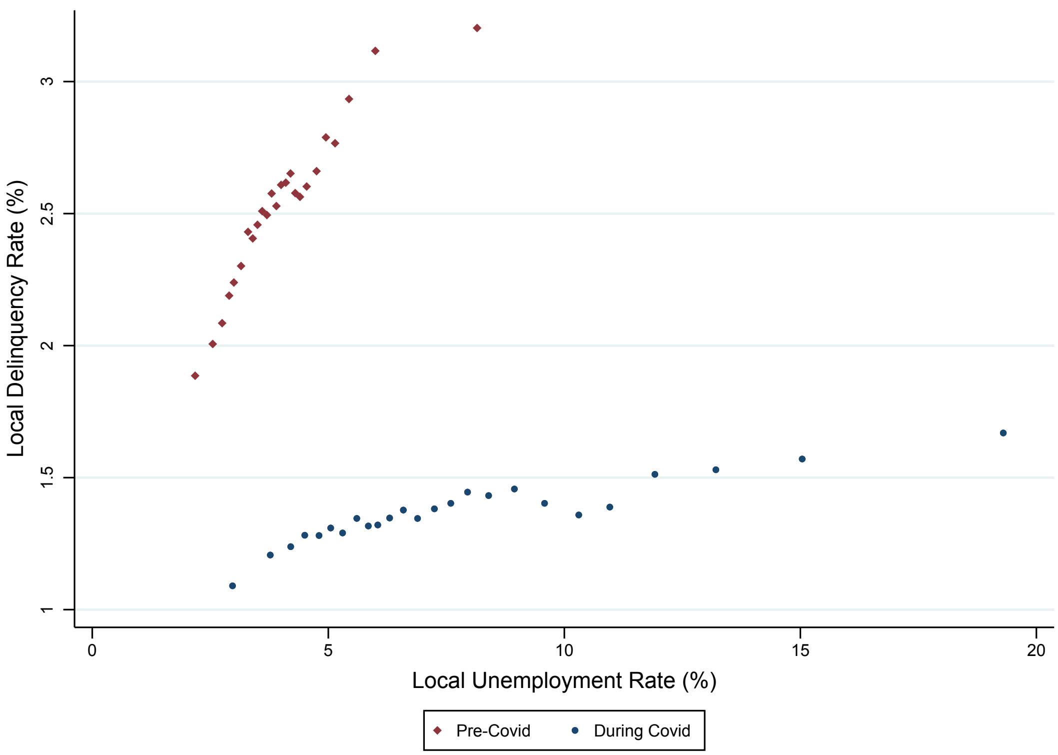

Notes: This figure displays a large attenuation in local responsiveness to unemployment rate shocks following the introduction of Covid policies. We perform a binned scatterplot of county-level any-loan delinquency rates against county-level unemployment rates, separately using county-months from January 2018 to February 2020 (red) and again using March 2020 to August 2021 (blue). Delinquency rates are constructed using our county-month aggregation of credit bureau microdata, and county unemployment rates are taken from the LAUS. More details on data and interpretation can be found in Sections 3.3 and 3.4 respectively.

Figure 3.6: Delinquency-Unemployment Sensitivity Over Time

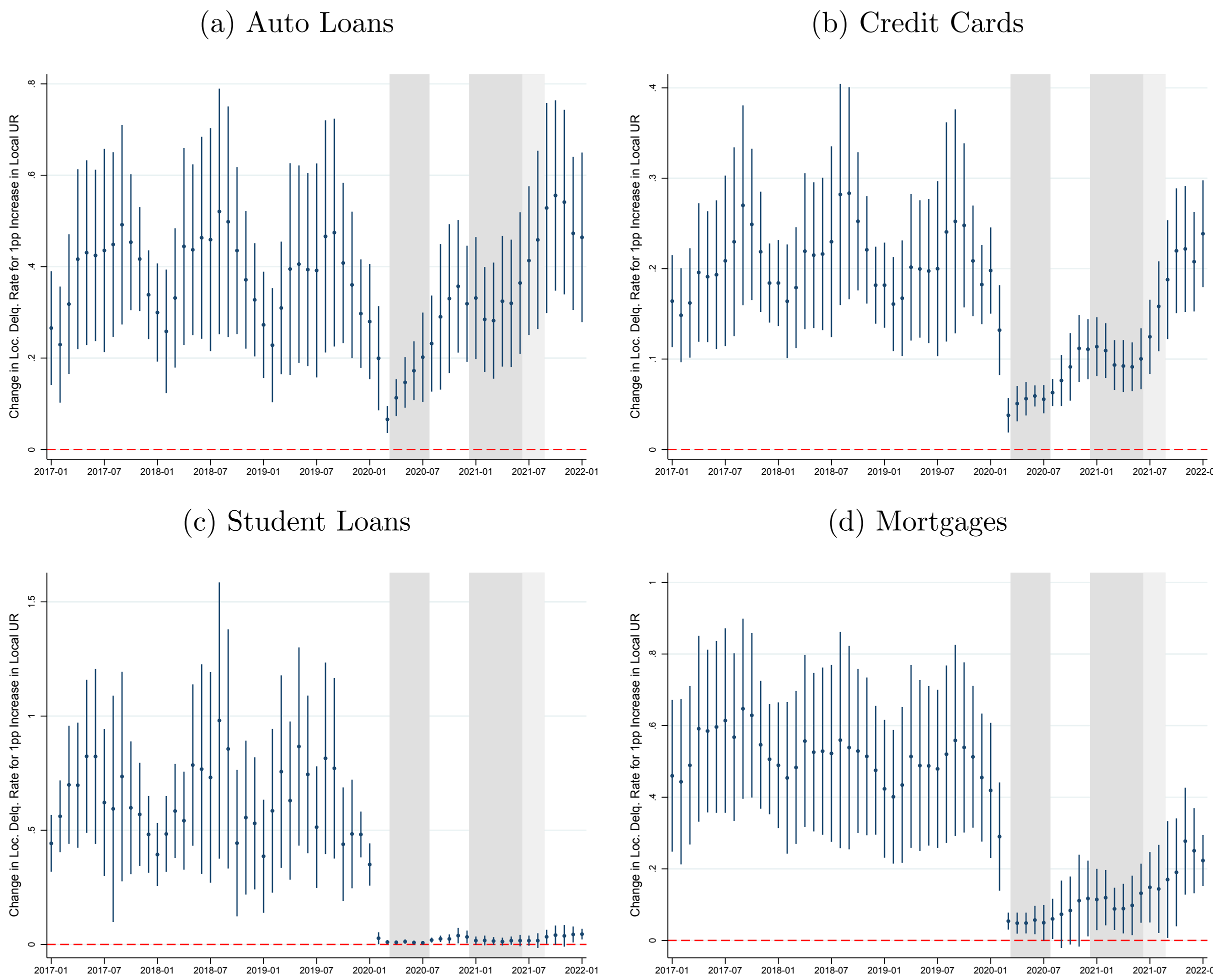

Notes: This figure shows the evolution of the estimated delinquency-unemployment sensitivity for each month between January 2017 and March 2022, separately for different credit types. Each panel is a separate regression, plotting coefficients $\{\beta_t\}_{\tau}^{T}$ from our estimation of Equation 3.1. Delinquency rates are constructed using our county-month aggregation of credit bureau microdata, and county unemployment rates are taken from the LAUS. More details on data and interpretation can be found in Sections 3.3 and 3.4 respectively.

Figure 3.7: Short-Term Delinquency-Unemployment Sensitivity Over Time

(a) Auto Loans, 30-89 Day Delinquencies

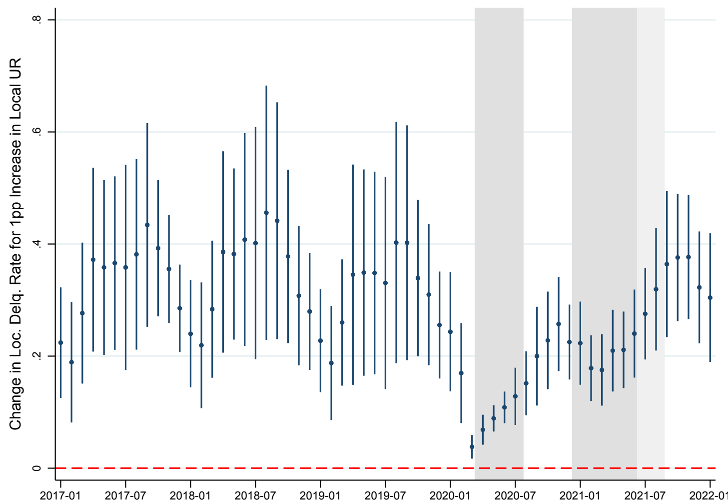

(b) Credit Cards, 30-89 Day Delinquencies

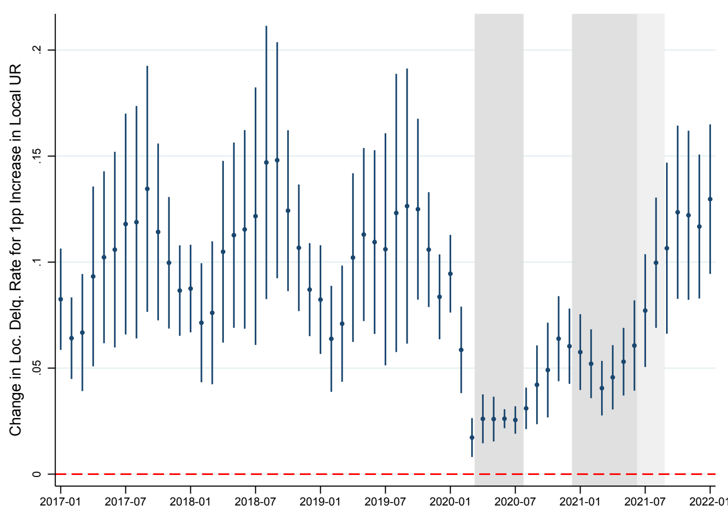

Notes: This figure shows the evolution of the estimated delinquency-unemployment sensitivity for each month between January 2017 and March 2022, separately for different credit types. In comparison to the previous figure, here we use the short-term 30-89 day delinquency rate as the dependent variable. Each panel is a separate regression, plotting coefficients $\{\beta_t\}_{\tau}^{T}$ from our estimation of Equation 3.1. Delinquency rates are constructed using our county-month aggregation of credit bureau microdata, and county unemployment rates are taken from the LAUS. More details on data and interpretation can be found in Sections 3.3 and 3.4 respectively.

Figure 3.8: New Loan Responses to Local Unemployment Shocks

(a) New Auto Loans

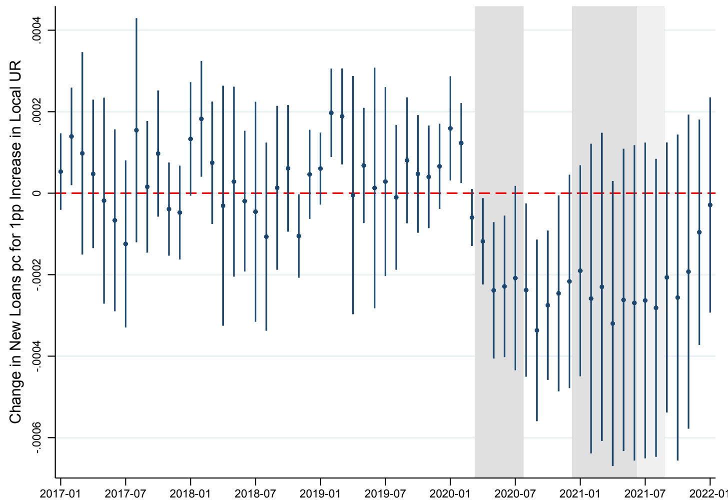

(b) New Credit Cards

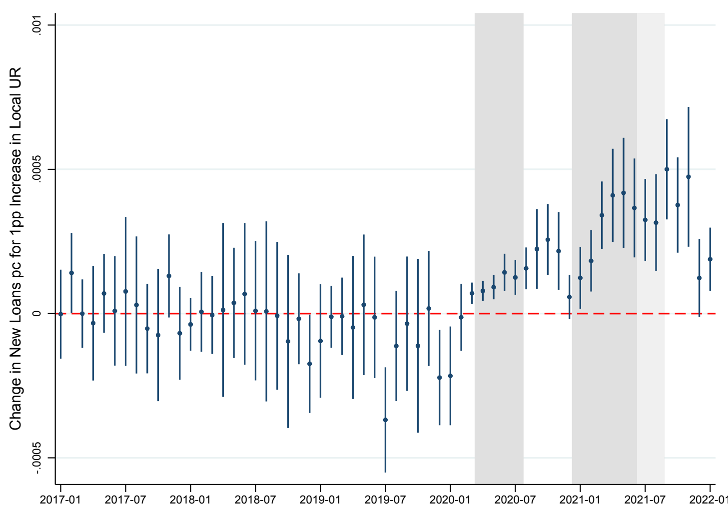

Notes: This figure assesses potential demand-side responses to local unemployment shocks, separately for each month between January 2017 and March 2022. In comparison to the previous figure, here we use the change in the number of per-capita loans (disaggregating into auto loans and credit cards) as the dependent variable. Each panel is a separate regression, plotting coefficients $\{\beta_t\}_{\tau}^{T}$ from our estimation of Equation 3.1. New loans for each credit type are constructed using our county-month aggregation of credit bureau microdata, and county unemployment rates are taken from the LAUS. More details on data and interpretation can be found in Sections 3.3 and 3.4 respectively.

Figure 3.9: Effects of UI Phase-Out on Delinquency-Unemployment Sensitivity

(a) Auto Loans

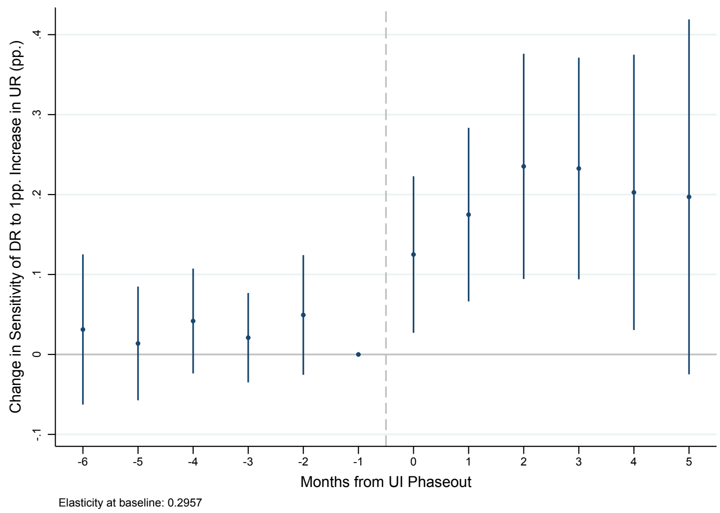

(b) Credit Cards

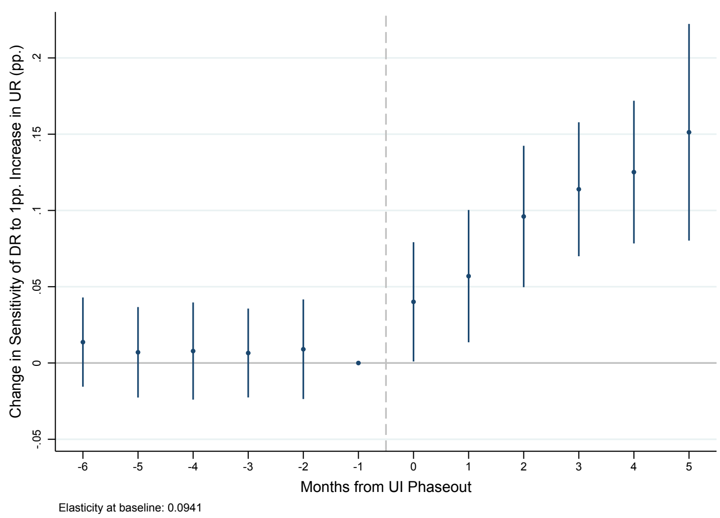

Notes: This figure assesses the impacts of federal UI withdrawals using a staggered event study design, leveraging the fact that different states withdrew at different times. In comparison to the previous monthly sensitivity graphs, here we estimate the effect of withdrawal on the delinquency-unemployment sensitivity (normalized to 0 in the period before withdrawal). More details on the estimation procedure and interpretation can be found in Section 3.5.

Figure 3.10: Effects of UI Phase-Out on Short-Term Delinquency-Unemployment Sensitivity

(a) Auto Loans, 30-89 Day Delinquencies

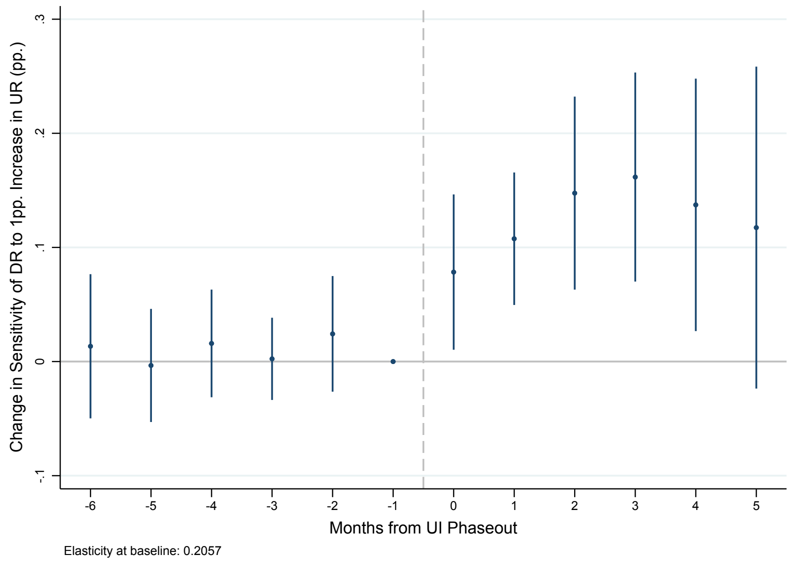

(b) Credit Cards, 30-89 Day Delinquencies

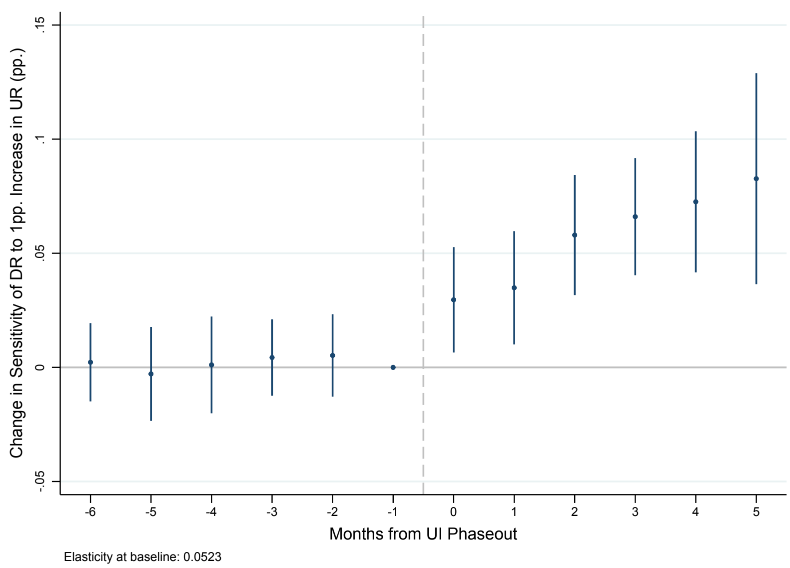

Notes: This figure assesses the impacts of federal UI withdrawals using a staggered event study design. In comparison to the previous monthly sensitivity graphs, here we estimate the effect of withdrawal on short-term delinquency-unemployment sensitivity (normalized to 0 in the period before withdrawal). More details on the estimation procedure and interpretation can be found in Section 3.5.

Figure 3.11: UI Phase-Out: New Loan Responses

(a) Auto Loans

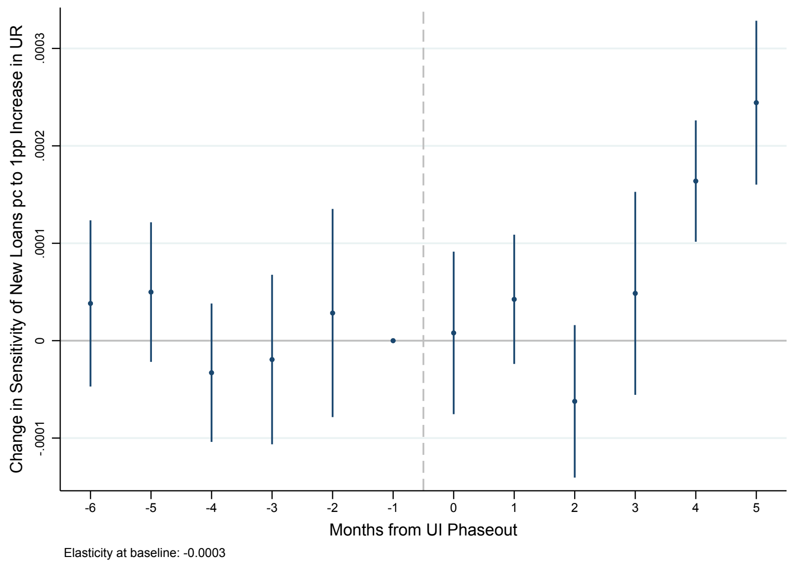

(b) Credit Cards

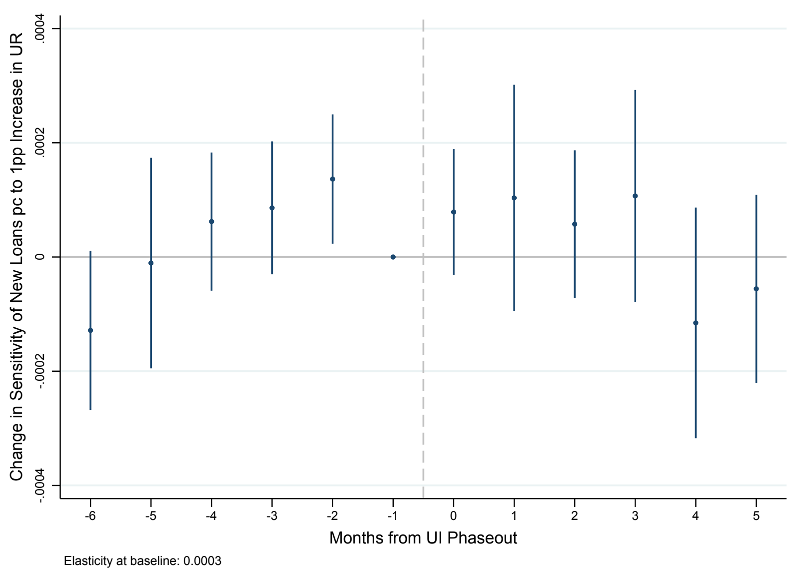

Notes: This figure assesses potential demand-side responses to state-level UI withdrawals. The outcome variable is the change in per capita new loans per percentage point change in the unemployment rate, relative to the period before withdrawal (-1, normalized to 0). More details on the estimation procedure and interpretation can be found in Section 3.5.

Figure 3.12: Construction of Counterfactual

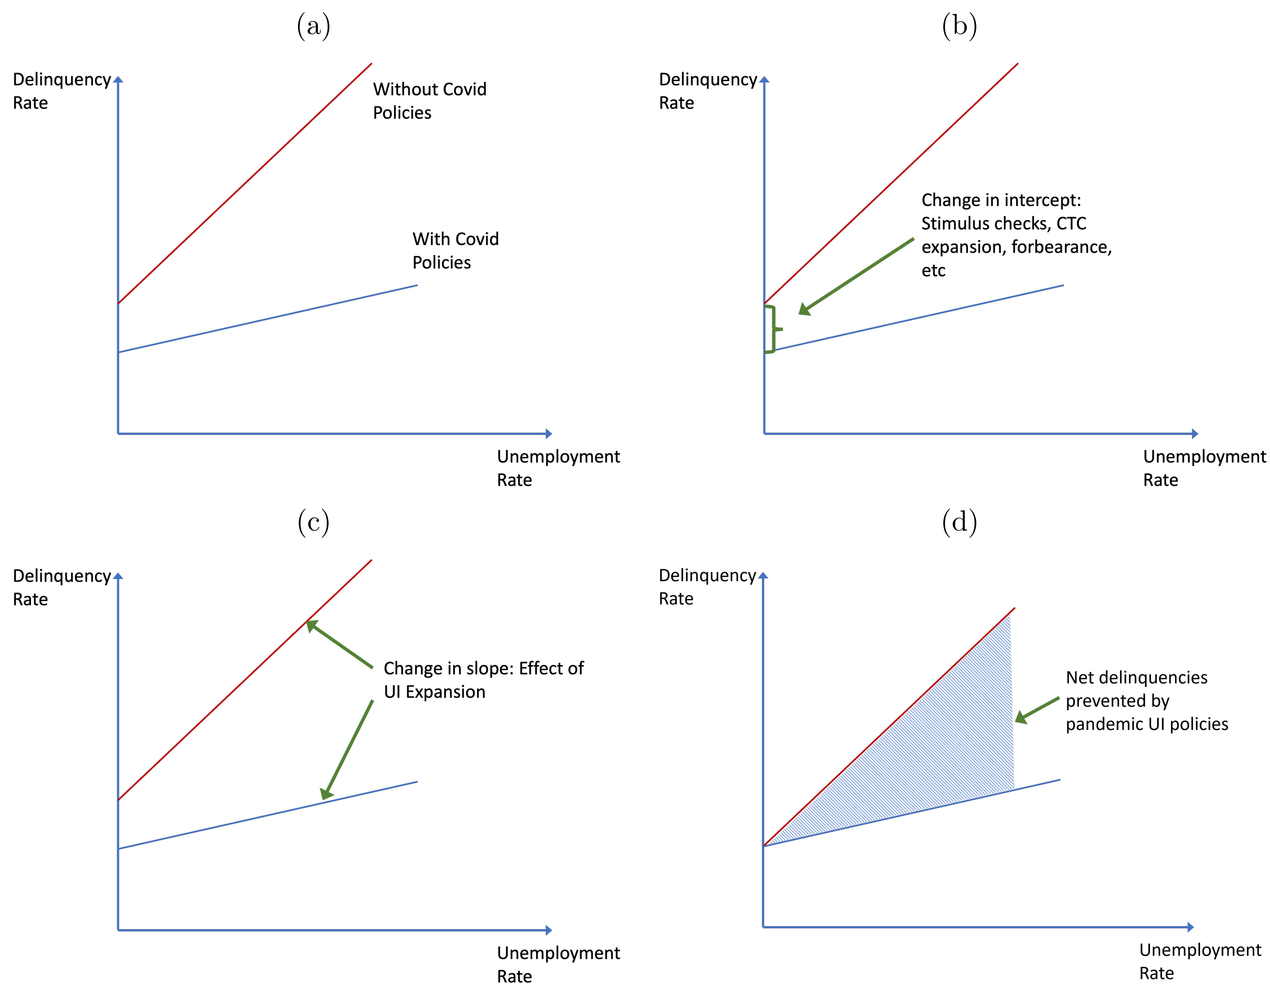

Notes: This figure presents a simplified visual aid to guide intuition for our counterfactual estimation procedure. Panel (a) starts with a reproduction of Figure 3.5, the empirical delinquency-unemployment relationship before and during the pandemic. As represented in panels (b) and (c), under a simplified attenuation framework UI policies can only impact delinquencies through a change in the curves' slope. Differences in intercepts thus reflect other existing Covid policies, such as stimulus checks or CTC expansion. Panel (d) illustrates our proposed calculation for aggregate delinquency effects: after removing intercept differences, the distance between the pre-Covid and during-Covid curves represent the prevented delinquencies at each value of the unemployment rate. We can thus sum across unemployment rates to yield the total number of delinquencies prevented. An expanded discussion of this figure and corresponding results can be found in Section 3.6.

Figure 3.13: Counterfactual Estimates, State-Month FE

(a) Auto Loans

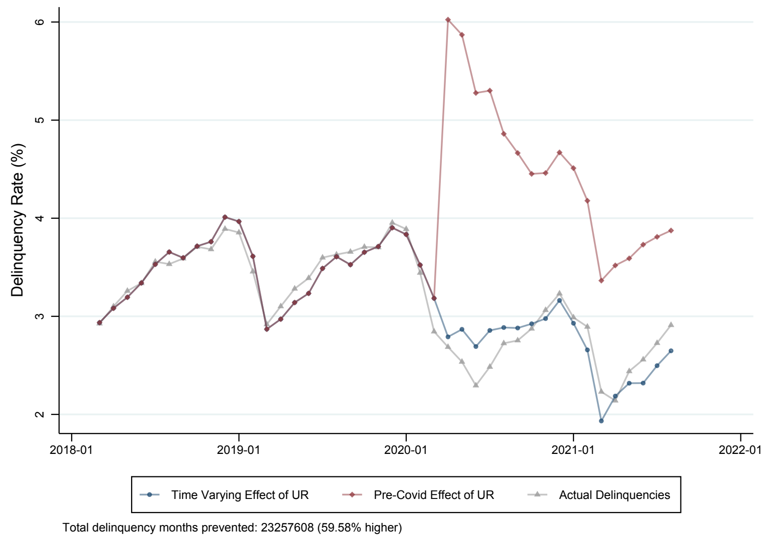

(b) Credit Cards

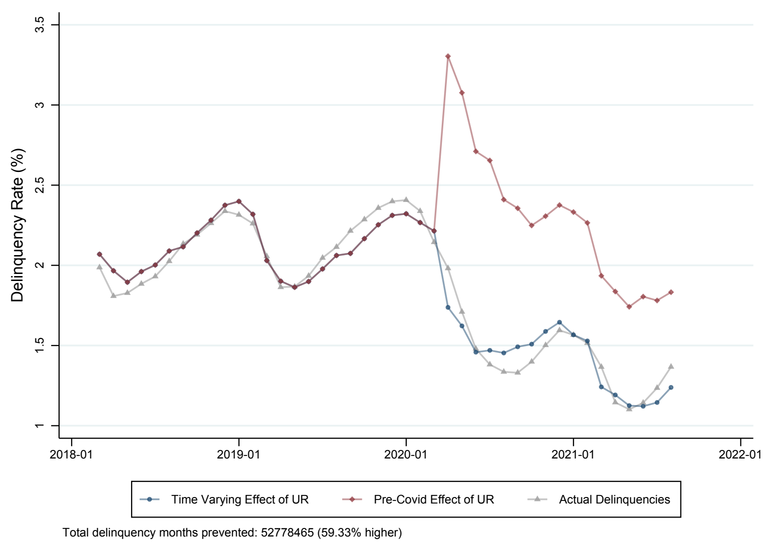

Notes: This figure presents empirical, predicted, and counterfactual delinquency time series. The blue line plots fitted values from Equation 3.3 as an estimated evolution of the delinquency rate. As a reassuring check, this series roughly matches the dynamics of actual observed delinquencies over time (grey line). The red line, however, instead plots fitted values where the level and shift effects are removed from the blue line between March 2020 and August 2021. This second series thus represents a designed counterfactual where we have removed the effect of federal Covid policies. We can then calculate the number of monthly prevented delinquency-months as the difference between our estimated counterfactual (red) and estimated status quo (blue) series for each month, multiplied by the number of loans for that credit type in our data. We sum this measure over the Covid UI period to arrive at a total sum for delinquency-months prevented.

Table 3.1: Distribution of Key County-Level Variables

Variable	**Mean**	**p5**	**p10**	**p25**	**p50**	**p75**	**p90**	**p95**	**p99**
Labor Force									
2019	52,181	1,463	2,393	5,016	11,824	32,051	106,554	229,631	699,903
Pct Change, 2019-2020	.15	-3.2	-2.2	-.97	.12	1.2	2.5	3.3	6.7
Number of People									
2019	1331	37	63	140	329	873	2737	5761	16270
Pct Change, 2019-2020	-.9	-6.4	-4.9	-2.7	-.92	.7	2.7	4.6	9.2
Unemployment Rate									
2019	4.7	2.4	2.7	3.3	4.3	5.6	7.3	8.4	11
Change, 2019-2020	-.48	-1.2	-.98	-.72	-.45	-.22	-.025	.12	.49
DQ Share: Any Loan									
All Term: 2019	2.3	.73	.97	1.4	2.1	2.9	3.9	4.7	6.7
All Term: Change, 2019-2020	.075	-1.3	-.83	-.28	.082	.45	.98	1.4	2.5
Short Term: 2019	1.2	.4	.53	.77	1.1	1.5	2	2.3	3.2
Short Term: Change, 2019-2020	.075	-1.3	-.83	-.28	.082	.45	.98	1.4	2.5
DQ Share: Auto Loan									
All Term: 2019	3.4	.21	.94	1.9	2.9	4.4	6.2	7.8	12
All Term: Change, 2019-2020	-.069	-2.6	-1.6	-.64	-.019	.56	1.4	2.3	4.5
Short Term: 2019	2.7	0	.7	1.5	2.4	3.6	5.1	6.3	9.8
Short Term: Change, 2019-2020	-.069	-2.6	-1.6	-.64	-.019	.56	1.4	2.3	4.5
DQ Share: CC									
All Term: 2019	2	.44	.78	1.3	1.8	2.5	3.3	4	5.8
All Term: Change, 2019-2020	.15	-1.4	-.88	-.27	.13	.56	1.2	1.7	3.6
Short Term: 2019	1.1	.25	.44	.71	.99	1.3	1.7	2.1	3.1
Short Term: Change, 2019-2020	.15	-1.4	-.88	-.27	.13	.56	1.2	1.7	3.6
Number of Counties: 3,107									

Notes: This table displays summary statistics for the balanced panel of counties in our analysis sample. The labor force size and unemployment rate are taken from the LAUS; person counts and delinquency shares are taken from our county-month aggregation of credit bureau microdata. See Section 3.3 for more information on the underlying data construction.

3.9 Appendix

Figure 3.A1: Comparison to CFPB Mortgage Delinquency Data

(a) Under 90 Days Delinquent (b) Over 90 Days Delinquent

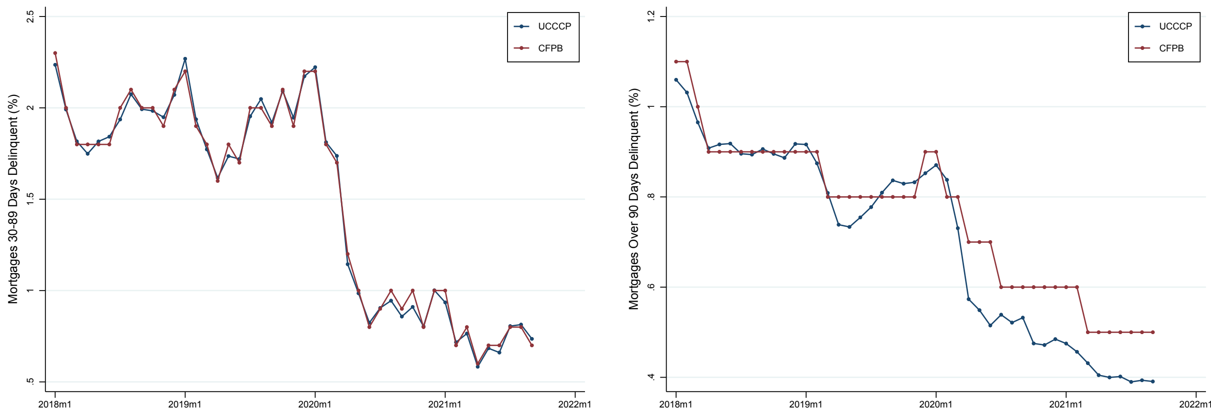

Notes: Our main analysis sample is a monthly reconstruction—using retrospective monthly loan payment status identifiers—of quarterly credit bureau archives. To validate our constructed data, we compare average mortgage delinquency rates in our microdata to public aggregates from the Consumer Financial Protection Bureau (other credit types are not available from the CFPB for a similar analysis). Small differences in Panel (b) are partially attributable to the CFPB's rounding of delinquency rates.

Figure 3.A2: Distribution of Observed County Sizes

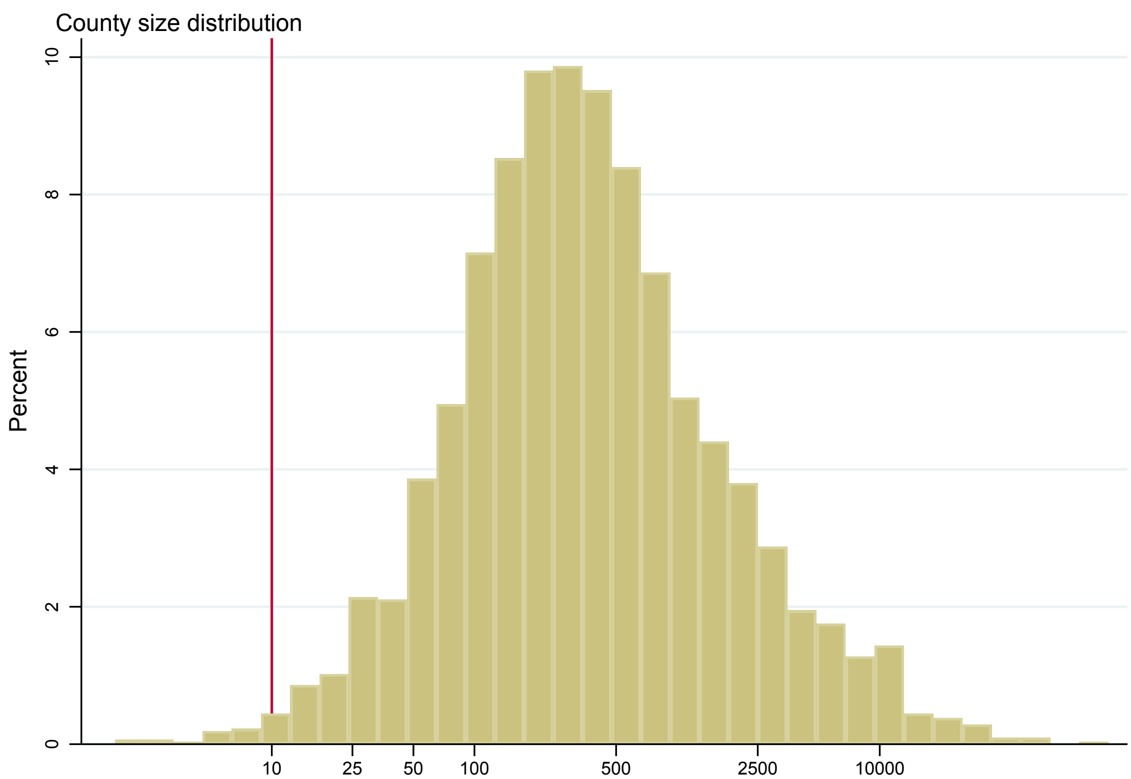

Notes: This figure shows the distribution of average county sizes (between 2017 and 2021) in our aggregated county-month analysis sample, before imposing a county size restriction. The horizontal axis is displayed in log scale (with corresponding level tick values). For our main analysis sample, we drop counties at the far left tail with less than 10 observed people on average between 2017 and 2021 (red line).

Figure 3.A3: Delinquency-Unemployment Sensitivity Over Time

(a) Auto Loans

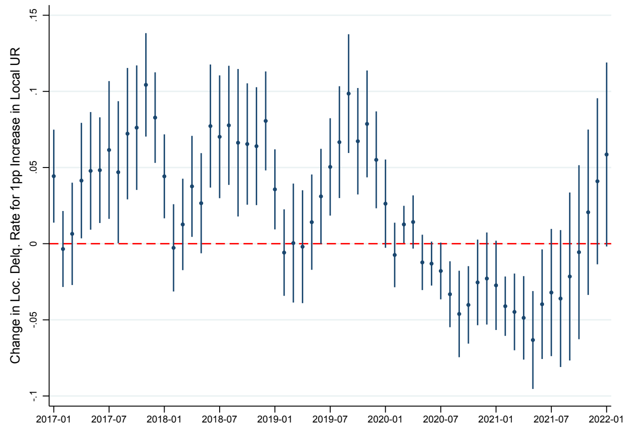

(b) Credit Cards

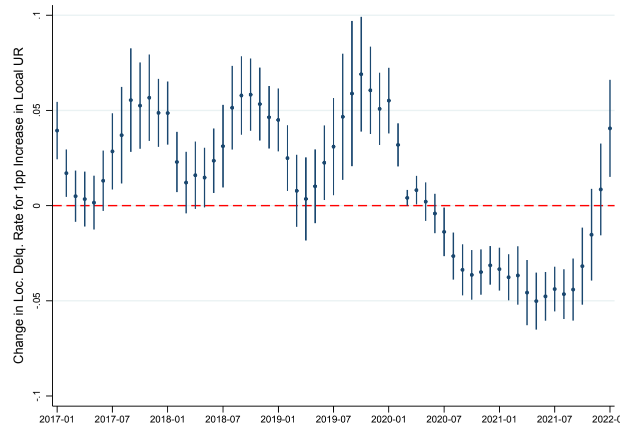

(c) Student Loans

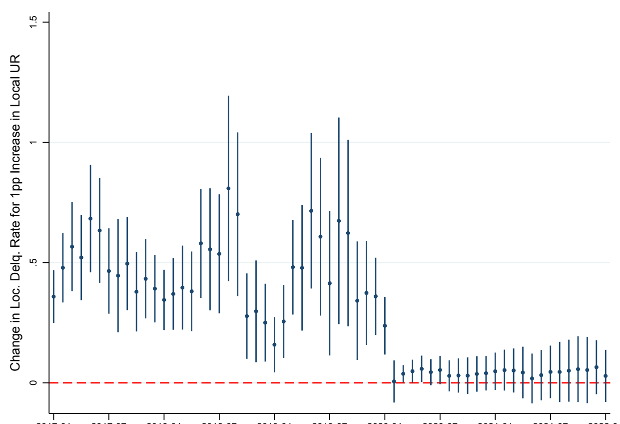

(d) Mortgages

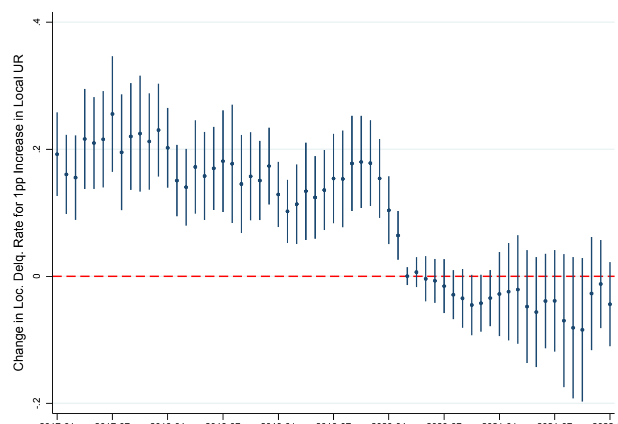

Notes: This figure shows the evolution of the estimated delinquency-unemployment sensitivity for each month between January 2017 and March 2022, separately for different credit types. Each panel is a separate regression, plotting coefficients $\{\beta_t\}_{\tau}^{T}$ from a version of Equation 3.1 that replaces the state-month fixed effect with separate county and month fixed effects. Delinquency rates are constructed using our county-month aggregation of credit bureau microdata, and county unemployment rates are taken from the LAUS. More details on data and interpretation can be found in Sections 3.3 and 3.4 respectively.

Figure 3.A4: Short-Term Delinquency-Unemployment Sensitivity Over Time

(a) Auto Loans, 30-89 Day Delinquencies

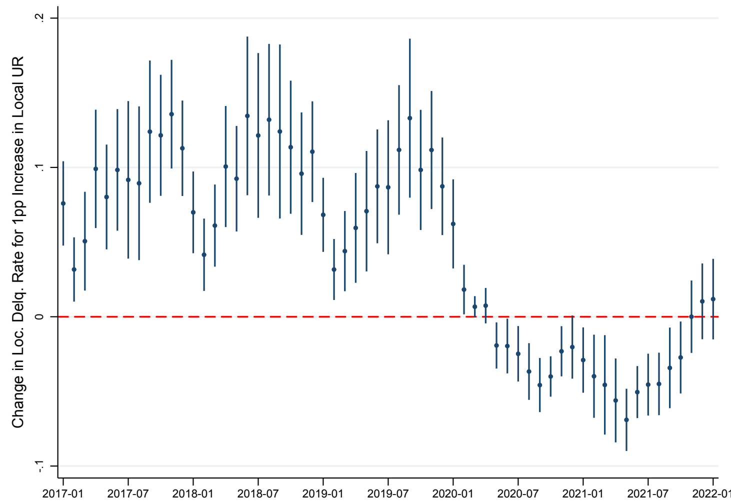

(b) Credit Cards, 30-89 Day Delinquencies

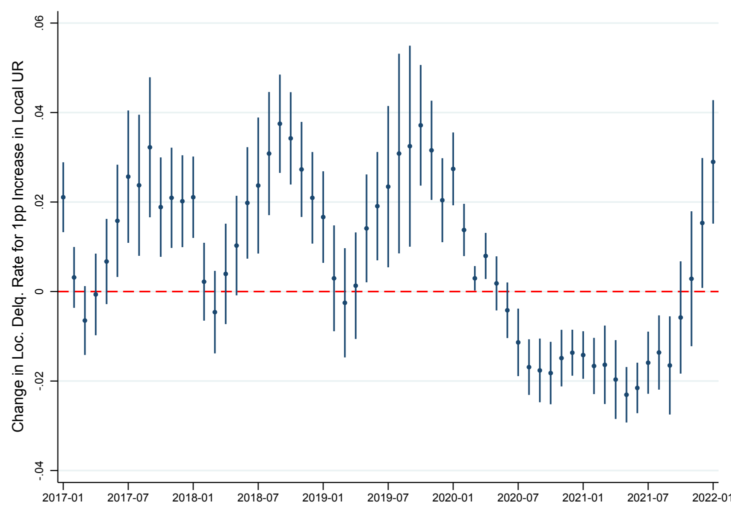

Notes: This figure shows the evolution of the estimated delinquency-unemployment sensitivity for each month between January 2017 and March 2022, separately for different credit types, now using the short-term 30-89 day delinquency rate as the dependent variable. Each panel is a separate regression, plotting coefficients $\{\beta_t\}_{\tau}^{T}$ from a version of Equation 3.1 that replaces the state-month fixed effect with separate county and month fixed effects. Delinquency rates are constructed using our county-month aggregation of credit bureau microdata, and county unemployment rates are taken from the LAUS. More details on data and interpretation can be found in Sections 3.3 and 3.4 respectively.

Figure 3.A5: New Loans Responses to Local Unemployment Shocks

(a) New Auto Loans

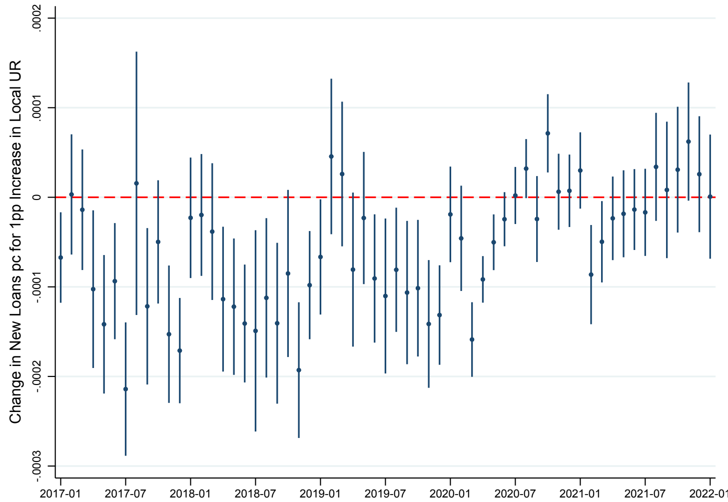

(b) New Credit Cards

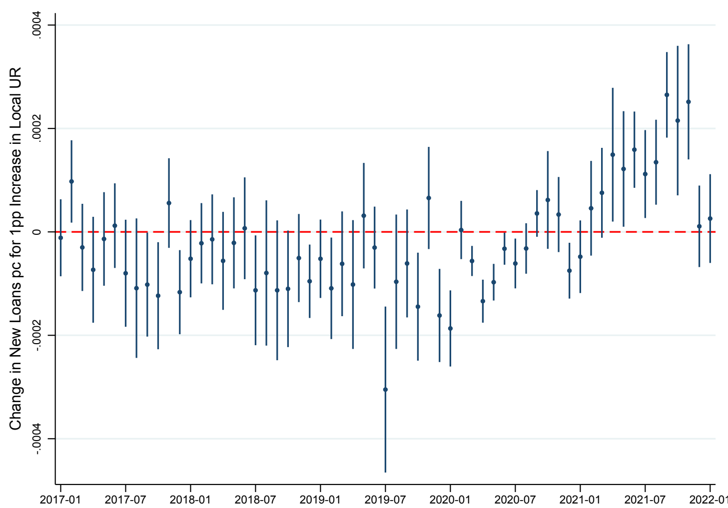

Notes: This figure assesses potential demand-side responses to local unemployment shocks, separately for each month between January 2017 and March 2022. In comparison to the previous figure, here we use the change in the number of per-capita loans (disaggregating into auto loans and credit cards) as the dependent variable. Each panel is a separate regression, plotting coefficients $\{\beta_t\}_{\tau}^{T}$ from a version of Equation 3.1 that replaces the state-month fixed effect with separate county and month fixed effects. New loans for each credit type are constructed using our county-month aggregation of credit bureau microdata, and county unemployment rates are taken from the LAUS. More details on data and interpretation can be found in Sections 3.3 and 3.4 respectively.

Figure 3.A6: Effects of UI Phase-Out on Delinquency-Unemployment Sensitivity

(a) Auto Loans

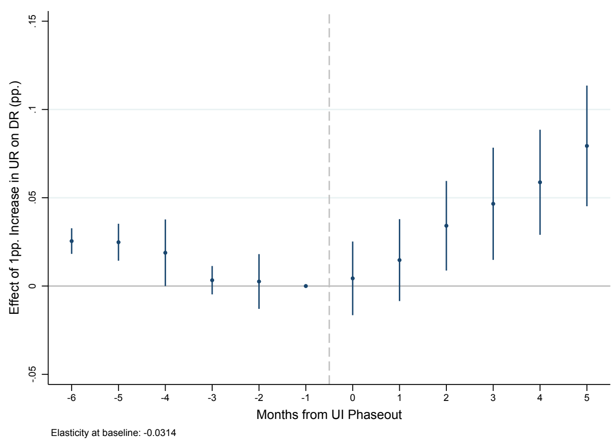

(b) Credit Cards

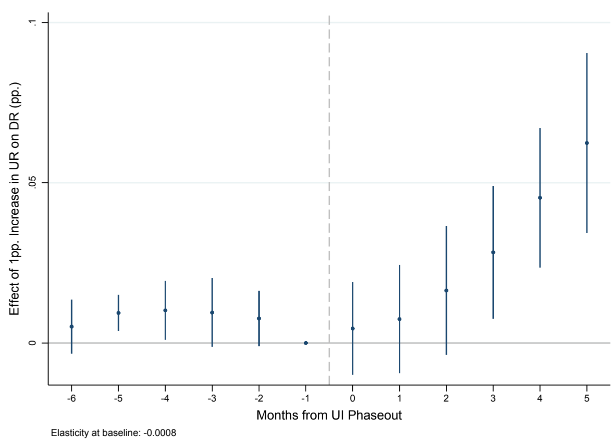

Notes: This figure assesses potential demand-side responses to local unemployment shocks, separately for each month between January 2017 and March 2022. In comparison to the previous figure, here we use the change in the number of per-capita loans (disaggregating into auto loans and credit cards) as the dependent variable. Each panel is a separate regression, plotting coefficients $\{\beta_t\}_{\tau}^{T}$ from estimating a version of Equation 3.1 that replaces the state-month fixed effect with separate county and month fixed effects. New loans for each credit type are constructed using our county-month aggregation of credit bureau microdata, and county unemployment rates are taken from the LAUS. More details on data and interpretation can be found in Sections 3.3 and 3.4 respectively.

Figure 3.A7: Effects of UI Phase-Out on Short-Term Delinquency-Unemployment Sensitivity

(a) Auto Loans, 30-89 Day Delinquencies

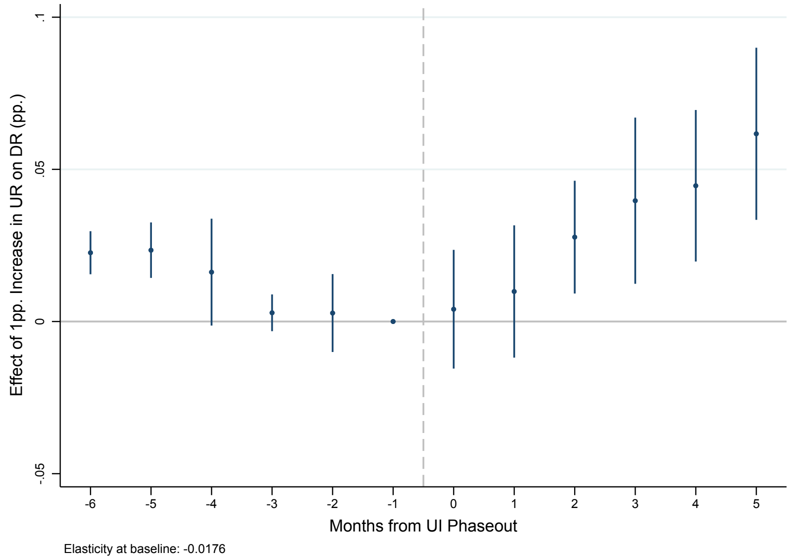

(b) Credit Cards, 30-89 Day Delinquencies

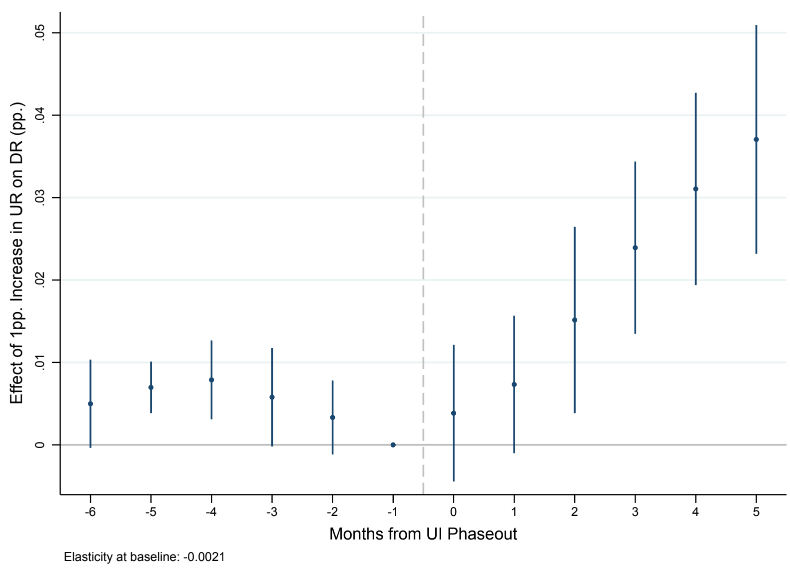

Notes: This figure assesses the impacts of federal UI withdrawals using a staggered event study design. In comparison to the previous monthly sensitivity graphs, here we estimate the effect of withdrawal on the short-term delinquency-unemployment sensitivity (normalized to 0 in the period before withdrawal). More details on the estimation procedure and interpretation can be found in Section 3.5.

Figure 3.A8: Counterfactual Estimates

(a) Auto Loans

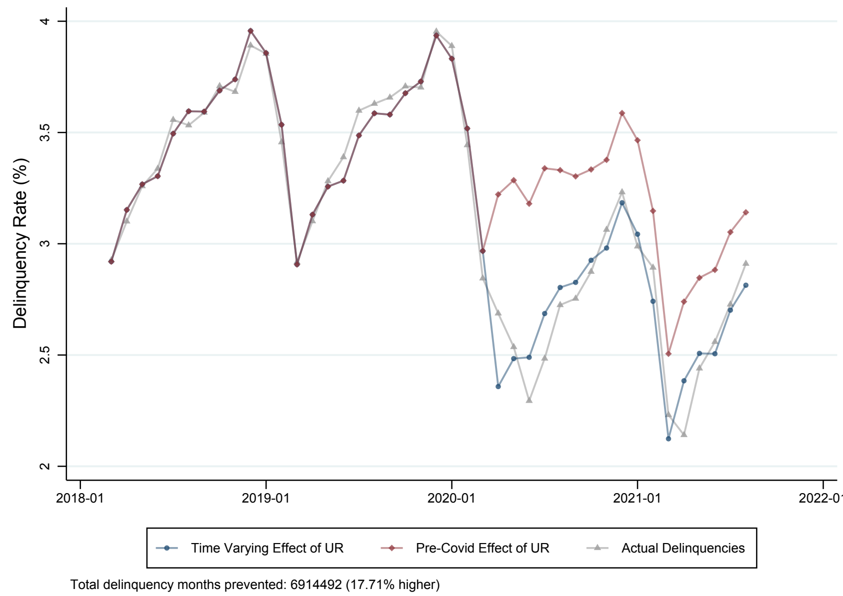

(b) Credit Cards

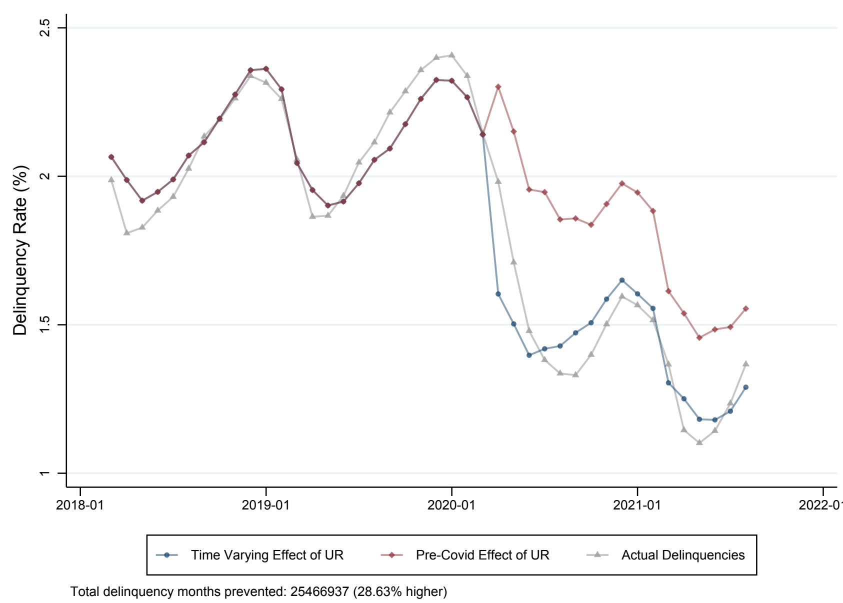

Notes: This figure presents empirical, predicted, and counterfactual delinquency rate time series during the pandemic, reproducing Figure 3.13 by replacing the state-month fixed effect with separate county and month fixed effects.

Chapter 4

Conclusion and Outlook

The Euler equation lies at the heart of modern macroeconomics. In many cases, the efficacy of monetary and fiscal policy depends on the fraction of households with large MPCs out of transitory changes in income and their sensitivity to interest rates. My findings suggest that the empirical consumption behavior of households is not fully captured by a standard Euler equation. In fact, even sophisticated models of consumption like the buffer-stock savings model or a model with uninsurable income risk and heterogeneous asset returns is unable to explain at least some consumption behavior. This suggests that we need to think about further modifications to state-of-the-art consumption models. While behavioral economists have often suggested that a model of heterogeneous discounting with some present-focused and some time-consistent households is very successful in explaining both consumption and job search behaviors, my overall findings suggest that even these models may not fully capture how households make consumption and savings choices. Specifically, a very valuable path for future research is to write down models of consumption that allow for mental accounting so that households have consumption rules that place some weight on current income by itself (not just through its effect on permanent income).

In Chapter 2, I find that households not only cut their expenditure at the onset of retirement, they also cut their *consumption*. Declines in expenditure are neither fully nor mostly explained by declines in household-level prices or changes in the composition of the consumption basket. Strikingly, even households with substantial savings cut their consumption and the size of the income drop is highly predictive of the magnitude of the consumption drop. This latter finding is hard to reconcile even with models of consumption with heterogeneous discount factors and points toward mental accounting as a potential mechanism in household consumption behavior.

In Chapter 3, I find that UI system expansions that led to unemployment insurance benefits often exceeding 100% of prior income resulted in local financial conditions being almost fully insulated from local economic conditions. While county-level delinquency rates are highly sensitive to county-level unemployment rates in "normal" times, the sensitivity of delinquencies to local unemployment rates was greatly attenuated while Covid-era UI expansions were in effect. Levering the politically motivated staggered phase-out of expanded

UI, we find that county-level delinquency rates become sensity of county-level unemployment rates precisely after the increases in UI generosity end. This suggests that liquidity–or current income–plays an important role in explaining delinquency dynamics, a finding consistent with the results in Chapter 2 where consumption is highly sensitive to current income.

These findings suggest many interesting avenues for future research. In investigating consumption smoothing behavior, researchers may want to distinguish consumption responses to four different kinds of shocks: Unanticipated temporary shocks, anticipated temporary shocks, unanticipated permanent shocks, and anticipated permanent shocks. In this dissertation, I have focused on retirement (an anticipated shock that permanently lowers income flows) and unemployment (an unanticipated shock that mostly reduces current income while leaving permanent income unchanged). In the existing literature, a lot of attention has been devoted to anticipated temporary shocks (e.g., tax rebates, tax refunds, payments from the Alaska permanent fund, or the exhaustion of unemployment benefits) and unanticipated temporary shocks (e.g., job loss or lottery winnings). However, there are many questions yet to be explored. For example, many seasonal workers face highly predictable and recurring transitory income shocks: Their income is high during the high season and low during the low season. Another potentially interesting setting is fully antipated future *increases* in income. For example, medical and law students (and students more generally) face the prospect of permanently higher incomes once they graduate. To my knowledge, researchers have yet to explore consumption behavior around these anticipated increases in income.

Bibliography

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

Aguiar, M. and Hurst, E. (2007). Life-Cycle Prices and Production. *American Economic Review*, 97(5):39.

Baily, M. N. (1978). Some Aspects of Optimal Unemployment Insurance. *Journal of Public Economics*, 10(3):379–402.

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

Bernheim, B. D. (1987). Dissaving after Retirement: Testing the Pure Life Cycle Hypothesis. In *NBER Chapters*, pages 237–280. National Bureau of Economic Research, Inc.

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.

Blanchet, T., Saez, E., and Zucman, G. (2022). Real-Time Inequality. Working Paper 30229, National Bureau of Economic Research.

Blank, R. M. and Card, D. E. (1991). Recent Trends in Insured and Uninsured Unemployment: Is There An Explanation? *The Quarterly Journal of Economics*, 106(4):1157–1189.

Boone, C., Dube, A., Goodman, L., and Kaplan, E. (2021). Unemployment Insurance Generosity and Aggregate Employment. *American Economic Journal: Economic Policy*, 13(2):58–99.

Borusyak, K., Jaravel, X., and Spiess, J. (2022). Revisiting Event Study Designs: Robust and Efficient Estimation. arXiv:2108.12419 [econ].

Braxton, J. C., Herkenhoff, K. F., and Phillips, G. M. (2020). Can the Unemployed Borrow? Implications for Public Insurance. Working Paper 27026, National Bureau of Economic Research.

Callaway, B. and Sant'Anna, P. H. C. (2021). Difference-in-Differences with multiple time periods. *Journal of Econometrics*, 225(2):200–230.

Card, D., Chetty, R., and Weber, A. (2007). The Spike at Benefit Exhaustion: Leaving the Unemployment System or Starting a New Job? *American Economic Review*, 97(2):113–118.

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

Carroll, C. D. and Summers, L. H. (1991). Consumption Growth Parallels Income Growth: Some New Evidence. In *National Saving and Economic Performance*, pages 305–348. University of Chicago Press.

Cattaneo, M. D., Crump, R. K., Farrell, M., and Feng, Y. (2021). On Binscatter. SSRN Scholarly Paper 3344739, Rochester, NY.

Chetty, R. (2006). A General Formula for the Optimal Level of Social Insurance. *Journal of Public Economics*, 90(10):1879–1901.

Chetty, R. (2008). Moral Hazard Versus Liquidity and Optimal Unemployment Insurance. *Journal of Political Economy*, 116(2):173–234.

Chetty, R., Friedman, J. N., Hendren, N., Stepner, M., and the Opportunity Insights Team (2020). The Economic Impacts of COVID-19: Evidence from a New Public Database Built Using Private Sector Data. Working Paper 27431, National Bureau of Economic Research.

Chodorow-Reich, G. and Coglianese, J. (2019). *Unemployment Insurance and Macroeconomic Stabilization*, pages 153–179. Brookings Institution, Washington, D.C.

Chodorow-Reich, G., Coglianese, J., and Karabarbounis, L. (2019). The Macro Effects of Unemployment Benefit Extensions: a Measurement Error Approach. *The Quarterly Journal of Economics*, 134(1):227–279.

Coibion, O., Gorodnichenko, Y., and Hong, G. H. (2015). The Cyclicality of Sales, Regular and Effective Prices: Business Cycle and Policy Implications. *American Economic Review*, 105(3):993–1029.

Coombs, K., Dube, A., Jahnke, C., Kluender, R., Naidu, S., and Stepner, M. (2022). Early Withdrawal of Pandemic Unemployment Insurance: Effects on Employment and Earnings. *AEA Papers and Proceedings*, 112:85–90.

de Chaisemartin, C. and D'Haultfœuille, X. (2020). Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects. *American Economic Review*, 110(9):2964–2996.

Deaton, A. (1991). Saving and Liquidity Constraints. *Econometrica*, 59(5):1221–1248.

DellaVigna, S. and Kim, W. (2022). Policy Diffusion and Polarization across U.S. States. Working Paper 30142, National Bureau of Economic Research.

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

Friedman, M. (1957). *Theory of the Consumption Function*. Princeton University Press.

Ganong, P., Greig, F. E., Noel, P. J., Sullivan, D. M., and Vavra, J. S. (2022). Spending and Job-Finding Impacts of Expanded Unemployment Benefits: Evidence from Administrative Micro Data. Working Paper 30315, National Bureau of Economic Research.

Ganong, P., Jones, D., Noel, P. J., Greig, F. E., Farrell, D., and Wheat, C. (2020a). Wealth, Race, and Consumption Smoothing of Typical Income Shocks. Working Paper 27552, National Bureau of Economic Research.

Ganong, P. and Noel, P. (2019). Consumer Spending During Unemployment: Positive and Normative Implications. *American Economic Review*, 109(7):2383–2424.

Ganong, P. and Noel, P. (2020). Liquidity versus Wealth in Household Debt Obligations: Evidence from Housing Policy in the Great Recession. *American Economic Review*, 110(10):3100–3138.

Ganong, P. and Noel, P. (2022). Why do Borrowers Default on Mortgages?*. *The Quarterly Journal of Economics*, page qjac040.

Ganong, P., Noel, P., and Vavra, J. (2020b). US Unemployment Insurance Replacement Rates During the Pandemic. *Journal of Public Economics*, 191:104273.

Gerard, F. and Naritomi, J. (2021). Job Displacement Insurance and (the Lack of) Consumption-Smoothing. *American Economic Review*, 111(3):899–942.

Gerardi, K., Herkenhoff, K. F., Ohanian, L. E., and Willen, P. S. (2017). Can't Pay or Won't Pay? Unemployment, Negative Equity, and Strategic Default. *The Review of Financial Studies*, 31(3):1098–1131.

Gourinchas, P.-O. and Parker, J. A. (2002). Consumption Over the Life Cycle. *Econometrica*, 70(1):47–89.

Griffith, R., Leibtag, E., Leicester, A., and Nevo, A. (2009). Consumer Shopping Behavior: How Much Do Consumers Save? *Journal of Economic Perspectives*, 23(2):99–120.

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

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

Hagedorn, M., Karahan, F., Manovskii, I., and Mitman, K. (2013). Unemployment Benefits and Enemployment in the Great Recession: The Role of Equilibrium Effects. *FRB of New York Staff Report*, (646).

Hall, R. E. (1978). Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence. *Journal of Political Economy*, 86(6):971–987.

Hamermesh, D. S. (1984). Consumption During Retirement: The Missing Link in the Life Cycle. *The Review of Economics and Statistics*, 66(1):1–7. Publisher: The MIT Press.

Hausman, J. and Paquette, L. (1987). Involuntary early retirement and consumption. In *Work, Health, and Income Among the Elderly*, pages 151–175.

Housing Finance Policy Center (2020). Housing Finance at a Glance: a Monthly Chartbook (June 2020). Technical report, Urban Institute.

Hsu, J. W., Matsa, D. A., and Melzer, B. T. (2018). Unemployment Insurance as a Housing Market Stabilizer. *American Economic Review*, 108(1):49–81.

Hurd, M. D. and Rohwedder, S. (2005). The Retirement-Consumption Puzzle: Anticipated and Actual Declines in Spending at Retirement. Publisher: RAND Corporation.

Hurd, M. D. and Rohwedder, S. (2006). Economic Well-Being at Older Ages: Income- and Consumption-Based Poverty Measures in the HRS. Working Paper 12680, National Bureau of Economic Research.

Hurd, M. D. and Rohwedder, S. (2010). Effects of the Financial Crisis and Great Recession on American Households. Working Paper 16407, National Bureau of Economic Research.

Hurst, E. (2008). The Retirement of a Consumption Puzzle. Technical Report w13789, National Bureau of Economic Research.

Iacus, S. M., King, G., and Porro, G. (2012). Causal Inference Without Balance Checking: Coarsened Exact Matching. *Political Analysis*, 20(1):1–24.

Jacobson, L. S., LaLonde, R. J., and Sullivan, D. G. (1993). Earnings Losses of Displaced Workers. *The American Economic Review*, 83(4):685–709. Publisher: American Economic Association.

Jaravel, X. (2019). The Unequal Gains from Product Innovations: Evidence from the U.S. Retail Sector. *The Quarterly Journal of Economics*, 134(2):715–783.

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.

Johnston, A. C. and Mas, A. (2018). Potential Unemployment Insurance Duration and Labor Supply: The Individual and Market-Level Response to a Benefit Cut. *Journal of Political Economy*, 126(6):2480–2522.

Kaplan, G. and Menzio, G. (2015). The Morphology of Price Dispersion. *International Economic Review*, 56(4):1165–1206.

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. (2021). Unemployment Insurance in Macroeconomic Stabilization. Working Paper 29505, National Bureau of Economic Research.

Keys, B. J. (2018). The Credit Market Consequences of Job Displacement. *The Review of Economics and Statistics*, 100(3):405–415.

Koustas, D. (2018). Consumption Insurance and Multiple Jobs: Evidence from Rideshare Drivers. Working Paper.

Krueger, A. B. and Mueller, A. (2010). Job search and unemployment insurance: New evidence from time use data. *Journal of Public Economics*, 94(3):298–307.

Laibson, D. (1997). Golden Eggs and Hyperbolic Discounting. *Quarterly Journal of Economics*, 112(2):443–478.

Landais, C., Michaillat, P., and Saez, E. (2018). A Macroeconomic Approach to Optimal Unemployment Insurance: Applications. *American Economic Journal: Economic Policy*, 10(2):182–216.

Landais, C. and Spinnewijn, J. (2021). The Value of Unemployment Insurance. *Review of Economic Studies*, forthcoming.

McKay, A. and Reis, R. (2021). Optimal Automatic Stabilizers. *The Review of Economic Studies*, 88(5):2375–2406.

Modigliani, F. and Brumberg, R. (1954). Utility Analysis and the Consumption Function: An Interpretation of Cross-Section Data. In *Post-Keynesian Economics*, pages 388–436. Rutgers University Press, New Brunswick, NJ.

Nevo, A. and Wong, A. (2019). The Elasticity of Substitution Between Time and Market Goods: Evidence from the Great Recession. *International Economic Review*, 60(1):25–51.

Olafsson, A. and Pagel, M. (2018a). The Liquid Hand-to-Mouth: Evidence from Personal Finance Management Software. *The Review of Financial Studies*, 31(11):4398–4446.

Olafsson, A. and Pagel, M. (2018b). The Retirement-Consumption Puzzle: New Evidence from Personal Finances. Working Paper 24405, National Bureau of Economic Research.

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.

Schmieder, J. F. and von Wachter, T. (2017). A Context-Robust Measure of the Disincentive Cost of Unemployment Insurance. *American Economic Review*, 107(5):343–48.

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

Sun, L. and Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, 225(2):175–199.