Wealth, Race, and Consumption Smoothing of Typical Labor Income Shocks*

Peter Ganong, University of Chicago and NBER
Damon Jones, University of Chicago and NBER
Pascal Noel, University of Chicago and NBER
Diana Farrell, JPMorgan Chase Institute
Fiona Greig, JPMorgan Chase Institute
Chris Wheat, JPMorgan Chase Institute

April 18, 2023

Abstract

We study the consumption response to typical labor income shocks and investigate how these vary by wealth and race. First, we develop an instrument based on firm-wide changes in labor income. Household income volatility stems mostly from fluctuations in labor income and this research design therefore studies the sort of income fluctuations that households typically experience from month to month. Using administrative banking data, we find an average elasticity of 0.21, with a much higher elasticity for low-liquidity households and close to zero elasticity for high-liquidity households. In a stylized model calibrated to our estimates, this degree of sensitivity implies that temporary income volatility has a large welfare cost for the average household, and especially large costs for low-liquidity households. Second, we use this instrument to study how wealth shapes racial inequality. Although an extensive body of work documents the long-term persistence of the racial wealth gap, less is known about its consequences on households' lives from month to month. We find that Black and Hispanic households are twice as sensitive to typical income shocks as White households. Nearly all of this difference is explained in a statistical sense by racial wealth inequality. Because of racial disparities in consumption smoothing, the welfare cost of temporary income volatility is twice as high for Black and Hispanic households than for White households.

*We thank Fenaba Addo, Richard Blundell, Kerwin Charles, Erik Hurst, Koichiro Ito, Greg Kaplan, Dmitri Koustas, Bruce Meyer, Emi Nakamura, Scott Nelson, Matthew Notowidigdo, Cormac O'Dea, Jonathan Parker, Jon Roth, Jesse Shapiro, Jón Steinsson, Amir Sufi, Joe Vavra, and Rob Vishny for helpful comments, and Adi Aladangady, Scott Baker, Karen Dynan, Stephen Hansen, Peter Henry, David Johnson, and Jialan Wang for discussing the paper. We thank Marta Lachowska and Stephen Woodbury for discussions about and analysis of the sources of pay volatility in the Washington state data. We thank numerous seminar participants for their helpful comments. We thank Therese Bonomo, Timotej Cejka, Guillermo Carranza Jordan, Rupsha Debnath, Maxwell Liebeskind, Lei Ma, Roshan Mahanth, Michael Meyer, Liam Purkey, Peter Robertson, Tanya Sonthalia, and John Spence for outstanding research assistance. This research was made possible by a data-use agreement between three of the authors and the JPMorgan Chase Institute (JPMCI), which has created de-identified data assets that are selectively available to be used for academic research. All statistics from JPMCI data, including medians, reflect cells with multiple observations. The opinions expressed are those of the authors alone and do not represent the views of JPMorgan Chase & Co. We gratefully acknowledge funding from the Washington Center for Equitable Growth, the University of Chicago Center for Data and Computing, the University of Chicago Center for the Study of Race, Politics, and Culture, and the Kathryn and Grant Swick and Fujimori/Mou Faculty Research Funds at the University of Chicago Booth School of Business. First draft: April 2020.

1 Introduction

Many households face substantial fluctuations in labor income from month to month. For example, in the Survey of Income and Program Participation, the standard deviation of monthly labor income changes is 30 percent. At the same time, households report holding few liquid assets that could be used to buffer such volatility. According to the Survey of Consumer Finances, 40 percent of Americans have less than two weeks worth of income in liquid assets. The volatility of monthly income in the first survey combined with the lack of liquid assets in the second survey raises concerns that it may be difficult for households to smooth consumption from month to month. Yet, perhaps surprisingly, there is little empirical evidence on how typical month-to-month fluctuations in labor income affect consumption, and even less evidence about how this consumption smoothing varies with wealth.

Concerns about the ability to weather monthly labor income shocks are particularly pronounced for Black households. Fifty-eight percent of Black households have less than two weeks worth of income in liquid assets, as compared to 33 percent of White households. This Black-White disparity is one manifestation of the racial wealth gap in the U.S. This gap has deep historical roots. At the end of the Civil War, General Sherman promised "forty acres and a mule" to formerly enslaved Black people. Yet his order was reversed shortly thereafter, and many federal policies to promote asset accumulation since then have deliberately excluded Black households (Lui et al. 2006). Today, the median Black household has only one-sixth the liquid wealth of the median White household. A similar gap in liquid wealth exists between Hispanic and White households. The existence and long-term persistence of the racial wealth gap have been carefully studied (see e.g., Derenoncourt et al. 2022). However, there is less quantitative research about the consequences of this gap on households' lives from month to month, such as the ability to weather income volatility.

The goal of this paper is to construct estimates of the consumption response to typical labor income shocks and investigate how these vary by wealth and race. A core challenge in achieving this goal is isolating exogenous variation in month-to-month labor income changes. To overcome this challenge, we instrument for changes in individual labor income using employer-wide changes in monthly pay, which we hereafter refer to as a firm pay shock design. The key idea is that employer-wide changes in pay capture shocks that are unrelated to an individual worker's endogenous labor supply decision (Koustas 2018). The three most common reasons for income volatility reported by employed households in the U.S. are, in descending order, irregular work schedules, bonuses and commissions, and seasonal employment. Firm pay shocks capture all these common sources of pay variation, and are therefore likely to be externally valid for the type of income volatility that workers regularly face. Furthermore, they can be identified using a generalization of the classic model of firm effects by Abowd, Kramarz and Margolis (1999).

For our empirical analysis, we use administrative records to build a de-identified dataset with information on income, consumption, liquid wealth, and race. The first three variables come from a monthly panel of Chase bank account records over five years. Self-reported race and Hispanicity information comes from voter registration records.[1] These two datasets were matched using personal identifiers and then de-identified.[2] The two key strengths of this dataset are a large sample size of nearly two million households and employer identifiers that enable our main empirical strategy. To assess external validity, we compare the distributions by race of income, liquid wealth, age, and sex in the analysis sample to nationally representative surveys; racial inequality in the new dataset mirrors racial inequality in the U.S. as a whole.

We implement the firm pay shocks design in this dataset and present four reduced-form empirical results. We then examine their welfare implications using a simple model.

First, we find that typical labor income fluctuations meaningfully affect consumption for the average household. We instrument for an individual's pay change with the change in average pay of their coworkers. This instrument has a strong first stage.[3] Our central estimate using this instrument is that the elasticity of consumption of nondurables with respect to income is 0.21.

Second, we find that the consumption of Black and Hispanic households is twice as sensitive to monthly labor income shocks as the consumption of White households. Specifically, we estimate elasticities of 0.34 for Black households, 0.29 for Hispanic households, and 0.16 for White households. To be clear, these estimates capture heterogeneity by race in a causal estimate, rather than the "causal effect" of race, which is an ill-formed concept given the endogenous and socially constructed nature of racial identity.

Third, we find that the extent of consumption smoothing also varies sharply with liquid assets. We find that the elasticity declines monotonically as liquid assets increase, ranging from 0.50 for the lowest-asset households to 0.08 for the highest-asset households.

Fourth, we find that wealth inequality by race can largely account, in a statistical sense, for the racial disparities in consumption smoothing that we document. To investigate this link, we estimate elasticities by race and control for liquid assets. When we control for liquid assets at the household level, about eighty percent of the Black-White and Hispanic-White differences in elasticities are eliminated, and the gaps that remain are not statistically

[1] We use 2018 voter registration records in the three states that had Chase branches in 2018 and that collect race information during voter registration. The voter registration forms ask a single question that offers "Black" and "Hispanic" (or "Hispanic/Latino" in one case) as mutually exclusive categories. Going forward, we use the shorthand "race" to describe responses to the question. We therefore use "Hispanic" to refer to respondents who have selected "Hispanic" (or "Hispanic/Latino") in response to the question, acknowledging that they may in practice self-identify with a variety of related terms. Although the Census treats "Hispanic" solely as an ethnic group, some members of Hispanic/Latino/Latinx groups see their identity as a racial category (Parker et al. 2015). See Almaguer (2012) for further discussion of the complex and evolving racialization of Latino populations in the U.S. over time.

[2] See Section 2.1 for more detail on the matching process and steps taken to protect privacy.

[3] Indeed, we demonstrate the general applicability of the firm pay shock instrument by documenting the presence of economically and statistically significant first stages in two additional types of datasets: time clock data from Homebase and administrative earnings data.

significant.

One way to interpret this fourth finding is as a neutrality result: when faced with the same income shocks *and* financial buffers, households of all races react similarly. This suggests that non-wealth channels through which differences in elasticities by race might arise are either quantitatively small, cancel each other out, or are correlated with wealth. These channels may include differences in access to credit, debt obligations, family structure, family labor supply, transfers from family and friends, access to transfer programs, expectations, and preferences.

These reduced-form findings continue to hold in three complementary research designs, each of which isolates the effect of a different source of labor income volatility. The primary research design uses variation in monthly pay per paycheck. This design captures variation in income which (i) has a recurring seasonal pattern (and is therefore possibly anticipated), (ii) is partly driven by hours worked, (iii) is concentrated in one month before gradually abating over several months, and (iv) is only relevant within employment spells. We explore three additional specifications of the instrument with different characteristics. The first isolates idiosyncratic non-seasonal (and therefore likely unanticipated) variation in income. The second relies on the receipt of additional checks in a month due to the pay calendar. This specification captures variation that is not driven by hours worked and lasts exactly one month. The third specification uses unemployment spells, capturing volatility between jobs rather than within jobs. Across all three research designs, we find substantially less consumption smoothing by Black and Hispanic households, less consumption smoothing by low-asset households, and that the racial wealth gap in liquid assets can account for most of the differences in consumption smoothing by race.

In the final section of the paper we consider the welfare implications of our results by calculating the cost of temporary income volatility. We combine our reduced-form empirical estimates with a widely-used model and standard preference assumptions based on Lucas (1987). The model is deliberately simple. Three sufficient statistics are required to calculate the welfare cost: the elasticity of consumption with respect to temporary income shocks, the coefficient of relative risk aversion, and the variance of temporary income shocks.

The model implies that income volatility has a substantial welfare cost for all groups, and further that the cost is substantially larger for Black and Hispanic households than it is for White households. We calculate the welfare gain from eliminating transitory income volatility in a model that is calibrated to our empirical estimates and uses common assumptions for preferences. We find a willingness to pay to eliminate volatility of 0.7 percent of lifetime consumption for White households, 1.4 percent for Hispanic households, and 1.8 percent for Black households. These estimates may be an upper bound on the true cost of such fluctuations for two reasons. First, as would be the case with any modern consumption-savings model, they assume that temporary income fluctuations do not directly affect utility (i.e., through changes in hours worked). Second, they assume that the elasticity of consumption with respect to income for all temporary labor income changes is the same as the Local Average Treatment Effect we estimate using temporary changes in firm pay. Nevertheless, relative to the benchmark view in Lucas (1987) that a cost of 0.5 percent of lifetime consumption is "large," these estimates suggest there could be a substantial welfare cost for all groups, and an especially large cost for Hispanic and Black households.

This paper makes two primary contributions. First, we provide estimates of monthly consumption sensitivity by race and connect racial inequality in *consumption* sensitivity to racial *wealth* inequality. These estimates relate to a rich literature on the extent of the racial wealth gap (e.g., Browne 1974; Oliver and Shapiro 2006); statistical correlates of this gap (Blau and Graham 1990; Smith 1995; Barsky et al. 2002; Altonji and Doraszelski 2005); its interplay with factors such as home ownership (Charles and Hurst 2002; Aaronson, Hartley and Mazumder 2021; Kermani and Wong 2022), marriage (Addo and Lichter 2013), inter-household transfers (Chiteji and Hamilton 2002), inter-generational transfers (Toney and Robertson 2021), and retirement savings incentives (Choukhmane et al. 2022); as well as potential policy solutions (Hamilton and Darity 2010). Most wealth gap studies focus on differences in levels of wealth and longer term trends in these gaps.

Less is known about the consequences of the racial wealth gap on households' lives from month to month. However, a newly emerging literature leverages administrative data—from banks, credit reports, financial aggregators, and payment platforms—to study households' financial lives (e.g., Aladangady et al. Forthcoming; Baker 2018; Baugh et al. 2021; Gross, Notowidigdo and Wang 2020; Keys, Mahoney and Yang 2022; Olafsson and Pagel 2018). Unfortunately, race is rarely observed in such datasets.[4] We bring these developments in the household finance literature to bear on questions of racial inequality by examining the relationship between the racial wealth gap and the transmission of income volatility to consumption.

Our neutrality finding—the finding that controlling for assets largely eliminates racial differences in the *sensitivity* of consumption—builds on a famous racial neutrality result on the *level* of consumption in Friedman (1957). In contrast to studies that had documented differences in consumption and saving rates between Black and White households conditional on annual income (Mendershausen 1940; Duesenberry 1949; Tobin 1951), Friedman showed instead that these consumption differences can be explained by differences in permanent income. In a more recent neutrality result, Charles, Hurst and Roussanov (2009) showed that Black-White differences in the distribution of permanent income can explain racial differences in the *composition* of consumption. Such results suggest that when faced with the same circumstances, households of all races react similarly (Hamilton and Darity 2017), posing challenges for so-called essentialist theories of racial differences in economic outcomes Darity (2022).

[4]One notable exception is Argyle et al. (2023), which documents a large role for discrimination in explaining racial disparities in consumer bankruptcy outcomes.

These neutrality results should not be interpreted to mean that race is irrelevant for understanding consumption inequality. Rather, they indicate channels which mediate such racial inequality. In the case of inequality in consumption sensitivity, our results indicate that such inequality is likely mediated through the racial wealth gap, which is caused in part (and possibly entirely) by a history of barriers to asset accumulation for Black and Hispanic households. Under the assumption that liquid wealth has a causal impact on consumption smoothing, these results suggest that the racial wealth gap leaves Black and Hispanic households more vulnerable to labor income fluctuations than White households.

The paper's second contribution is to provide causal evidence on the consumption-smoothing of *typical labor* income fluctuations. This evidence complements prior estimates from a leading empirical strategy in the consumption smoothing literature which documents substantial consumption responses to windfalls of *unusual non-labor* income.[5] Consumption researchers have adopted two alternative interpretations of this empirical finding with conflicting out-of-sample predictions for frequently occurring income shocks. The first interpretation rationalizes high windfall spending through two-asset consumption-savings models in which either preferences (Laibson, Repetto and Tobacman 2007) or technology (Kaplan and Violante 2014) lead to low liquidity as well as a substantial response to income fluctuations, even frequently occurring ones. The second interpretation treats windfall spending as "near rational" behavior (e.g., Kueng 2018; Fuchs-Schundeln and Hassan 2016). This argument is based on the observation first developed in Cochrane (1989) that there is close to zero welfare cost of failing to smooth consumption in the face of small or infrequent income changes.[6] This interpretation implies that households *would* smooth frequently-occurring income shocks, in line with Friedman's Permanent Income Hypothesis, because the cost of failing to smooth such shocks is large.

The design in this paper addresses the Cochrane (1989) critique by studying a source of frequently-occurring income fluctuations that, if not smoothed, have a substantial welfare cost: typical labor income shocks. Most working-age households get the vast majority of their income from labor, and as a consequence most income volatility comes from labor income instead of non-labor income.[7] Yet we find that households are still highly sensitive even to typical labor income volatility. The results in this paper thus favor the low-liquidity

[5]Examples of such windfalls include life insurance dividends, restitution payments, tax rebates, veterans' bonuses, lottery winnings, and Alaska Permanent Fund payments. See Bodkin (1959), Kreinin (1961), Johnson, Parker and Souleles (2006), Hausman (2016), Fagereng, Holm and Natvik (2021), and Kueng (2018), respectively.

[6]This interpretation is bolstered by mental accounting theories of consumption behavior. For example, when the first stimulus checks were sent out in July 2001, Time Magazine reported that White House cabinet members "spent their time on the Sunday shows essentially calling for a mass national shopping spree" (Pellegrini 2001). In fact, anomalously high spending out of windfalls is used to motivate the canonical mental accounting papers (Thaler 1985; Thaler 1990). Recent evidence shows that labels applied to government transfers can have dramatic effects on how funds are spent (Hastings and Shapiro 2018; Beatty et al. 2014).

[7]We calculate using the SIPP that the standard deviation of monthly changes in labor income is four times larger than the standard deviation of monthly changes in non-labor income.

interpretation of consumption sensitivity over the near rationality interpretation.[8]

Our finding that households are highly sensitive to typical labor income volatility leads us to revisit a classic view in macroeconomics that temporary income shocks are irrelevant for welfare. Indeed, a standard approach in this literature is simply to assume that these welfare costs are zero (e.g., Heathcote, Storesletten and Violante 2008). We show instead that households are highly sensitive even to routine transitory shocks, implying that such volatility has large welfare costs. Furthermore, this cost is highest for groups with low average liquid assets such as Black and Hispanic households. This reinforces the conclusion in Hardy et al. (2018) that policies which mitigate temporary income volatility may also reduce inequality by race.

Finally, our focus on typical income shocks also builds on a second strand of the consumption-smoothing literature pioneered by Blundell, Pistaferri and Preston (2008). This "semi-structural" literature uses *all* variation in income and defines income shocks by modeling the earnings process.[9] Although this method captures a broad set of income changes, it may suffer from concerns about the endogeneity of labor supply, which can vary with consumption needs. In contrast, this paper's primary research design isolates *exogenous* variation arising from changes in firm-wide labor demand. We thus seek to combine the strength of the semi-structural approach (capturing typical income variation) with the strength of the quasi-experimental windfall approach (capturing exogenous variation). Consistent with endogenous labor supply dampening the relationship between consumption and income, we find that the consumption elasticity is about 50 percent higher in a specification that uses exogenous variation from firm pay shocks than in a specification that uses all variation in labor income.

The paper proceeds as follows. Section 2 describes the data and compares the analysis sample to several public use datasets. Section 3 describes the empirical results using firm-wide variation in pay and unemployment. Section 4 quantifies the welfare cost of income volatility using a simple model. Section 5 concludes.

2 Data and External Validity

We study consumption smoothing by using a link between two administrative datasets: bank account data from the JPMorgan Chase Institute (JPMCI) and race data from voter registration files in three states. This de-identified dataset has two advantages over existing data sources: it includes nearly two million households and plausible instruments for income

[8] Our results are consistent with models such as Laibson, Repetto and Tobacman (2007) and Kaplan and Violante (2014) in which low liquidity arising in part from negative income shocks *causes* high sensitivity. They are also consistent with models where an alternative channel that is correlated with wealth (e.g., inter-household transfers in Chiteji and Hamilton (2002) or permanent heterogeneity as in Aguiar, Bils and Boar (2021), Epper et al. (2020), or Ganong et al. (2022)) contributes to both low liquidity and high MPCs.

[9] See also Hall and Mishkin (1982), Arellano et al. (2023), Attanasio and Pavoni (2011), Choi, McGarry and Schoeni (2016), Etheridge (2015), and Guvenen and Smith (2014). Earlier work which uses all variation in income includes Deaton and Paxson (1994), Guvenen (2007), and Townsend (1994).

fluctuations. We describe the source datasets, show that the matching procedure is accurate, and then use external benchmarks to explore external validity. It appears that racial gaps in household finances in the matched data are either similar to, or a lower bound on, national gaps by race.

2.1 Source Datasets

The first dataset for this paper is drawn from the monthly data on 20 million households with a Chase checking account. The unit of observation is household-by-month, from January 2018 through November 2022.

Our analysis centers on three variables from this dataset: income, spending, and assets. JPMCI categorizes each checking account inflow by whether it appears to reflect income. Among payments received by direct deposit, we are able to observe several types of income, including labor income, income from government sources, and capital income. Most of our analysis focuses on labor income, for which we observe an encrypted version of the transaction description. This field enables us to identify other workers with Chase bank accounts who are paid by the same employer. This is useful because, as we describe in Section 3, we use fluctuations in firm pay to instrument for changes in labor income.

The second key variable is a measure of spending on nondurable goods and services. Spending is measured from debit and credit card transactions, cash withdrawals, and electronic transactions captured through the bank account. Ganong and Noel (2019) constructs a measure of nondurable spending in the JPMCI data, and we adopt that definition in this paper. Examples of nondurable spending include groceries, food away from home, fuel, utilities, clothing, medical co-pays, and payments at drugstores. Altogether, transactions that can be affirmatively categorized as nondurable comprise an average of 44 percent of checking account outflows.

Ganong and Noel (2019) shows that this measure of nondurable spending is more comprehensive than what can be measured in U.S. survey data, but nevertheless has some limitations which are relevant for this paper. First, we only observe data on bank account and credit card spending at a single financial institution. We estimate using the Survey of Consumer Finances (SCF) that about two-thirds of households have checking accounts at only one financial institution; for the remaining one-third of households, we miss some spending through non-Chase bank accounts. In addition, although the data include spending on Chase credit cards, they do not include spending on non-Chase credit cards. Second, even if we had data on spending from all financial institutions, we still would not observe in-kind transfers, which may be an important source of consumption smoothing.

The third key variable is the sum of balances in the household's Chase checking accounts. We supplement this with external data from the Survey of Consumer Finances (SCF) on comprehensive measures of total liquid assets (using the definition in Kaplan and Violante 2014) and total financial assets. Analysis of the SCF indicates that checking account balance is a useful proxy for total assets. Although households frequently hold most of their financial assets outside the checking account, average log financial assets increase one-for-one with average log checking account balances, as shown in Figure A-1. The same figure also demonstrates that, for a given checking account balance, Black and Hispanic households have less than White households in total liquid assets. For example, consider households with a checking balance near the national median of $2,100: the median White household has liquid assets of $3,100, while the median Black and Hispanic households still have total liquid assets of this same amount ($2,100). We thus use the SCF to impute assets from checking account balances and race. See Table A-1 for the imputation factors. We also observe age, sex, and ZIP code of residence for the primary bank account holder. Chase has a footprint via physical bank branches in 22 states and nearly all Chase bank account holders live in these states.

The second dataset is voter registration files available to the public, which we use to measure race. Although the Fifteenth Amendment gave Black citizens the right to vote in 1870, Jim Crow laws and practices in Southern states prevented most Black citizens from voting for nearly a century thereafter. To address this discrimination, the U.S. Congress passed the Voting Rights Act of 1965, which seeks to ensure equal voting rights for Black citizens and other minorities. Today, eight Southern states collect data on race as part of voter registration. Hersh (2015) finds that the origins of this practice in these specific states are unclear, but notes that the set of states is similar to the set subject to additional scrutiny under the Act as well as the set that had White-only primaries in the Jim Crow era. Figure A-2 shows a map of these states. The voter files also contain name, address, and date of birth, which we use for matching.[10]

Our analysis sample is drawn from three states: Florida, Georgia, and Louisiana. These are states where the voter files contain race and there are a substantial number of people with Chase bank accounts. The voter registration forms in these states ask respondents to provide a single answer to a single question, which does not separately capture race and Hispanic identity. This formulation is in line with how the Census asked about race in 1960 and 1970, but not how the Census has asked about race since 1980. This means that we are unable to separately analyze race and Hispanicity. For example, we cannot distinguish Hispanic individuals who identify as White from Hispanic individuals who identify as Black. The exact wording of the question in each state, and in the SCF, is shown in Table A-2 and the distribution of responses is shown in Table A-3.

We also use two supplementary datasets to understand the degree to which changes in

[10]Voter registration records contain personal identifiers, including name, address, and birthdate, which were matched to bank records and then delivered to the JPMorgan Chase Institute stripped of these personal identifiers. Voter registration data was obtained in 2018 for the exclusive purpose of enabling the JPMorgan Chase Institute to conduct this research examining financial outcomes by race and not to identify party affiliation. The matched file that contains personal identifiers, banking records, and self-reported race has been deleted. The remaining de-identified file that contains banking records and self-reported race is only available to the JPMorgan Chase Institute and is not being maintained by or made available to JPMC business units or other parts of the firm.

coworker pay are correlated with own pay. The first data set is the Continuous Wage and Benefit History (CWBH) which captures quarterly wage records for UI claimants from seven U.S. states. We are able to construct the change in coworker pay because the dataset includes information on total firm pay and total number of employees at the firm for each UI claimant. The second dataset is time clock data from Homebase, detailing when workers start and stop their shifts. The company's software is primarily used by small service-sector firms such as coffee shops.

2.2 Matched Dataset with Bank Account Data and Voter Registration

A JPMorgan Chase data team used a fuzzy matching algorithm to link bank records to voter records. For the purposes of measuring race, this procedure treats the name listed first on the bank account as the household head. The algorithm finds 56 percent of potential matches.[11] We discuss details of the algorithm in Appendix B.1. The matched file with bank records and self-reported race was then delivered to the JPMorgan Chase Institute stripped of personal identifiers.

The analysis sample consists of 1.8 million matched bank-voter records. The sample includes 870,000 White households, 462,000 Black households, and 414,000 Hispanic households. Until now, administrative data with information on household finances and race have largely been limited to the analysis of mortgages.[12] Thus, the number of observations in the analysis sample is orders of magnitude larger than previously available datasets used to study household finance and race. The analysis sample is concentrated in urban areas, reflecting the geographic distribution of Chase branches within these three states.[13]

We validate the quality of the matches using a companion sample where JPMCI directly observes race for a subset of Chase customers. Although in general banks do not and cannot collect information on the race and ethnicity of their customers, the one exception is that they are required by law to ask mortgage applicants one question about their race and a

[11] This 56 percent estimate is the ratio of the fraction of bank customers with a match (40 percent) to the estimated fraction of bank customers registered to vote (72 percent). We estimate the 72 percent voter registration rate for the bank customers based on CPS data which shows that 72 percent of people with bank accounts in the three states are registered to vote. A 40 percent match rate is at the lower end of similar studies that link administrative data to voter files (Kempf and Tsoutsoura 2021; Akee et al. 2018; Baicker and Finkelstein 2019). One reason the match rate is lower than prior studies is that we impose stringent similarity standards for name, address and age.

[12] Federal law requires that mortgage lenders collect data on race. Fuster et al. (2022), Bartlett et al. (2022), and Bhutta and Hizmo (2020) analyze discrimination in mortgage lending. We are aware of only one prior study (Avery, Brevoort and Canner 2009) that links data on race to administrative data on a broad set of household financial measures. Argyle et al. (2023) imputes race from names and addresses using a deep-learning model to analyze racial disparities in bankruptcy outcomes. In related work, Dobbie et al. (2021) documents discrimination against immigrants in the UK.

[13] Table A-4 and Figure A-3 demonstrate this concentration in urban areas (as compared to the overall distribution of Blacks and Hispanics shown in Figure A-4). Although one might have expected Black and Hispanic households to be *under*-represented because they are less likely to report having a bank account in the Current Population Survey, we find the opposite. Black and Hispanic households are *over*-represented in the analysis sample relative to the universe of voters in the three states. Table A-5 shows that much of this over-representation is explained by two facts: (1) Black and Hispanic households tend to live in urban areas and (2) Chase branches tend to be in urban areas.

second about their Hispanicity. There are 194,000 de-identified Chase mortgage applications that JPMorgan Chase is able to match to the voter files.[14]

Race in the voter files agrees with the mortgage application files for 99 percent of people, which gives us confidence in the accuracy of the matching procedure. Appendix B.2 describes how we assess whether race agrees between the two datasets. A 99 percent agreement rate is remarkably high, considering that even when questions are worded consistently, some individuals' self-reported race and ethnicity changes over time (Liebler et al. 2017).

2.3 Comparison of Matched Dataset to External Benchmarks

One natural question is whether people in the dataset have different economic and demographic patterns by race than the U.S. as a whole. To better understand the external validity of estimates using these data, we decompose this question into two parts. First, how do patterns by race compare between the nation as a whole and people with bank accounts in these three states? Second, how does this three-state banked population compare to the analysis sample, which is composed of people who have a bank account with Chase and whose race can be ascertained from the voter files?

We find that differences by race for banked households in Florida, Georgia, and Louisiana generally track differences by race nationally. The best dataset for capturing a representative banked sample from these three states is the Current Population Survey (CPS). We use the CPS to analyze household heads who have a bank account and live in Florida, Georgia and Louisiana. We refer to people who meet these criteria as being in the "sample frame." The CPS also captures four statistics which we can easily compare to the matched JPMCI sample: race, age, sex, and income.[15]

It is unclear a priori whether Black-White economic gaps in the sample frame should be larger or smaller than national gaps. One reason they might be larger is that voter data with race is specifically available in places where the Voting Rights Act was instituted to combat a history of discrimination against Black households. This history includes under-investment in public goods such as education (Cascio and Washington 2014). However, there are also some reasons why Black-White gaps might be smaller. One reason is that White households in the Deep South also have low average incomes: Georgia, Florida, and Louisiana all rank in the bottom quarter of U.S. states in 2017 in terms of average household income for White households. Further, 17 percent of Black households were unbanked in 2017 (compared to 3 percent of White households) and the unbanked are excluded from the sample frame by construction.

[14] We do not use the mortgage application sample for our primary analysis both because the share of Black and Hispanic households in this sample is small and because applicants for mortgages tend to be economically advantaged, which attenuates racial gaps relative to population averages.

[15] To align with the mutually exclusive concept of race and Hispanicity from the voter files, we allocate people in the CPS to one of four mutually exclusive categories: non-Hispanic White, non-Hispanic Black, Hispanic, or other. This choice follows Chetty et al. (2019), Card and Rothstein (2007), and Federal Reserve Board (2017).

In practice, we find that Black-White income inequality in the sample frame is quite similar to national patterns. Figure A-5a shows that nationally, the median Black household earns 65 cents for each dollar earned by the median White household; in the sample frame, the ratio is similar but slightly higher: 77 cents. This similarity arises because, for both Black and White households, every quintile of the national income distribution is represented in approximately equal proportions in the sample frame (see Figure A-6a). Remarkably, this pattern holds even for the bottom quintile of the national income distribution: 19 percent of Black households and 21 percent of White households in the sample frame have incomes below the national 20th percentile race-specific thresholds. This pattern arises because of two offsetting forces: a geographic screen that increases the share of low-income households and a banked screen that reduces the share of low-income households.

National differences between Black and White households in terms of age and sex also hold in the sample frame. One distinctive feature of the U.S. is that Black household heads are more likely to be female than male; Figure A-5b shows that this also holds in the sample frame. In addition, Figure A-5b shows that Black household heads are younger than White household heads, both nationally and in the sample frame.

The story for Hispanic households is more nuanced, owing to the unique history of Hispanics in Florida. The majority of Hispanics in the U.S. are of Mexican origin, but fewer than one in five Hispanics in the three states are Mexican. This pattern arises in large part because Florida has attracted a diverse set of Hispanic immigrants from across Latin America. Figure A-7 compares the countries of origin for Hispanics across the U.S. to the three states. In the three states, the most common ancestry group is Cubans and there are as many Puerto Ricans as there are Mexicans. This presumably arises because of Florida's relative proximity to Cuba and Puerto Rico. In addition, Florida has attracted a substantial number of Hispanics from Colombia, Venezuela, and other countries in South America. Both Cubans and South Americans tend to be more affluent and older than other Hispanic immigrants (Pew Hispanic Center 2013).

We find that the broad economic and demographic contours of differences between Hispanic and White households nationally are preserved in the sample frame. Figure A-5a shows that the Hispanic-White income ratio is similar in the sample frame to the nation and Figure A-6a shows that the distribution of income is similar. Figure A-5b shows that when using median age, Hispanic households are the youngest of the three groups, both in the sample frame and nationally. However, older Hispanic households are over-represented (see Figure A-6b).

National racial wealth gaps are also preserved in the sample frame. One common way to measure the racial wealth gap is to compare the median Black (Hispanic) household to the median White household. Using this metric, Figure A-5a shows that Black (Hispanic) liquid wealth is 18 (14) cents for each dollar of White liquid wealth nationally. A slightly smaller racial wealth gap is apparent for checking balances. These national gaps are preserved in the sample frame.[16]

One striking fact visible in Figure A-5a is the extent to which racial wealth gaps far exceed racial income gaps. The fact that racial wealth gaps exceed racial income gaps is not driven by a low level of savings among all races by households with low annual income. Rather, Black and Hispanic households have less liquid wealth than White households at every level of the income distribution. For example, among households who say that their income in a normal year is about $55,600 (the national household median in the SCF), Black, Hispanic, and White households have $2,000, $1,800, and $4,000 respectively in liquid wealth.[17] These racial wealth gaps mean that, for an equivalent-sized income loss, a Black or Hispanic family will have less financial buffer to smooth consumption than a White family.

Finally, data on self-reported financial vulnerability suggest that racial differences in the sample frame will be a *lower bound* on national differences by race. The CPS asks questions about financial vulnerability as part of its Unbanked and Underbanked Supplement. Forty-two percent of Americans say that they do not have money set aside that could be used for unexpected expenses or emergencies. This masks heterogeneity by race: 55 percent of Black households do not have savings for unexpected shocks, as compared to 38 percent of White households, for a Black-White gap of 17 percentage points (see Figure A-8b). In the sample frame, the gap is 11 percentage points. For the share of households who report being behind on bills in the past year, the Black-White gap is 17 percentage points nationally versus 11 percentage points in the sample frame (see Figure A-8a). Similar patterns are apparent for Hispanic households.

Having shown that racial income and wealth gaps are similar in the sample frame to the nation, we show that similar economic and demographic patterns by race hold in the sample frame and in the analysis sample. Figure A-5a shows that racial gaps in income and assets are similar. Figure A-5b shows that patterns of gender and age by race are similar as well. The median age in the analysis sample is about five years younger than the sample frame, but the ranking of median age by race is the same. Figure A-6 compares the distribution of income and age in the sample frame and in the analysis sample.

To summarize, we construct a novel, large, de-identified dataset with information on

[16] Unfortunately, the public use SCF has no geographic information and so we use the Health and Retirement Study (HRS) to compute a regional adjustment factor. To compute the three-state banked estimates in Figure A-5a, we rescale the national estimates from the SCF by the ratio of regional asset medians to national asset medians by race from the HRS. Although the HRS is the best available public use dataset with both detailed information on assets and place of residence, it has three limitations for this purpose. First, the survey focuses on older people and we therefore analyze survey respondents ages 50-65. Second, it does not distinguish between checking and savings accounts. Third, the geographic variable corresponds to the Census division (state of residence is suppressed). The division which most closely corresponds to the three-state sample frame is the South Atlantic Census division; we calculate that 68 percent of the observations in this division in the HRS live in Florida or Georgia.

[17] One open question is the extent to which differences in permanent income (as opposed to annual income) by race can explain the racial wealth gap. See Barsky et al. (2002), Scholz and Levine (2003), Altonji and Doraszelski (2005), and Aliprantis, Carroll and Young (2019). The SCF's question about "income in a normal year" is an attempt to measure permanent income in the context of a cross-sectional survey.

income, consumption, liquid wealth, and race. The well-known and large racial economic gaps in the U.S. also appear in this new dataset. The available evidence indicates that differences in consumption smoothing by race will be similar to—or perhaps a lower bound on—national differences by race.

3 Consumption Smoothing by Race and Assets

We assess whether and how much consumption smoothing of typical labor income volatility varies by race and assets. We first analyze this relationship using Ordinary Least Squares (OLS). We then describe our identification strategy for estimating the causal impact of labor income variation on consumption using firm pay shocks. We report results for the pooled sample, separately by race, separately by assets, and separately by race controlling for assets. Finally, we report results from an alternative identification strategy that uses unemployment.

3.1 Ordinary Least Squares

For our descriptive analysis, we use a first-difference OLS specification

$$\Delta c_{it} = \alpha + \beta \Delta y_{it} + \varepsilon_{it} \tag{1}$$

where i indexes households, t indexes months, c is the log of monthly nondurable consumption, and y is the log of monthly labor income. The use of logs means that β can be interpreted as an elasticity of consumption with respect to income as in Blundell, Pistaferri and Preston (2008, henceforth BPP). In Section 4 we show that this elasticity is one of the sufficient statistics needed for measuring the welfare cost of income volatility. Table A-6 provides summary statistics for our analysis sample.

Equation (1) identifies β if

Assumption EO (Earnings Orthogonality)

$$E(\varepsilon_{it} | \Delta y_{it}) = 0$$

holds. This assumption says that income changes are strictly exogenous with respect to the unobserved error ε_{it}. This is the type of assumption that is made in the literature that uses a statistical decomposition of income shocks such as BPP. Meghir and Pistaferri (2011) describes this assumption as imposing that "the individuals and the econometrician have the same information set."

Changes in labor income and nondurable consumption are correlated. Table 1 column (1) reports an estimate for $\hat{\beta}$ of 0.08, which means that a 10 percent change in income is associated with a change in consumption of 0.8 percent.

This correlation between income and consumption is about twice as strong for Black households and 40 percent stronger for Hispanic households as it is for White households.

We estimate

$$\Delta c_{it} = \alpha + \beta_1 \Delta y_{it} + \beta_2 \Delta y_{it} \times Black_{it} + \beta_3 \Delta y_{it} \times Hispanic_{it} + Black_{it} + Hispanic_{it} + \varepsilon_{it}, \quad (2)$$

with White as the omitted category. Table 1 column (2) reports estimates for $\hat{\beta}_1$ of 0.057, $\hat{\beta}_2$ of 0.048, and for $\hat{\beta}_3$ of 0.023.

However, this correlation is hard to interpret because the $\hat{\beta}$ estimates from an OLS regression will not have a causal interpretation if Assumption EO does not hold. For example, labor supply changes can lead to a violation of Assumption EO because of reverse causality. $\hat{\beta}$ will be biased downward if the worker takes time away from paid work during a period of higher-than-normal consumption (such as a vacation or a health emergency) or biased upward if a worker increases her labor supply to finance a specific consumption need.

3.2 Firm Pay Shocks as Instrument for Labor Income

To measure the causal effect of typical income variation on consumption, we use variation in firm-level pay. This strategy avoids the contamination from labor supply discussed above. It is inspired by the labor economics model of firm effects pioneered by Abowd, Kramarz and Margolis (henceforth AKM, 1999) and more recently by a research design used in Koustas (2018). Consider the data-generating process

$$y_{it} = \alpha_i + \psi_{j(i,t),t} + \eta_{it}. \quad (3)$$

where α_i is an individual effect, $j(i,t)$ is worker i's firm in month t, $\psi_{j(i,t),t}$ is the firm effect at time t, and η_{it} is an idiosyncratic individual shock. Equation (3) follows Lachowska et al. (2023) in generalizing AKM to allow for *time-varying* firm effects ($\psi_{j(i,t),t}$), while most of the AKM literature assumes *time-invariant* firm effects ($\psi_{j(i,t)}$).[18]

In this section, we estimate a modified version of equation (3), investigate the existence of pre-trends and the persistence of firm pay shocks, discuss the sources of variation in firm pay shocks, describe how we estimate the causal effect of income on consumption using firm pay shocks, and defend the identifying assumptions needed to estimate this causal effect. All of this discussion defines y as pay per paycheck (log $\frac{Pay}{NumberOfChecks}$); in Section 3.3 we consider alternative definitions of y.

We find that there is an economically and statistically significant relationship between own pay and the change in coworker pay. We estimate

$$y_{i,t} - y_{i,t-1} = \phi + \rho \Delta y_{j(-i,t),t} + \nu_{it} \quad (4)$$

where

$$\Delta y_{j(-i,t),t} \equiv \frac{\sum_{i' \in j, i' \neq i} y_{i't} - y_{i',t-1}}{\sum_{i'} \mathbf{1}(i' \in j, i' \neq i)}$$

[18] Lachowska et al. (2023) seek to measure the stability of firm effects from one *year* to the next, while we allow firm effects to vary by *month* and use them to instrument for individual income.

is a leave-out mean that uses the coworkers at firm j of worker i. We use the change in coworkers' pay to avoid a mechanical correlation between the left and right sides of the equation.[19] To improve precision, coworker pay is estimated using everyone with a Chase bank account, not just the workers for whom we have matched race. Firm pay shocks estimated using coworkers have strong predictive power for individual pay, as shown in Figure 1a. Nearly every one of the binscatter points falls on the linear regression line, indicating that the relationship in means is almost exactly linear. We estimate ρ of 0.38 with a standard error of 0.01 (Table A-7 column 2), which means that a 10% increase in coworker pay per paycheck is associated, on average, with a 3.8% increase in own pay.

This strong relationship between changes in own pay and changes in firm pay is a robust finding across many different datasets. We re-estimate equation (4) in time clock data and in U.S. tax data, both of which are described in more detail in Section 2.1. Although the JPMCI estimate of the change in coworker pay relies on the subset of coworkers who also happen to have bank accounts with Chase, both the Homebase time clock data and the CWBH tax data capture the change in pay for the universe of coworkers at the firm. We report these estimates in Table A-8. In the time clock data, we estimate ρ of 0.40, which is very similar to the estimate from the JPMCI data. The tax data are only available at the quarterly level. We estimate ρ separately for the seven states in the tax data and find estimates of 0.42 to 0.61, as compared to 0.64 at the quarterly level in the JPMCI data. Other researchers have also found a strong first stage when estimating similar regressions: Lachowska et al. (2023) report a coefficient of 0.55 to 0.65 for earnings per hour using quarterly data while Koustas (2018) reports an F-statistic of 104 for earnings per fortnightly pay period (no coefficient or standard error reported). These strong relationships suggest that firm pay may be a useful instrument for own pay in any employer-employee matched dataset.

To illustrate the relationship between own pay and coworker pay, we regress leads and lags of the change in own pay on the change in coworker pay. Specifically, we estimate

$$y_{i,t+k} - y_{i,t-2} = \phi + \rho_k \Delta y_{j(-i,t),t} + \nu_{it}. \tag{5}$$

We choose the base period for income to be $k = -2$ so that we can study pay changes from $t-2$ to $t-1$; $\hat{\rho}_{-2}$ is therefore mechanically zero. Figure A-9a shows that an increase in coworker pay from $t-1$ to t is associated with a *decrease* in own pay from $t-2$ to $t-1$. This pattern of a negative correlation in first differences is the one that we would expect to

[19] There are two technical issues which have been found to be important in other contexts but are not important in our setting. First, we use all workers at firm j in months t and $t-1$, including workers who leave after $t-1$ or join at t. Because the monthly turnover rates at the firms are low, conditioning on stayers would have little impact on $\Delta y_{j(-i,t),t}$. Second, we could in principle use an empirical Bayes procedure to adjust for the fact that $\Delta y_{j(-i,t),t}$ is a noisier estimate of $\psi_{j(i,t),t}$ for small firms than large firms. However, in this context, such an adjustment would have little impact upon our estimate of ρ. The estimate for ρ changes little when we limit the sample to firms with above 50 employees. Because 73% of workers are employed by firms which meet this size threshold, the adjustment from empirical Bayes for $\Delta y_{j(-i,t),t}$ would affect only a small share of workers.

see if there are transitory shocks to $\psi_{j(i,t),t}$. The fact that transitory shocks imply a negative correlation of first differences has been documented for individual income since at least Hall and Mishkin (1982); our results simply show that this pattern extends to firm pay shocks. For example, if $\psi_{j,t} = 0.1$ and $\psi_{j,t'} = 0 \quad \forall t' \neq t$, a specification that did not control for lags would give $\Delta y_{j(-i,t),t} = 0.1$ and $\Delta y_{j(-i,t),t+1} = -0.1$. The negative income "shock" at $t+1$ would simply be the reversion of the positive innovation from the prior period. Thus, when analyzed unconditionally, changes in coworker pay capture times when month t firm pay was unusually high as well as times when month $t-1$ firm pay was unusually low. To extract just the innovations to ψ (and not the mechanical reversion of those innovations), we control for lags of the change in coworker pay.[20]

The relationship between leads and lags of own pay and coworker pay in our preferred specification suggests that this instrument captures transitory seasonal innovations to firm pay. Two patterns in Figure A-9b support this interpretation. First, the coefficient is of almost the same magnitude for $\hat{\rho}_0$, $\hat{\rho}_{-12}$, and $\hat{\rho}_{12}$, indicating either an annual bonus payment or a short-lived labor demand shock which recurs annually.[21] Second, the pattern of coefficients shows that firm pay shocks gradually dissipate over the following months. In addition, we show later in Section 3.4 that by applying panel data methods we can explicitly isolate the temporary component of the firm pay shock; this procedure yields a very similar estimate to our baseline specification, thereby suggesting that the baseline specification is already capturing a temporary shock. We also note that firm pay shocks may be predictable from the consumer's perspective and discuss interpretation in more detail below. Having established that firm pay shocks can be used to predict temporary shocks to earnings, we next ask "where do firm pay shocks come from?"

Based on the limited data that is available, it appears that firm pay shocks are a mix of changes to hours and changes to pay per hour in roughly equal proportions. The best available information that we are aware of to explain the link between pay and hours comes from Washington state tax data studied by Lachowska et al. (2022). They find that 73% of quarterly changes in pay at the firm level are explained by changes in average hours at the firm level. However, this 73% estimate is biased upward because many workers are paid fortnightly and many employers appear to report the number of work hours *paid* in a quarter rather than the number of hours actually *worked*.[22] Our best guess based on simulations which account

[20]We control for five lags in our baseline specification, though our key quantitative results continue to hold with other choices of lags and also without any lags.

[21]Supplementary analysis of time clock data suggests that the labor demand shock channel is quantitatively important. In those data—which by definition include only hours and never bonuses—we find similar estimates for $\hat{\rho}_0$ and $\hat{\rho}_{12}$.

[22]Lachowska, Mas and Woodbury (2022) notes that there is excess mass in the Washington quarterly hours distribution at 480, 520, and 560 hours per quarter. A worker who worked exactly 40 hours per week for 13 weeks would record 520 hours worked. However, a worker who is paid fortnightly will receive pay for 12 weeks of work (6 paychecks) or 14 weeks of work (7 paychecks), which can explain the excess mass at 480 and 560 hours per quarter. This 73% estimate is from personal correspondence with Lachowska and Woodbury based on new regressions using the same dataset from Lachowska et al. (2022).

for this time aggregation bias is that—after abstracting from paycheck timing—about half of firm pay shocks are from changes in hours and half are from changes in pay per hour. We return to the issue of paycheck timing in the JPMC data in Section 3.3.

The importance of these two types of pay fluctuations—hours and pay per hour—is consistent with three other datasets. First, among households with volatile income, the Survey of Household and Economic Decisionmaking (SHED) asks about specific reasons for volatility. We tabulate all the responses to this question in Table A-9. The most commonly-cited reason for income volatility is "irregular work schedule." A second source of volatility is changes in compensation such as bonuses and commissions. Grigsby, Hurst and Yildirmaz (2021) document using ADP and Lemieux, MacLeod and Parent (2009) document using the PSID that performance pay plays an important role in compensation for some workers. If we combine the fraction of workers who cite "bonuses" or "commissions" as reasons for volatility in the SHED, this is the second most commonly-cited reason for income volatility.

We measure the effect of transitory income changes on consumption using a linear two-stage least squares estimator that uses a modified equation (1) together with equation (5):

$$\Delta c_{it} = \alpha + \beta \Delta y_{it} + X_{it} + \varepsilon_{it} \tag{6}$$

$$\Delta y_{it} = \phi + \rho \Delta y_{j(-i,t),t} + X_{it} + \nu_{it}. \tag{7}$$

where X_{it} is a vector of control variables for five lags of the change in coworker pay. Appendix B.3 contains additional details on how we handle households with more than one worker and workers who have newly joined a firm. Our identification assumptions are that $\rho \neq 0$ (see Figure 1a), residualized $\Delta y_{j(-i,t),t}$ is an unbiased estimator of $\psi_{j(i,t),t}$, and the condition

Assumption FO (Firm-pay Orthogonality)

$$E(\varepsilon_{it} | \psi_{j(i,t),t}) = 0 \tag{8}$$

holds. The economic content of Assumption FO can be best understood as an independence assumption and an exclusion restriction.

The independence assumption is that the shocks to firm pay ψ are exogenous with respect to individual worker preferences over the level of consumption. It would be violated if, for example, workers who had a pre-existing preference for high consumption in a given month were to sort to firms which have high pay in that month. While we think this condition is plausible, we also show that our results are similar in an alternative specification (discussed in Section 3.3) which uses only variation in firm pay that is idiosyncratic relative to a firm's normal pay calendar, where this condition is even more likely to be satisfied.

The exclusion restriction is that firm pay shocks only affect consumption through their impact on a worker's income.[23] One way this restriction could be violated is if fluctuations

[23] One way that this assumption can fail is if a change in the average pay of all workers at a firm affects consumption through peer effects. De Giorgi, Frederiksen and Pistaferri (2020) document that idiosyncratic

in ψ affect hours worked and then hours worked directly affect consumption. Three findings documented later are consistent with the restriction: a similar estimate from a specification where the source of variation in ψ is not hours worked but instead idiosyncrasies of the pay schedule (discussed in Section 3.3), a finding that even consumption categories that are insensitive to hours worked respond to ψ (discussed in Section 3.4), and a finding of much smaller consumption responses for households with high assets (discussed in Section 3.6).

It is useful to emphasize that Assumption FO is silent regarding whether firm pay shocks are anticipated. Indeed, it allows both that workers may be able to predict pay changes before they occur, and also that once a shock is realized that workers may anticipate the dynamic effects on income in subsequent months. Assumption FO therefore follows the prior consumption literature in two ways. First, similar to the tax rebate literature (Johnson, Parker and Souleles 2006; Parker et al. 2013), it allows households to know in advance when their income is likely to change. As these papers note, if households do anticipate the income changes then the research design captures the causal effect of liquidity alone, rather than the combined effect of both news and liquidity. Second, similar to the semi-structural literature (as in BPP) this research design captures the contemporaneous spending response in one period to a transitory shock that may dissipate over several subsequent periods. In Appendix B.4 we discuss a more demanding assumption under which the coefficient we estimate identifies the consumption response to an unanticipated income shock lasting exactly one period. While there are reasons to believe this more demanding assumption is plausible, we do not require it for our main analysis.

Our interest in firm pay shocks is inspired by Koustas (2018) and we expand on this paper by providing additional econometric detail and considering a range of specifications.[24] In work subsequent to ours, Lachowska et al. (2023) uses a similar design to study the effect of quarterly firm wage shocks on labor supply at secondary jobs. The firm pay shocks strategy—which uses all changes in coworkers' earnings as an instrument—complements another strand in the consumption smoothing literature which instruments for earnings changes using changes in observable firm characteristics such as changes in union contracts (Shea 1995), value added (Fagereng, Guiso and Pistaferri 2017) and stock prices (Baker 2018).

Firm pay shocks enable estimation of consumption sensitivity under weaker assumptions

changes in the income of one peer (e.g an income shock to a "co-worker of the spouse of my co-worker") affects consumption via peer effects. One standard story motivating such a behavioral channel is a "keeping up with the Joneses" logic where my consumption choices are driven in part by a desire to keep up with the consumption of my higher-consuming peers. It is unclear if this story applies to the firm pay shocks we study, where nearly everyone in the firm experienced a similar income shock.

[24]More specifically, we expand on Koustas (2018) in four ways. First, we document that "raw" changes in coworker pay capture both the effect of innovations to firm pay and their predictable reversion. Second, we provide orthogonality conditions and discuss the identifying assumptions. Third, we report a nonparametric first stage (Figure 1a), a first stage regression coefficient, and the correlation of the instrument with leads and lags. This last exercise reveals that firm pay shocks primarily capture seasonal fluctuations in labor income. Fourth, we study a range of specifications, including ones which rely on changes in pay schedules (rather than changes in hours) and ones which isolate the idiosyncratic (non-seasonal) component of firm pay shocks, which we discuss in Section 3.3.

than the prior semi-structural literature. More specifically, Assumption FO is weaker than any economically reasonable version of Assumption EO, which identifies the OLS model discussed in Section 3.1. To understand how the two are related, note that first-differencing of equation (3) yields $\Delta y_{it} = \Delta\psi_{j(i,t),t} + \Delta\eta_{it}$ and a sufficient condition for Assumption EO to hold is

$$E(\varepsilon_{it}|\psi_{j(i,t),t}, \eta_{it}) = 0. \tag{9}$$

Indeed, it is hard for us to think of any economic model where $E(\varepsilon_{it}|\Delta\psi_{j(i,t),t} + \Delta\eta_{it}) = 0$ would hold but equation (9) would not hold. Equation (9) requires both Assumption FO—that firm pay is orthogonal to ε_{it}—and also that individual earnings component η_{it} is orthogonal to ε_{it}. This second assumption is not satisfied in a model with endogenous labor supply. For example, it will fail if individual consumption needs or tastes drive individual labor supply choices.

3.3 Alternative Definitions of the Firm Pay Instrument

While our baseline specification captures variation in firm-wide pay per paycheck, we also consider both broader and narrower definitions of firm pay as the instrument. These alternative definitions isolate the effect of different sources of labor income volatility.

The first specification uses total monthly pay to coworkers and is motivated by the observation that most workers experience some timing mismatch between the periodicity of income and the periodicity of expenses. This mismatch arises because many expenses are due once a month while about three-quarters of U.S. workers are paid either every week or every fortnight. In the modal month, a worker paid every fortnight will receive two paychecks; however, this worker will receive an extra third paycheck two months a year. The definition of the instrument in the baseline specification—which only uses variation in pay per paycheck—therefore understates the magnitude of firm-induced income volatility that the worker faces in solving their monthly consumption-savings problem. We note that at least three prior consumption smoothing papers have used research designs based on check timing (Baugh and Wang 2021; Zhang 2022; Baugh and Correia 2022). Because variation in this version of the instrument is primarily driven by the number of paychecks, it captures income variation that lasts only one month and is largely independent of hours worked or pay per hour. This definition of the instrument is not our leading specification due to the potential bias from time aggregation, but as we discuss below the bias from time aggregation appears to be small.[25]

The second specification isolates non-seasonal pay fluctuations. We do this by adding

[25] While some expenses are monthly, others are likely to be sub-monthly. If workers time their sub-monthly expenses to coincide with paycheck receipt, an instrument utilizing variation in the number of paychecks received in a month may lead to bias in the estimated consumption sensitivity. For example, a worker paid every week who follows a rule to go shopping every payday will have more income and more spending in months with five paydays, even if both their income and expenditure are actually smooth at the weekly level.

firm-by-calendar-month fixed effects $\upsilon_{j(i,t),m(t)}$ to equations (6) and (7) where $m(t)$ refers to the calendar month (e.g., January, February, etc.) associated with t. These fixed effects remove firm-specific seasonal fluctuations in income as a source of identifying variation. One benefit of this specification from an identification perspective is that it bolsters the case for the independence assumption in Assumption FO. It is unlikely that workers would sort to firms based on the timing of idiosyncratic future pay changes. On the other hand, this specification captures only a small fraction of workers' total income volatility and so is less representative of the kinds of fluctuations that workers typically face from month to month.

3.4 Causal Impact of Income on Consumption

We use the firm pay shock as an instrument to measure the causal effect of income on consumption using the instrumental variables (IV) specification in equations (6) and (7). Figure 1b shows a binscatter of the reduced-form and 1c shows the second stage. The relationship is precisely estimated. To provide a quantitative interpretation of the evidence in Figure 1, Table 1 column (3) reports an IV estimate of an elasticity of consumption to income of 0.21 ($\hat{\beta} = 0.21$). The IV estimate is more than twice as large as the OLS estimate of 0.08. The larger IV estimate might reflect the instrument capturing longer-lasting income shocks than OLS or omitted variables bias due to endogenous labor supply.

To put IV and OLS on a common footing in having both procedures isolate a temporary income shock, we use panel data methods from BPP. BPP, using a method which dates back to Hall and Mishkin (1982), isolates temporary shocks by instrumenting for contemporaneous income changes using the income change for one period forward. When we implement BPP's estimator in the JPMCI data using all (potentially endogenous) own income changes, we estimate a passthrough of 0.12 (Table A-10).[26] When we implement BPP's estimator using coworker pay changes, we estimate a passthrough of 0.18. Comparing the BPP estimator applied to own pay changes (0.12) versus coworker pay changes (0.18) is consistent with an endogenous labor supply channel where when workers reduce hours either voluntarily (e.g., vacation) or involuntarily (e.g., health shock) they also have higher-than-normal consumption. However, the lower estimate when using own pay changes is also consistent with serially uncorrelated measurement error in income. We are unable to distinguish between these two explanations.

We implement three tests which indicate that the consumption sensitivity we document is not driven by changes in time use or by high frequency substitution of spending across months, and is robust to alternative specifications of firm pay shocks. First, we find little evidence that spending fluctuations are being driven by work-related expenses. We define a spending category as work related if it exhibits a larger-than-median drop at retirement

[26]For comparison, BPP report a slightly smaller passthrough of temporary shocks of 0.05. Other potential sources of differences are that (a) there could be differences between the Panel Study of Income Dynamics (PSID) data and JPMCI data (e.g., income is measured with error in the PSID per Bound et al. (1994) and total consumption is imputed from food), and (b) our analysis is monthly whereas the analysis in BPP is yearly.

(Aguiar and Hurst 2013). We find that the spending response for non-work-related expenses is 0.20, nearly identical to the overall spending response (Table A-11). This is consistent with the view that changes in spending are driven by changes in income rather than changes in time use. Second, we cumulate the income and spending changes to firm pay shocks over one quarter rather than one month and find a similar elasticity (0.29 quarterly versus 0.21 monthly, see Table A-12), indicating a sustained impact of income changes on spending. This shows that households are not just responding to firm pay shocks by reallocating spending from one month to the next. Furthermore, it suggests that households are not simply dipping into their stockpiles while keeping true consumption smooth, since Baker, Johnson and Kueng (2021) shows that such inventory management happens at a time horizon shorter than one quarter.[27]

Third, we find estimates of the consumption smoothing elasticity $\hat{\beta}$ of 0.17-0.18 using the two alternative definitions of pay shocks from Section 3.3. Table A-13 column (3) uses total monthly pay as an instrument (and, because the IV estimate in column (3) is so similar to Table 1 column (3), there is little evidence of bias from time aggregation in this specification), while Table A-14 column (3) uses pay per paycheck with firm-by-calendar-month fixed effects. The next three sections examine heterogeneity in $\hat{\beta}$ by race, by asset, and by race controlling for assets. For all three specifications, we find less consumption smoothing by Black and Hispanic households, less consumption smoothing by low-asset households, and that the racial wealth gap explains most of the differences in consumption smoothing by race (also reported in Tables A-13 and A-14); our discussion therefore focuses on just the results from the baseline specification.

3.5 Consumption Smoothing by Race

Figures 2a and 2b provide non-parametric evidence that when faced with similarly sized labor income shocks, Black and Hispanic households increase or cut their consumption more than White households. To provide a quantitative interpretation of these estimates, we re-estimate equations (6) and (7), interacting Δy_{it} and $\Delta y_{j(-i,t),t}$ with $Black_{it}$ and $Hispanic_{it}$ and adding categorical dummies by race. This specification captures heterogeneity by race in the causal impact of income, not the causal impact of race itself, which is an ill-formed concept.

As part of this analysis, we reestimate the first stage equation separately by race. We report the first stage by race for each of our specifications in Table A-15. The first stage for total monthly pay is nearly identical for each racial group.[28] The first stage for pay per paycheck is weaker for Black and Hispanic workers. This appears to be driven by differences in income levels by race. Black and Hispanic workers have lower levels of income (apparent in Table A-6) and the first stage is weaker for low-income workers than high-income workers.

[27]Baker, Johnson and Kueng (2021) also shows that household stockpiling is mostly driven by sales, whose timing is determined by stores rather than by a consumer's individual financial position.

[28]This is not surprising since variation in this specification is mainly driven by paycheck timing and most workers at a firm are on the same pay schedule.

Indeed, when we re-estimate the first stage controlling for prior quarter income quartile, the first stage in pay per paycheck is similarly strong for each racial group (Figure A-10). We also investigate whether there is a difference in the persistence of the income shock between racial groups. Figure A-11 shows that the month-by-month dynamics of income after a firm pay shock are nearly identical for each of the three groups.

We find that the causal effect of an income shock on consumption is about twice as large for Black and Hispanic households as it is for White households. Table 1 column (4) estimates $\hat{\beta} \times Black$ of 0.18 and $\hat{\beta} \times Hispanic$ of 0.13, where $\hat{\beta}$ is 0.16, which captures the elasticity for the omitted race category of White.

3.6 Consumption Smoothing by Assets

As we discuss in the introduction, there are several mechanisms that might give rise to differences in consumption sensitivity by race. In this paper, we focus on just one candidate mechanism: differences in liquid assets. We focus on this variable because heterogeneity in consumption sensitivity by liquid assets is a central implication of many consumption models. This analysis follows in a tradition going back to Zeldes (1989) of measuring heterogeneity in consumption smoothing by asset holdings. We first discuss heterogeneity in consumption smoothing by assets and relate our estimates to those in the prior literature. Then, in Section 3.7, we investigate whether racial inequality in assets can account for racial inequality in consumption smoothing.

To analyze the role of liquid assets, we construct each household's financial buffer as the ratio of average checking balance to average nondurable consumption in the six months prior to the payroll shock. We normalize assets by nondurable consumption because the key state variable in models of household consumption is the ratio of assets to permanent income (which we proxy for using lagged consumption).[29]

Figure 3 shows non-parametric evidence that households with fewer liquid assets do less consumption smoothing. We stratify households into deciles by their checking asset buffer and re-construct Figure 1c for the bottom and top deciles. The slope is quite steep for the bottom decile (0.50 in Figure 3a) and shallow for the top decile (0.08 in Figure 3b). The fact that the slope varies so sharply with assets—and in particular that it is so flat for the high-asset households—is consistent with the exclusion restriction in Assumption FO that fluctuations in time use are not the main driver of the changes in spending that we observe. To compare our results to prior estimates, in the remainder of the section we report results in terms of the marginal propensity to consume (MPC) instead of the elasticity of consumption to income. We do this by scaling our elasticity estimates by the ratio of average consumption to average income.[30]

[29] Baker (2018) takes a similar approach to measuring assets; when measuring heterogeneity in consumption smoothing that paper stratifies households by the ratio of raw assets to income.

[30] Table A-16 shows the regression results and scaling factors by asset quartile. An alternative strategy would be to estimate equations (6) and (7) in levels instead of logs. When we implement this strategy we also find less consumption smoothing by Black and Hispanic households, less consumption smoothing

The extent of heterogeneity in consumption sensitivity by liquid assets has been difficult to precisely quantify in prior work using exogenous variation in windfall income. Figure 4 illustrates the statistical uncertainty of estimates in Parker et al. (2013), Fagereng, Holm and Natvik (2021), and Kueng (2018).[31] Both Parker et al. (2013) and Kueng (2018) have sufficient statistical uncertainty to be consistent with both an MPC-asset slope that is steep and no slope at all (e.g an MPC that is constant with respect to assets). The results from the JPMCI data show statistically precise estimates that the consumption response to firm pay shocks declines steeply as liquid wealth rises.

One limitation of both of our estimates and this prior work is that differences in assets can arise for many reasons. Thus, it may not be the case that causal variation in assets will generate the differences in consumption smoothing depicted in Figures 3 and 4. When high-asset households have lower MPCs than low-asset households, this could reflect a causal effect of access to assets on consumption smoothing or it could reflect a selection effect, if for example low-asset households have a weaker preference for consumption smoothing (Parker 2017). One way to assess whether these differences in liquid assets can causally generate the differences in consumption smoothing in the data is to evaluate a structural model which attempts to match both the liquid asset distribution and consumption sensitivity in the data. We pursue this path in Section 4.

3.7 Consumption Smoothing by Race and Assets

This subsection explores the link between our two prior results: can racial inequality in wealth account for racial inequality in consumption smoothing? As discussed in Section 2, Black and Hispanic households have fewer liquid assets than White households, even after controlling for income. These large racial wealth gaps are as apparent in our sample as they are in national data. Since we find a clear link between liquid wealth and consumption smoothing in Section 3.6, it is possible that racial wealth gaps contribute to the racial gaps in consumption smoothing that we document in Section 3.5.

by low-asset households, and that the racial wealth gap explains most of the differences in consumption smoothing by race. However, this is not our preferred strategy because a dollar-based specification upweights the consumption response of high-wage households (Yitzhaki 1996), whose shocks are larger in dollar terms (but not in percentage terms). This concern arises because our research design captures the average response to shocks of varying sizes. In contrast, previous strategies estimating shocks that are similar in dollar terms across groups do not have this bias in estimating dollar-based MPC measures.

[31] For tax rebates we use Parker et al. (2013) as the benchmark because the MPC heterogeneity estimates in Johnson, Parker and Souleles (2006) are less precisely estimated and the estimates in Parker (2017) apply to only expenditure categories captured in the Nielsen retail panel. Souleles (1999) shows clear heterogeneity by liquid assets in the consumption response to tax refunds. We cannot directly compare our point estimates to those in Souleles (1999) because that paper reports asset heterogeneity using a tax refund coefficient but not an MPC, and furthermore the refund coefficient for our comparable outcome (nondurable spending) is only reported for the low asset group. Jappelli and Pistaferri (2014) analyze heterogeneity in responses to questions about MPC hypotheticals and find estimates similar to Fagereng, Holm and Natvik (2021). Misra and Surico (2014) document heterogeneity in the MPC among people with high mortgage debt. BPP and Fisher et al. (2020) analyze MPC heterogeneity by total wealth, but not by liquid wealth, using the PSID. Both of those papers find that consumption sensitivity is lowest for high-asset households, which is consistent with our findings.

We use a simple specification to summarize the role of assets in consumption smoothing in columns (5)-(8) of Table 1. We reestimate columns (3) and (4), adding a linear control for the rank of the household's asset buffer ($Checking = AssetRank/N - 0.5$). This variable is scaled from -0.5 for the lowest asset household to 0.5 for the highest asset household. Thus, the coefficient on $\Delta y \times Checking$ captures the change in consumption smoothing from being the lowest-ranked household in terms of assets to being the highest-ranked household, while the coefficient on Δy captures the consumption response for the median household in terms of assets. Thus, column (5) implies that the consumption elasticity for the median asset household is 0.26 and that this elasticity ranges from 0.46 for the bottom asset decile of households to 0.06 for the top asset decile households. Although this linear specification is convenient for expositional purposes, it is useful to note that the estimates implied by this linear control are quite close to actual estimates for the bottom decile of 0.50 and top decile of 0.08.

Controlling for observed asset buffers reduces racial gaps in consumption smoothing by about 25 to 50 percent. Column (6) controls for checking account balance as a measure of asset buffer. The coefficients on $\Delta y \times Black$ and $\Delta y \times Hispanic$ both fall to about 0.10 (from 0.18 for Black and 0.13 for Hispanic in column (4)).

However, the estimates that use observed checking account balance understate the importance of assets in explaining racial differences in consumption smoothing. The problem is one of differential measurement error. The measurement error exists because the JPMCI data do not contain all of the assets that someone might use to smooth consumption. This measurement error differs by race because, as we show in Section 2.1, Black and Hispanic households disproportionately hold their assets in checking accounts. Thus, they have fewer total assets than White households, even after conditioning on a level of checking account balances.

To estimate how incorporating a broader measure of assets affects consumption smoothing, we control for imputed liquid assets. The imputation procedure is described in Section 2.1. It assumes that the relationship of checking account balances to liquid assets is the same by race in the matched JPMCI-voter sample (where we do not observe liquid assets) as in the SCF (where we do observe liquid assets). This is analogous to the procedure in Rothstein and Wozny (2013), which seeks to measure the Black-White test score gap while controlling for permanent income, but lacks a single dataset with both variables; the authors use one dataset with test scores and annual income, and a second supplementary dataset where they impute permanent income from annual income and race.

Indeed, we find that nearly all of the racial gaps in consumption smoothing are explained by controlling for differences in liquid assets. These results can be seen in column (8) of Table 1. Once we control for liquid assets, about 80 percent of the Black-White and Hispanic-White gaps in consumption sensitivity to income are eliminated. Furthermore, the gaps that remain are not statistically significant. This finding is depicted graphically in Figure 5a. This finding is robust to controlling instead for a narrower measure of liquid assets from the SCF summary file, which excludes credit card debt (unlike our baseline measure which follows Kaplan and Violante 2014), or the broader measure of total financial assets.

There are two related interpretations of these results. First, this finding suggests that when faced with the same financial circumstances (e.g., same income shocks, same financial buffer), households of all races react similarly. Second, the finding suggests that due to the racial wealth gap, which leads to systematic differences in financial circumstances, Black and Hispanic households are more vulnerable to income fluctuations than White households. We note that the first interpretation holds whether the relationship between assets and consumption smoothing is causal or is driven by selection. The second interpretation, however, relies on assets having a causal impact on consumption smoothing.

In both cases, a natural underlying question is why Black and Hispanic households have lower liquid wealth than White households. We do not have any new evidence on that question in this study, but prior studies have documented that racial gaps in assets have roots in the long-lasting effects of structural racism embedded within government and institutional policies and practices (Bayer, Ferreira and Ross 2018; Du Bois 1935; Lui et al. 2006; Collins et al. 2017; Kijakazi, Smith and Runes 2019; Rothstein 2017). These structural forces not only have a direct effect on wealth and wealth accumulation at a given point in time, but may also drive differences in the key determinants of wealth disparities over time and across generations, such as family transfers (e.g., Shapiro 2017; Chiteji and Hamilton 2002; McKernan et al. 2014), neighborhood conditions such as poverty rates, home values, delinquency rates, and access to banking (Chetty et al. 2019; Chernenko and Scharfstein 2021; Howell et al. 2022; Perry, Rothwell and Harshbarger 2018; Keys, Mahoney and Yang 2022; Stein and Yannelis 2020; Kermani and Wong 2022), geographic and financial barriers to human capital accumulation (Dobbie and Fryer Jr. 2011; Jackson and Reynolds 2013; Addo, Houle and Simon 2016), and racial segregation and discrimination in the labor market (e.g., Grodsky and Pager 2001; Bertrand and Mullainathan 2004). As such, we do not interpret our results as a horse race between liquid assets or race as the key driver of the differences in consumption smoothing that we document; rather, the differences in liquid assets may themselves be determined by historical factors that are also driven by racial discrimination.

3.8 Alternative Research Design: Consumption Sensitivity to Unemployment

While the firm pay shock design captures typical sources of labor income volatility *within* employment spells, income changes *between* employment spells are another important source of volatility for workers. This motivates a second research design studying unemployment. This is the same event studied in Ganong and Noel (2019). For this analysis we use a dataset covering unemployment spells between 2012 to 2018.[32]

[32]See our companion paper Farrell et al. (2020) for more discussion of this design and this dataset.

Figure 6a illustrates that the income shock around the onset of unemployment insurance receipt is much larger than the shock in our primary research design. Average monthly income drops by about 25 percent, and remains depressed for a full year following UI receipt. (In unreported results, we find that incomes recover substantially from month 12 to 24 after UI receipt.) Figure 6a also shows that the income path in the year after unemployment is similar for Black, Hispanic, and White households.[33]

The average elasticity of consumption with respect to income from unemployment is similar to the elasticity from typical labor income variation. We quantify this using a linear two-stage least squares specification:

$$c_{it} = \alpha + \beta y_{it} + \varepsilon_{it} \tag{10}$$

$$y_{it} = \phi + \rho Post_{it} + \nu_{it} \tag{11}$$

where $Post_{it}$ is the 12-month period from one month before receipt of the first UI check (when unemployment usually begins) to ten months after receipt of the first UI check. We estimate an elasticity of consumption to income of 0.18 ($\hat{\beta} = 0.18$, Table 2 column 1).

Faced with a similar-sized income shock from unemployment, Figure 6b shows that Black and Hispanic households cut spending more than White households. To evaluate this difference quantitatively, we reestimate equations (10) and (11), interacting y_{it} and $Post_{it}$ with $Race_i$ and adding a dummy $Race_i$. Table 2 column (2) shows that the elasticity is about 30 percent higher for Black and Hispanic households. This racial gap is smaller than the gap estimated in our prior research design. Our finding of a racial gap in consumption smoothing is qualitatively similar to but quantitatively smaller than the estimates in Patterson (2019) using the PSID.[34] The larger consumption loss during unemployment for Black households is consistent with the survey evidence in Davis and Krolikowski (2023) that Black workers would be more willing than White workers to accept pay cuts to avoid layoffs.

This difference in elasticities by race could reflect differences by race in assets, differences by race in expectations about the long-term path of income (beyond the time horizon we

[33]This might be surprising given the well-known fact that Black workers have longer average unemployment durations in the U.S. However, these patterns do not hold in our sample. Table A-17 shows that average unemployment durations for banked households in the three states are similar for Black and White workers (and lower for Hispanic workers). Figure 6a shows a similar income trajectory in the year after receipt of the first UI check. Thus, we conclude that in our sample frame—banked households in Florida, Georgia, and Louisiana from 2012 to 2018—the effect of insured job loss on household income in the year after job separation is similar by race.

[34]Patterson (2019) instruments for income using unemployment, calculates the elasticity of *food* consumption with respect to income in the PSID, and then imputes total consumption using the method of Blundell, Pistaferri and Preston (2008). This paper finds that Black households have an MPC that is roughly twice that of White households during unemployment. The larger average racial gap in the PSID may be attributable to two sources. First, unemployment is largely a permanent negative shock in the PSID, whereas in our dataset income eventually recovers. Second, there are differences in gender composition in the PSID. Patterson finds a large MPC gap between Black men and White men and essentially no gap between Black women and White women. The PSID data disproportionately capture men, reflecting the fact that the PSID data go back to 1969, when most prime earners were men. In contrast, Figure A-5b shows that more than half of Black households in the U.S. and in the JPMCI data are women.

can observe), or any of the other reasons that consumption smoothing might differ by race as discussed in the introduction.

However, just as with firm pay shocks, Figure 5b shows that liquid assets are the key channel for understanding racial differences in consumption smoothing during unemployment. Table 2 column (3) shows that assets are strongly correlated with consumption smoothing. Controlling for assets reduces or eliminates the racial gap in consumption smoothing. Table 2 column (4) controls for checking assets and column (6) controls for imputed liquid assets. In the latter column, the racial gap in consumption smoothing is economically small and statistically insignificant.

Across each of our research designs, we find less consumption smoothing by Black and Hispanic households, less consumption smoothing by low-asset households, and that the racial wealth gap explains most of the differences in consumption smoothing by race. These empirical findings suggest that income volatility will have particularly adverse consequences for households with few assets, such as Black and Hispanic households.

4 Welfare Costs of Volatility

We have documented that consumption is sensitive to typical labor income fluctuations, and that Black and Hispanic households are significantly more sensitive than White households. In this section we ask whether this sensitivity, and the racial gaps in this sensitivity, are economically meaningful. To evaluate this question, we calculate the cost of temporary income volatility implied by our reduced-form estimates and investigate how this cost differs by race. We combine our empirical estimates with a simple, widely used framework based on Lucas (1987). We find that temporary income volatility has a large welfare cost for all groups, and further that this cost is substantially larger for Black and Hispanic households than it is for White households. We then discuss how our finding that households are sensitive even to frequently occurring income fluctuations with large welfare costs helps to distinguish between alternative interpretations of prior consumption sensitivity estimates.

4.1 Framework for calculating welfare cost of income volatility

We specify a statistical model of consumption, in the spirit of Lucas (1987). Suppose that log income in each period is comprised of three orthogonal components: a lifecycle component (κ_t), a permanent idiosyncratic component (z_t), and a transitory idiosyncratic component (θ_t):

$$\begin{aligned} y_t &= \kappa_t + z_t + \theta_t \\ z_t &= z_{t-1} + \zeta_t, \end{aligned} \tag{12}$$

where θ_t and ζ_t are i.i.d. and normally distributed random variables with standard deviations σ_θ and σ_ζ, respectively. The firm component of income variation $\psi_{j,t}$ is now subsumed into the general transitory component θ_t. As in the empirical analysis, we do not require any assumption about the extent to which the transitory shocks θ_t are predictable. We assume $\mathbb{E}\left[e^{\theta}\right] = \mathbb{E}\left[e^{\zeta}\right] = 1$.

Let log consumption be a function of the different components of income:

$$c_t = \beta\theta_t + \tilde{c}_t(\kappa_t, z_t, \xi_t), \tag{13}$$

with an error term, ξ_t, orthogonal to (κ, z, θ). Kaplan and Violante (2010) show that a function where consumption is linear in temporary income shocks is a good approximation to predictions from a structural lifecycle model.

We now derive a formula, similar to that of Lucas (1987), to express the welfare gain from shutting off transitory income shocks. From (13) above, we can write the *level* of consumption as:

$$C_t = e^{\tilde{c}_t} e^{\beta\theta_t},$$

and define $\tilde{C}_t = e^{\tilde{c}_t}$ as the level of consumption, net of the impact of the transitory shocks. We can define the welfare gain λ as the percent increase in consumption in each period that would leave the household indifferent between facing transitory shocks or having them shut off. Assuming constant relative risk aversion (CRRA) preferences with risk parameter γ, we have:

$$\mathbb{E}\left[\sum_{t=1}^{T} \delta^t \frac{\left((1+\lambda)\, C_t\right)^{1-\gamma}}{1-\gamma}\right] \equiv \mathbb{E}\left[\sum_{t=1}^{T} \delta^t \frac{\tilde{C}_t^{1-\gamma}}{1-\gamma}\right]. \tag{14}$$

We show in Appendix C.1 that a closed form solution for λ exists and, after an approximation of the natural log function, we have:

$$\lambda \approxeq \left[\beta + \beta^2\left(\gamma - 1\right)\right] \frac{\sigma_\theta^2}{2}. \tag{15}$$

When $\beta = 1$, this result collapses to the Lucas (1987) result that $\lambda \approxeq \gamma \frac{\sigma_\theta^2}{2}$.

4.2 Welfare cost estimates

In order to apply this formula in our empirical setting, we need to choose a period length and we need values for β, γ, and σ_θ^2. We present results at both quarterly and monthly frequencies. Historically this type of calculation has most often been done at a relatively low frequency such as the quarter (e.g., Constantinides 2021), partly due to data availability and partly because it may be harder to measure the costs of inter-temporal substitution at a higher frequency. However, recent work has also performed this type of calculation at a monthly frequency (see e.g., Fuchs-Schundeln and Hassan 2016). For the elasticity β we use our firm pay shock estimates from columns (3) and (4) of Table 1 (monthly) and from columns (1) and (2) of Table A-12 (quarterly). We vary the values of γ between 1 and 4. Finally, we estimate the variance of transitory labor income shocks σ_θ^2 in the Chase data using the method of Carroll and Samwick (1997). We find monthly variances between 0.09

and 0.10 and quarterly variances between 0.19 and 0.22 (Table 3). Table A-18 shows that we find similar estimates of transitory variances in the SIPP.

Our estimates suggest that households face a substantial amount of income variation from month to month and quarter to quarter, consistent with prior work. For example, Hannagan and Morduch (2015) come to a similar conclusion analyzing data from the U.S. Financial Diaries. Their preferred volatility measure is the average coefficient of variation of total monthly income within a year. They measure an average monthly coefficient of variation of 39 percent among their sample of low and moderate-income households. We calculate that the equivalent statistic for all households is 37 percent in the Chase data and 34 percent in the SIPP. Furthermore, we find only moderate evidence of heterogeneity in income volatility by race at these frequencies. This conclusion is in line with Garber (2016), who measures volatility in the SIPP as the monthly percent change from a household's two-year average and finds only small differences between racial groups.

Combining these inputs with the formula in equation (15), Table 3 shows that the welfare gain from eliminating transitory income volatility is substantial: at least 1 percent overall and at least 0.7 percent for every racial group. These estimates may be an upper bound on the true cost of transitory fluctuations for two reasons. First, as is the case with any calculations within the standard Lucas (1987) framework, they assume that temporary income fluctuations do not directly affect utility (e.g., through changes in hours worked). Second, they assume that the elasticity of consumption with respect to income for all temporary labor income changes is the same as the Local Average Treatment Effect we estimate using temporary changes in firm pay. Nevertheless, relative to the benchmark view in Lucas (1987) that a cost of 0.5 percent of lifetime consumption is "large," these estimates suggest there could be a substantial welfare cost from transitory income volatility. Although Lucas (1987) conjectured that the welfare cost of transitory income volatility might be large if his formula were applied to individual volatility rather than aggregate volatility, we are unaware of any prior estimate of these costs.[35]

The exercise also uncovers a previously unstudied dimension of racial inequality. We see the largest welfare gain for Black households, and the smallest for White households. Looking at the case where $\gamma = 1$, the welfare gain is about two times larger for Black and Hispanic households than for White households. These differences grow modestly as we increase the value of γ.

4.3 Implications for alternative interpretations of consumption sensitivity

Our finding that households are sensitive even to frequently occurring labor income fluctuations with meaningful welfare costs helps to distinguish between alternative interpretations

[35] Hai, Krueger and Postlewaite (2020) and De Nardi, Fella and Paz-Pardo (2019) document a significant welfare gain from eliminating *all* income volatility (both temporary shocks and permanent shocks). In complementary work, Constantinides (2021) documents a significant welfare gain from eliminating all permanent idiosyncratic consumption shocks.

of prior consumption sensitivity estimates. A rich literature consistently documents that household consumption is sensitive to infrequent windfalls such as tax rebates (Johnson, Parker and Souleles 2006; Parker et al. 2013) and lottery winnings (Fagereng, Holm and Natvik 2021; Golosov et al. 2023). The consumption-smoothing literature has developed two alternative interpretations of this evidence, which have conflicting out-of-sample predictions for the frequently occurring income shocks that we study. The first interpretation rationalizes high windfall spending through two-asset consumption-savings models in which either preferences (Laibson, Repetto and Tobacman 2007) or technology (Kaplan and Violante 2014) lead to low liquidity as well as a substantial response to income fluctuations, even frequently occurring ones.

The second interpretation treats windfall spending as "near rational" behavior. The logic for this argument can be seen by examining the expression for welfare costs in equation (15). This cost is proportional to the *product* of sensitivity β and the variance of the shocks σ_θ^2. A line of argument dating back to Cochrane (1989) argues that when this product—which captures the upper bound on the welfare loss from setting consumption sub-optimally—is small, it is hard to learn much about individuals' true consumption model. By definition, σ_θ^2 is small for unusual windfalls; even if $\beta = 1$, $\gamma\sigma_\theta^2$ is still small. Indeed, Fuchs-Schundeln and Hassan (2016) conduct a meta-analysis of 18 studies and find that the welfare loss from sub-optimally consuming all of an unusual windfall in the month of receipt is less than one percent under commonly used preference assumptions. This line of reasoning suggests that sensitivity to windfalls may simply reflect a low-cost deviation from the true "rational" model. Under this interpretation, households *would* smooth frequently occurring income shocks, in line with Friedman's Permanent Income Hypothesis, because the cost of failing to smooth such shocks is large (i.e., σ_θ^2 is large).

In contrast, we find that households are sensitive even to frequently occurring labor income volatility. Although our study is similar to prior work in finding a sensitivity estimate (β) of around 0.2, our study differs in that it analyzes a type of income variation whose variance σ_θ^2 is much larger.[36] Finding a high degree of consumption sensitivity even in this setting thus favors the low-liquidity interpretation of consumption sensitivity over the near-rationality interpretation.

In Appendix C.2 we complement the calculations from our statistical model using the

[36]One other study which analyzes the consumption response to large pay fluctuations is Browning and Collado (2001, henceforth BC). Our finding that consumption is sensitive to frequently occurring labor income fluctuations is superficially at odds with the finding of no sensitivity in that paper; however, further inspection suggests the discrepancy between our estimates may not be large. BC compare the quarterly consumption of n = 341 workers who are paid in 12 monthly installments to n = 1877 workers who are paid in 14 installments, with double installments in June and December. The standard deviation of quarterly consumption in that survey is 29 percent, and the change in quarterly income from the installments is about 25 percent, so the design would only detect consumption fluctuations if the elasticity is 0.13 or greater (assuming a 95% confidence interval). The specification which most relies on predictable income variation in this paper yields an elasticity of 0.18. The elasticity for seasonal payments would only need to be a bit smaller in order for the design in BC to be unable to reject the null hypothesis of full consumption smoothing.

kind of modern two-asset structural model underlying the low-liquidity interpretation of consumption sensitivity to windfalls. We show that such a model can indeed match the sensitivity to typical labor income volatility that we document, and that the welfare cost of such volatility in this model is 0.8 percent, similar to the welfare cost we calculate using the simple statistical model. Moreover, while the statistical model isolates the effect of the passthrough of income volatility to consumption on welfare gains, a structural model also allows us to explicitly incorporate the racial wealth gap in our calculation. We estimate separate models for each racial group, choosing interest rates that allow us to match the elasticities estimated in Section 3 and race-specific wealth to income ratios, as measured in the SCF. We then use these estimated parameters to simulate behavior in the absence of transitory shocks, and compare the relative welfare gains across racial groups. Similar to the case of our statistical model, we find that the welfare gain of avoiding transitory income shocks is more than twice as large for Black and Hispanic households as it is for White households.

5 Conclusion

In this paper we study the consumption response to typical labor income shocks and investigate how these vary by wealth and race. We develop an instrument for labor income based on firm pay shocks. Household income volatility stems mostly from fluctuations in labor income and this research design therefore studies the sort of income fluctuations that households typically experience from month to month.

We have two main findings. First, we find an average elasticity of 0.21, with a much higher elasticity for low-liquidity households and close to zero elasticity for high-liquidity households. In a stylized model calibrated to our estimates, this degree of sensitivity implies that temporary income volatility has a large welfare cost for the average household, and especially large costs for low-liquidity households. Furthermore, this result helps to distinguish between competing explanations for consumption sensitivity. Our finding that consumption is sensitive even to frequently-occurring shocks with large welfare costs favors the low-liquidity interpretation of consumption sensitivity over the near-rationality interpretation.

Second, we use this instrument to study how wealth shapes racial inequality. Although an extensive body of work documents the long-term persistence of the racial wealth gap, less is known about its consequences on households' lives from month to month. We find that Black and Hispanic households are twice as sensitive to typical income shocks as White households. Nearly all of this difference is explained in a statistical sense by racial wealth inequality. Because of racial disparities in consumption smoothing, the welfare cost of temporary income volatility is twice as high for Black and Hispanic households than for White households

The results in this paper leave open questions about the sources of labor income volatility as well as policies to reduce income volatility. Although both survey and administrative data indicate that many workers experience substantial labor income volatility, it would be useful for researchers to describe and understand the sources of such volatility. Surveys with hours and pay are subject to recall bias as well as seam bias, while administrative data such as quarterly earnings records are subject to time aggregation bias. However, the increasing availability of administrative data from payroll processors should enable a more accurate picture of pay volatility.

Finally, the welfare loss from temporary income volatility suggests that policies and practices which mitigate volatility may have large welfare gains. This suggests that it would be useful for labor and public economists to learn more about the costs of reducing labor income volatility. One strand of such research involves learning about the costs to firms of practices that reduce pay fluctuations. Another strand involves measuring the willingness-to-pay for policies that mitigate the high degree of monthly income volatility, such as regulation of workers' schedules, short-term disability, paid sick leave, and worksharing. If the costs of such policies and practices are smaller than the welfare gains from reducing volatility, then volatility-reducing policies are likely to both decrease inequality and increase aggregate welfare.

References

Aaronson, Daniel, Daniel Hartley, and Bhashkar Mazumder. 2021. "The Effects of the 1930s HOLC "Redlining" Maps." *American Economic Journal: Economic Policy*, 13(4): 355–92.

Abowd, John M., Francis Kramarz, and David N. Margolis. 1999. "High Wage Workers and High Wage Firms." *Econometrica*, 67(2): 251–333.

Addo, Fenaba R., and Daniel T. Lichter. 2013. "Marriage, Marital History, and Black – White Wealth Differentials Among Older Women." *Journal of Marriage and Family*, 75(2): 342–362. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/jomf.12007.

Addo, Fenaba R., Jason N. Houle, and Daniel Simon. 2016. "Young, Black, and (Still) in the Red: Parental Wealth, Race, and Student Loan Debt." *Race and social problems*, 8(1): 64–76.

Aguiar, Mark A, Mark Bils, and Corina Boar. 2021. "Who Are the Hand-to-Mouth?" *Working Paper.*

Aguiar, Mark, and Erik Hurst. 2013. "Deconstructing Life Cycle Expenditure." *Journal of Political Economy*, 121(3): 437–492.

Akee, Randall, William Copeland, E. Jane Costello, John B Holbein, and Emilia Simeonova. 2018. "Family Income and the Intergenerational Transmission of Voting Behavior: Evidence from an Income Intervention." National Bureau of Economic Research Working Paper 24770. Series: Working Paper Series.

Aladangady, Aditya, Shifrah Aron-Dine, David Cashin, Wendy Dunn, Laura Feiveson, Paul Lengermann, Katherine Richard, and Claudia Sahm. Forthcoming. "Spending Responses to High-Frequency Shifts in Payment Timing: Evidence from the Earned Income Tax Credit." *American Economic Journal: Economic Policy.*

Aliprantis, Dionissi, Daniel Carroll, and Eric R. Young. 2019. "The Dynamics of the Racial Wealth Gap." Social Science Research Network SSRN Scholarly Paper ID 3467718, Rochester, NY.

Almaguer, Tomas. 2012. "Race, Racialization, and Latino Populations"." In *Racial Formation in the Twenty-First Century.* University of California Press.

Altonji, Joseph G., and Ulrich Doraszelski. 2005. "The Role of Permanent Income and Demographics in Black/White Differences in Wealth." *The Journal of Human Resources*, 40(1): 1–30.

Arellano, Manuel, Richard Blundell, Stèphane Bonhomme, and Jack Light. 2023. "Heterogeneity of Consumption Responses to Income Shocks in the Presence of Nonlinear Persistence." *Working Paper*.

Argyle, Bronson, Sasha Indarte, Benjamin Iverson, and Christopher Palmer. 2023. "Explaining Racial Disparities in Personal Bankruptcy Outcomes." *Working Paper*.

Attanasio, Orazio P., and Nicola Pavoni. 2011. "Risk Sharing in Private Information Models With Asset Accumulation: Explaining the Excess Smoothness of Consumption." *Econometrica*, 79(4): 1027–1068. Publisher: Econometric Society.

Avery, Robert B., Kenneth P. Brevoort, and Glenn B. Canner. 2009. "Credit Scoring and Its Effects on the Availability and Affordability of Credit." *Journal of Consumer Affairs*, 43(3): 516–537.

Baicker, Katherine, and Amy Finkelstein. 2019. "The Impact of Medicaid Expansion on Voter Participation: Evidence from the Oregon Health Insurance Experiment." *Quarterly Journal of Political Science*, 14(4): 383–400. Publisher: Now Publishers, Inc.

Baker, Scott R. 2018. "Debt and the Response to Household Income Shocks: Validation and Application of Linked Financial Account Data." *Journal of Political Economy*, 126(4): 1504–1557.

Baker, Scott R, Stephanie G Johnson, and Lorenz Kueng. 2021. "Financial Returns to Household Inventory Management." National Bureau of Economic Research Working Paper 27740.

Bank, Chase Consumer. 2019. "Chase Announces Major Branch Expansion in 2019."

Barsky, Robert, John Bound, Kerwin Kofi Charles, and Joseph P. Lupton. 2002. "Accounting for the Black-White Wealth Gap: A Nonparametric Approach." *Journal of the American Statistical Association*, 97(459): 663–673.

Bartlett, Robert, Adair Morse, Richard Stanton, and Nancy Wallace. 2022. "Consumer-lending discrimination in the FinTech Era." *Journal of Financial Economics*, 143(1): 30–56.

Baugh, Brian, and Filipe Correia. 2022. "Does paycheck frequency matter? Evidence from micro data." *Journal of Financial Economics*, 143(3): 1026–1042.

Baugh, Brian, and Jialan Wang. 2021. "When Is It Hard to Make Ends Meet?" *Working Paper*.

Baugh, Brian, Itzhak Ben-David, Hoonsuk Park, and Jonathan A. Parker. 2021. "Asymmetric Consumption Smoothing." *American Economic Review*, 111(1): 192–230.

Bayer, Patrick, Fernando Ferreira, and Stephen L. Ross. 2018. "What Drives Racial and Ethnic Differences in High-Cost Mortgages? The Role of High-Risk Lenders." *The Review of Financial Studies*, 31(1): 175–205. Publisher: Oxford Academic.

Beatty, Timothy K. M., Laura Blow, Thomas F. Crossley, and Cormac O'Dea. 2014. "Cash by any other name? Evidence on labeling from the UK Winter Fuel Payment." *Journal of Public Economics*, 118: 86–96.

Bertrand, Marianne, and Sendhil Mullainathan. 2004. "Are Emily and Greg More Employable Than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination." *American Economic Review*, 94(4): 991–1013.

Bhutta, Neil, and Aurel Hizmo. 2020. "Do Minorities Pay More for Mortgages?" *The Review of Financial Studies*, 34(2): 763–789.

Biggs, Andrew, and Glenn Springstead. 2008. "Alternate Measures of Replacement Rates for Social Security Benefits and Retirement Income." *Social Security Bulletin*, 68(2). Library Catalog: www.ssa.gov.

Blau, Francine D., and John W. Graham. 1990. "Black-White Differences in Wealth and Asset Composition." *The Quarterly Journal of Economics*, 105(2): 321–339. Publisher: Oxford University Press.

Blundell, Richard, Luigi Pistaferri, and Ian Preston. 2008. "Consumption Inequality and Partial Insurance." *American Economic Review*, 98(5): 1887–1921.

Board, Federal Reserve. 2017. "The Fed - Changes in U.S. Family Finances from 2013 to 2016: Evidence from the Survey of Consumer Finances."

Bodkin, Ronald. 1959. "Windfall Income and Consumption." *The American Economic Review*, 49(4): 602–614.

Bound, John, Charles Brown, Greg J. Duncan, and Willard L. Rodgers. 1994. "Evidence on the Validity of Cross-Sectional and Longitudinal Labor Market Data." *Journal of Labor Economics*, 12(3): 345–368. Publisher: [University of Chicago Press, Society of Labor Economists, NORC at the University of Chicago].

Browne, Robert S. 1974. "Wealth Distribution and its Impact on Minorities." *The Review of Black Political Economy*, 4(4): 27–38. Publisher: SAGE Publications Inc.

Browning, Martin, and M. Dolores Collado. 2001. "The Response of Expenditures to Anticipated Income Changes: Panel Data Estimates." *American Economic Review*, 91(3): 681–692.

Card, David, and Jesse Rothstein. 2007. "Racial segregation and the black-white test score gap." *Journal of Public Economics*, 91(11-12): 2158–2184.

Carr, Michael D, Robert A Moffitt, and Emily E Wiemers. 2020. "Reconciling Trends in Volatility: Evidence from the SIPP Survey and Administrative Data." National Bureau of Economic Research Working Paper 27672.

Carroll, Christopher D., and Andrew A. Samwick. 1997. "The nature of precautionary wealth." *Journal of Monetary Economics*, 40(1): 41–71. Publisher: Elsevier.

Cascio, Elizabeth U., and Ebonya Washington. 2014. "Valuing the Vote: The Redistribution of Voting Rights and State Funds following the Voting Rights Act of 1965." *The Quarterly Journal of Economics*, 129(1): 379–433. Publisher: Oxford Academic.

Center, Pew Hispanic. 2013. "Diverse Origins: The Nation's 14 Largest Hispanic-Origin Groups."

Charles, Kerwin Kofi, and Erik Hurst. 2002. "The Transition To Home Ownership And The Black-White Wealth Gap." *The Review of Economics and Statistics*, 84(2): 281–297. Publisher: MIT Press.

Charles, Kerwin Kofi, Erik Hurst, and Nikolai Roussanov. 2009. "Conspicuous Consumption and Race." *The Quarterly Journal of Economics*, 124(2): 425–467.

Chernenko, Sergey, and David S. Scharfstein. 2021. "Racial Disparities in the Paycheck Protection Program." *Working Paper*.

Chetty, Raj, Nathaniel Hendren, Maggie R Jones, and Sonya R Porter. 2019. "Race and Economic Opportunity in the United States: an Intergenerational Perspective." *The Quarterly Journal of Economics*, 135(2): 711–783.

Chiteji, N., and Darrick Hamilton. 2002. "Family connections and the black-white wealth gap among middle-class families." *The Review of Black Political Economy*, 30(1): 9–28. Publisher: Springer.

Choi, HwaJung, Kathleen McGarry, and Robert F. Schoeni. 2016. "Consumption and the extended family." *Economics Letters*, 140: 34–38.

Choukhmane, Taha, Jorge Colmenares, Cormac O'Dea, Jonathan Rothbaum, and Lawrence Schmidt. 2022. "Who Benefits from Retirement Saving Incentives in the U.S.? Evidence on Racial Gaps in Retirement Wealth Accumulation." *Working Paper*.

Cochrane, John H. 1989. "The Sensitivity of Tests of the Intertemporal Allocation of Consumption to Near-Rational Alternatives." *American Economic Review*, 79(3): 319–37.

Collins, Chuck, Dedrick Asante-Muhammed, Josh Hoxie, and Emanuel Nieves. 2017. "The Road to Zero Wealth."

Constantinides, George M. 2021. "Welfare Costs of Idiosyncratic and Aggregate Consumption Shocks." National Bureau of Economic Research Working Paper 29009.

Darity, William A., Jr. 2022. "Position and Possessions: Stratification Economics and Intergroup Inequality." *Journal of Economic Literature*, 60(2): 400–426.

Davis, Stephen J., and Pawel M. Krolikowski. 2023. "Sticky Wages on the Layoff Margin." *Working Paper.*

Deaton, Angus, and Christina Paxson. 1994. "Intertemporal Choice and Inequality." *Journal of Political Economy*, 102(3): 437–467. Publisher: University of Chicago Press.

De Giorgi, Giacomo, Anders Frederiksen, and Luigi Pistaferri. 2020. "Consumption Network Effects." *The Review of Economic Studies*, 87(1): 130–163.

De Nardi, Mariacristina, Giulio Fella, and Gonzalo Paz-Pardo. 2019. "Nonlinear Household Earnings Dynamics, Self-Insurance, and Welfare." *Journal of the European Economic Association.*

Derenoncourt, Ellora, Chi Hyun Kim, Moritz Kuhn, and Moritz Schularick. 2022. "When Is It Hard to Make Ends Meet?" *Working Paper.*

Dobbie, Will, Andres Liberman, Daniel Paravisini, and Vikram Pathania. 2021. "Measuring Bias in Consumer Lending." *The Review of Economic Studies*, 88(6): 2799–2832.

Dobbie, Will, and Roland G. Fryer Jr. 2011. "Are High-Quality Schools Enough to Increase Achievement among the Poor? Evidence from the Harlem Children's Zone." *American Economic Journal: Applied Economics*, 3(3): 158–187.

Druedahl, Jeppe, and Thomas Høgholm Jørgensen. 2017. "A general endogenous grid method for multi-dimensional models with non-convexities and constraints." *Journal of Economic Dynamics and Control*, 74: 87–107.

Du Bois, W. E. B. 1935. *Black reconstruction; an essay toward a history of the part which black folk played in the attempt to reconstruct democracy in America, 1860-1880.* Harcourt, Brace and Company.

Duesenberry, James S. 1949. *Income, saving, and the theory of consumer behavior.* Harvard University Press.

Epper, Thomas, Ernst Fehr, Helga Fehr-Duda, Claus Thustrup Kreiner, David Dreyer Lassen, Søren Leth-Petersen, and Gregers Nytoft Rasmussen. 2020. "Time Discounting and Wealth Inequality." *American Economic Review*, 110(4): 1177–1205.

Etheridge, Ben. 2015. "A test of the household income process using consumption and wealth data." *European Economic Review*, 78(C): 129–157. Publisher: Elsevier.

Fagereng, Andreas, Luigi Guiso, and Luigi Pistaferri. 2017. "Firm-Related Risk and Precautionary Saving Response." *American Economic Review*, 107(5): 393–97.

Fagereng, Andreas, Martin B. Holm, and Gisle J. Natvik. 2021. "MPC Heterogeneity and Household Balance Sheets." *American Economic Journal: Macroeconomics*, 13(4): 1–54.

Farrell, Diana, Fiona Greig, and Chenxi Yu. 2019. "Weathering Volatility 2.0A Monthly Stress Test to Guide Savings." JPMorgan Chase Instiute.

Farrell, Diana, Fiona Greig, Chris Wheat, Peter Ganong, Damon Jones, and Pascal Noel. 2020. "Racial Gaps in Financial Outcomes: Big Data Evidence." JPMorgan Chase Instiute.

Fisher, Jonathan D., David S. Johnson, Timothy M. Smeeding, and Jeffrey P. Thompson. 2020. "Estimating the marginal propensity to consume using the distributions of income, consumption, and wealth." *Journal of Macroeconomics*, 65: 103218.

Friedman, Milton. 1957. *Theory of the Consumption Function.* Princeton University Press.

Fuchs-Schundeln, N., and T.A. Hassan. 2016. "Chapter 12 - Natural Experiments in Macroeconomics." In . Vol. 2 of *Handbook of Macroeconomics*, , ed. John B. Taylor and Harald Uhlig, 923–1012. Elsevier.

Fuster, Andreas, Paul Goldsmith-Pinkham, Tarun Ramadorai, and Ansgar Walther. 2022. "Predictably Unequal? The Effects of Machine Learning on Credit Markets." *The Journal of Finance*, 77(1): 5–47.

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

Ganong, Peter, Fiona Greig, Pascal Noel, Daniel M. Sullivan, and Joseph Vavra. 2022. "Spending and Job-Finding Impacts of Expanded Unemployment Benefits: Evidence from Administrative Micro Data." *Working Paper.*

Garber, Judith. 2016. "Is There a Racial Volatility Gap?" *Aspen Institute.*

Golosov, Mikhail, Michael Graber, Magne Mogstad, and David Novgorodsky. 2023. "How Americans Respond to Idiosyncratic and Exogenous Changes in Household Wealth and Unearned Income." *Working Paper.*

Grigsby, John, Erik Hurst, and Ahu Yildirmaz. 2021. "Aggregate Nominal Wage Adjustments: New Evidence from Administrative Payroll Data." *American Economic Review*, 111(2): 428–71.

Grodsky, Eric, and Devah Pager. 2001. "The Structure of Disadvantage: Individual and Occupational Determinants of the Black-White Wage Gap." *American Sociological Review*, 66(4): 542–567. Publisher: [American Sociological Association, Sage Publications, Inc.].

Gross, Tal, Matthew J. Notowidigdo, and Jialan Wang. 2020. "The Marginal Propensity to Consume over the Business Cycle." *American Economic Journal: Macroeconomics.*

Guvenen, Fatih. 2007. "Learning Your Earning: Are Labor Income Shocks Really Very Persistent?" *American Economic Review*, 97(3): 687–712.

Guvenen, Fatih, and Anthony A. Smith. 2014. "Inferring Labor Income Risk and Partial Insurance From Economic Choices." *Econometrica*, 82(6): 2085–2129. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA9446.

Hai, Rong, Dirk Krueger, and Andrew Postlewaite. 2020. "On the welfare cost of consumption fluctuations in the presence of memorable goods." *Quantitative Economics*, 11(4): 1177–1214.

Hall, Robert E., and Frederic S. Mishkin. 1982. "The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data on Households." *Econometrica*, 50(2).

Hamilton, Darrick, and William A. Darity. 2017. "The Political Economy of Education, Financial Literacy, and the Racial Wealth Gap." *Federal Reserve Bank of St. Louis Review.*

Hamilton, Darrick, and William Darity. 2010. "Can 'Baby Bonds' Eliminate the Racial Wealth Gap in Putative Post-Racial America?" *The Review of Black Political Economy*, 37(3): 207–216.

Hannagan, Anthony, and Jonathan Morduch. 2015. "Income Gains and Month-to-Month Income Volatility: Household evidence from the US Financial Diaries." NYU Wagner Research Paper.

Hardy, Bradley, Jonathan Morduch, William Darity, and Darrick Hamilton. 2018. "Wealth Inequality, Income Volatility, and Race." *Working Paper.*

Hastings, Justine, and Jesse M. Shapiro. 2018. "How Are SNAP Benefits Spent? Evidence from a Retail Panel." *American Economic Review*, 108(12): 3493–3540.

Hausman, Joshua K. 2016. "Fiscal Policy and Economic Recovery: The Case of the 1936 Veterans' Bonus." *American Economic Review*, 106(4): 1100–1143.

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante. 2008. "Insurance and opportunities: A welfare analysis of labor market risk." *Journal of Monetary Economics*, 55(3): 501–525.

Hersh, Eitan D. 2015. *Hacking the Electorate: How Campaigns Perceive Voters.* Cambridge University Press.

Howell, Sabrina T, Theresa Kuchler, David Snitkof, Johannes Stroebel, and Jun Wong. 2022. "Lender Automation and Racial Disparities in Credit Access." National Bureau of Economic Research Working Paper 29364.

Jackson, Brandon A., and John R. Reynolds. 2013. "The Price of Opportunity: Race, Student Loan Debt, and College Achievement." *Sociological Inquiry*, 83(3): 335–368. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/soin.12012.

Jappelli, Tullio, and Luigi Pistaferri. 2014. "Fiscal Policy and MPC Heterogeneity." *American Economic Journal: Macroeconomics*, 6(4): 107–136.

Johnson, David S, Jonathan A Parker, and Nicholas S Souleles. 2006. "Household Expenditure and the Income Tax Rebates of 2001." *American Economic Review*, 96(5): 1589–1610.

Kaplan, Greg, and Giovanni L. Violante. 2010. "How Much Consumption Insurance beyond Self-Insurance?" *American Economic Journal: Macroeconomics*, 2(4): 53–87.

Kaplan, Greg, and Giovanni L Violante. 2014. "A Model of the Consumption Response to Fiscal Stimulus Payments." *Econometrica*, 82(4): 1199–1239.

Kaplan, Greg, and Giovanni L. Violante. 2022. "The Marginal Propensity to Consume in Heterogeneous Agent Models." *Annual Review of Economics*, 14(1): 747–775.

Kempf, Elisabeth, and Margarita Tsoutsoura. 2021. "Partisan Professionals: Evidence from Credit Rating Analysts." *The Journal of Finance*, 76(6): 2805–2856.

Kermani, Amir, and Francis Wong. 2022. "Racial Disparities in Housing Returns." *Working Paper.*

Keys, Benjamin J, Neale Mahoney, and Hanbin Yang. 2022. "What Determines Consumer Financial Distress? Place- and Person-Based Factors." *The Review of Financial Studies*, 36(1): 42–69.

Kijakazi, Kilolo, Karen E. Smith, and Charmaine Runes. 2019. "African American Economic Security and the Role of Social Security." Library Catalog: www.urban.org.

Koustas, Dmitri. 2018. "Consumption Insurance and Multiple Jobs : Evidence from Rideshare Drivers." *Working Paper.*

Kreinin, Mordechai E. 1961. "Windfall Income and Consumption: Additional Evidence." *The American Economic Review*, 51(3): 388–390.

Kueng, Lorenz. 2018. "Excess Sensitivity of High-Income Consumers." *The Quarterly Journal of Economics*, 133(4): 1693–1751.

Lachowska, Marta, Alexandre Mas, and Stephen A. Woodbury. 2022. "How reliable are administrative reports of paid work hours?" *Labour Economics*, 75: 102131.

Lachowska, Marta, Alexandre Mas, Raffaele Saggio, and Stephen A. Woodbury. 2022. "Wage Posting or Wage Bargaining? A Test Using Dual Jobholders." *Journal of Labor Economics*, 40(S1): S469–S493.

Lachowska, Marta, Alexandre Mas, Raffaele Saggio, and Stephen A. Woodbury. 2023. "Do firm effects drift? Evidence from Washington administrative data." *Journal of Econometrics*, 233(2): 375–395.

Laibson, David, Andrea Repetto, and Jeremy Tobacman. 2007. "Estimating Discount Functions with Consumption Choices over the Lifecycle." National Bureau of Economic Research Working Paper 13314.

Lemieux, Thomas, W. Bentley MacLeod, and Daniel Parent. 2009. "Performance Pay and Wage Inequality." *The Quarterly Journal of Economics*, 124(1): 1–49.

Liebler, Carolyn A., Sonya R. Porter, Leticia E. Fernandez, James M. Noon, and Sharon R. Ennis. 2017. "America's Churning Races: Race and Ethnicity Response Changes Between Census 2000 and the 2010 Census." *Demography*, 54(1): 259–284.

Lucas, Robert E. 1987. *Models of business cycles.* Oxford [Oxfordshire] ; New York, NY : B. Blackwell.

Lui, Meizhu, Barbara Robles, Betsy Leondar-Wright, Rose Brewer, and Rebecca Adamson. 2006. *The Color of Wealth: The Story Behind the U.S. Racial Wealth Divide.* . Copyright 2006 edition ed., New York:The New Press.

McKernan, Signe-Mary, Caroline Ratcliffe, Margaret Simms, and Sisi Zhang. 2014. "Do Racial Disparities in Private Transfers Help Explain the Racial Wealth Gap? New Evidence From Longitudinal Data." *Demography*, 51(3): 949–974.

Meghir, Costas, and Luigi Pistaferri. 2011. "Earnings, Consumption and Life Cycle Choices." Elsevier Handbook of Labor Economics.

Mendershausen, Horst. 1940. "Differences in Family Savings between Cities of Different Size and Location, Whites and Negroes." *The Review of Economics and Statistics*, 22(3): 122–137.

Misra, Kanishka, and Paolo Surico. 2014. "Consumption, Income Changes, and Heterogeneity: Evidence from Two Fiscal Stimulus Programs." *American Economic Journal: Macroeconomics*, 6(4): 84–106.

Mitchell, Olivia S., and John W. R. Phillips. 2006. "Social Security Replacement Rates for Alternative Earnings Benchmarks." Social Science Research Network SSRN Scholarly Paper ID 1094839, Rochester, NY.

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

Oliver, Melvin L., and Thomas M. Shapiro. 2006. *Black Wealth, White Wealth: A New Perspective on Racial Inequality.* Taylor & Francis. Google-Books-ID: 4ksJuX02DNwC.

Parker, Jonathan A. 2017. "Why Don't Households Smooth Consumption? Evidence from a $25 Million Experiment." *American Economic Journal: Macroeconomics*, 9(4): 153–183.

Parker, Jonathan A., Nicholas S. Souleles, David S. Johnson, and Robert McClelland. 2013. "Consumer spending and the economic stimulus payments of 2008." *American Economic Review*, 103(6): 2530–53.

Parker, Kim, Julia Menasce Horowitz, Rich Morin, and Mark Hugo Lopez. 2015. "Multiracial in America: Proud, Diverse and Growing in Numbers."

Patterson, Christina. 2019. "The Matching Multiplier and the Amplification of Recessions." Society for Economic Dynamics 95.

Pellegrini, Frank. 2001. "Breaking News, Analysis, Politics, Blogs, News Photos, Video, Tech Reviews." *Time.*

Perry, Andre, Jonathan Rothwell, and David Harshbarger. 2018. "The devaluation of assets in black neighborhoods." Library Catalog: www.brookings.edu.

Rothstein, Jesse, and Nathan Wozny. 2013. "Permanent Income and the Black-White Test Score Gap." *Journal of Human Resources*, 48(3): 510–544. Publisher: University of Wisconsin Press.

Rothstein, Richard. 2017. *The Color of Law: A Forgotten History of How Our Government Segregated America.* . 1 edition ed., New York ; London:Liveright.

Scholz, John Karl, and Kara Levine. 2003. "U.S. Black-White Wealth Inequality: A Survey." *Working Paper.*

Shapiro, Thomas M. 2017. *Toxic Inequality: How America's Wealth Gap Destroys Mobility, Deepens the Racial Divide, and Threatens Our Future.* . 1 edition ed., New York:Basic Books.

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

Smith, James P. 1995. "Racial and Ethnic Differences in Wealth in the Health and Retirement Study." *The Journal of Human Resources*, 30: S158–S183. Publisher: [University of Wisconsin Press, Board of Regents of the University of Wisconsin System].

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

Squires, Munir. 2016. "Kinship taxation as a constraint on microenterprise growth." PhD diss. The London School of Economics and Political Science (LSE).

Stein, Luke C D, and Constantine Yannelis. 2020. "Financial Inclusion, Human Capital, and Wealth Accumulation: Evidence from the Freedman's Savings Bank." *The Review of Financial Studies*, 33(11): 5333–5377.

Thaler, Richard. 1985. "Mental Accounting and Consumer Choice." *Marketing Science*, 4(3): 199–214.

Thaler, Richard H. 1990. "Anomalies: Saving, Fungibility, and Mental Accounts." *Journal of Economic Perspectives*, 4(1): 193–205.

Tobin, James. 1951. *Relative Income, Absolute Income and Saving.* University of Cambridge.

Toney, Jermaine, and Cassandra L. Robertson. 2021. "Intergenerational Economic Mobility and the Racial Wealth Gap." *AEA Papers and Proceedings*, 111: 206–10.

Townsend, Robert M. 1994. "Risk and Insurance in Village India." *Econometrica*, 62(3): 539–591. Publisher: [Wiley, Econometric Society].

Yitzhaki, Shlomo. 1996. "On Using Linear Regressions in Welfare Economics." *Journal of Business & Economic Statistics*, 14(4): 478–486.

Zedlewski, Sheila R, and Austin Nichols. 2012. "What Happens to Families' Income and Poverty after Unemployment?" Urban Institute.

Zeldes, Stephen. 1989. "Consumption and Liquidity Constraints: An Empirical Investigation." *Journal of Political Economy*, 97(2): 305–46.

Zhang, C. Yiwei. 2022. "Consumption Responses to Pay Frequency: Evidence From 'Extra' Paychecks." *Working Paper.*

Figure 1: Consumption Sensitivity to Labor Income Fluctuations

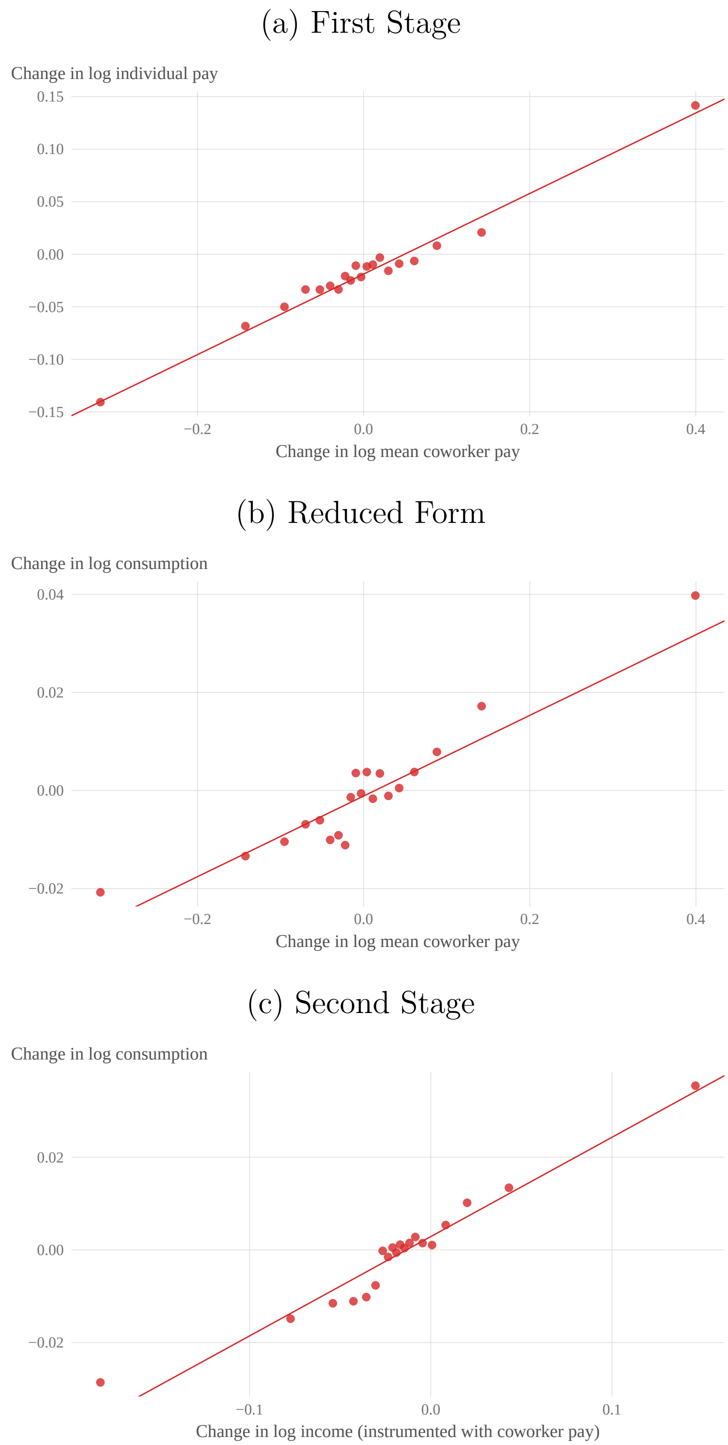

Note: The top two panels in this figure show the effect of changes in firm pay on labor income (panel (a)) and nondurable consumption (panel (b)) in the matched Chase-voter data. The x-axis shows a dot for each vingtile of the change in log mean pay of all other workers at the same firm. The y-axis shows conditional means of the change in log labor income and the change in log nondurable consumption. Panel (c) shows an instrumental variables interpretation of the first two panels, where the x-axis is the change in labor income (instrumented using the change in pay of all other workers at the same firm) and the y-axis is the change in nondurable consumption. The slope in panel (c) is the elasticity of consumption with respect to income, which is $\hat{\beta}$ from equation (6). See Section 3.4 for details.

Figure 2: Consumption Sensitivity by Race

(a) Black and White

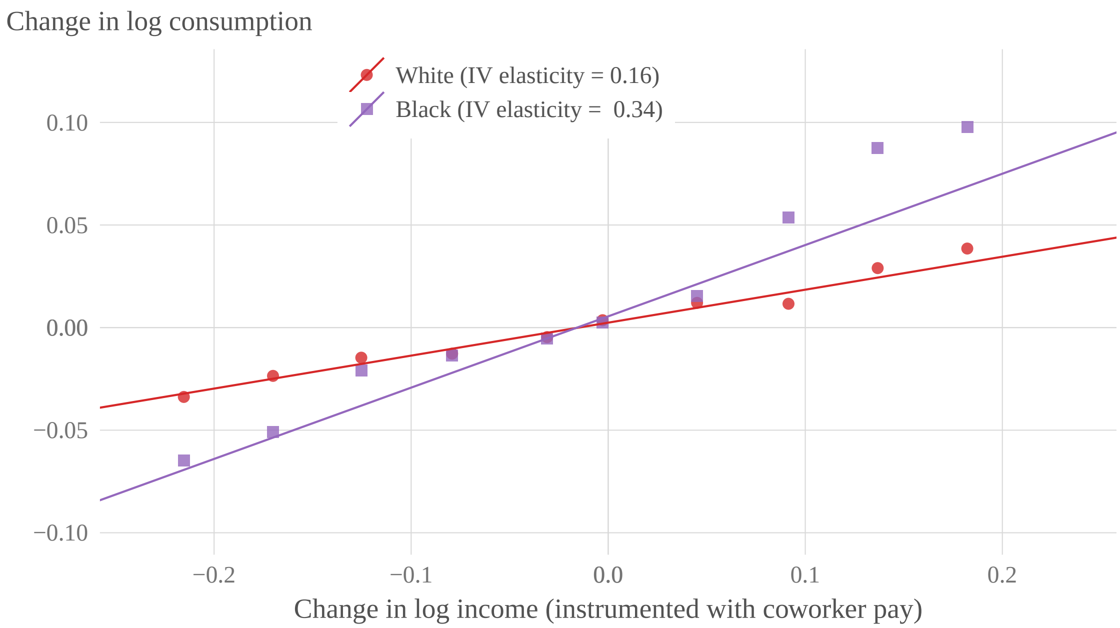

(b) Hispanic and White

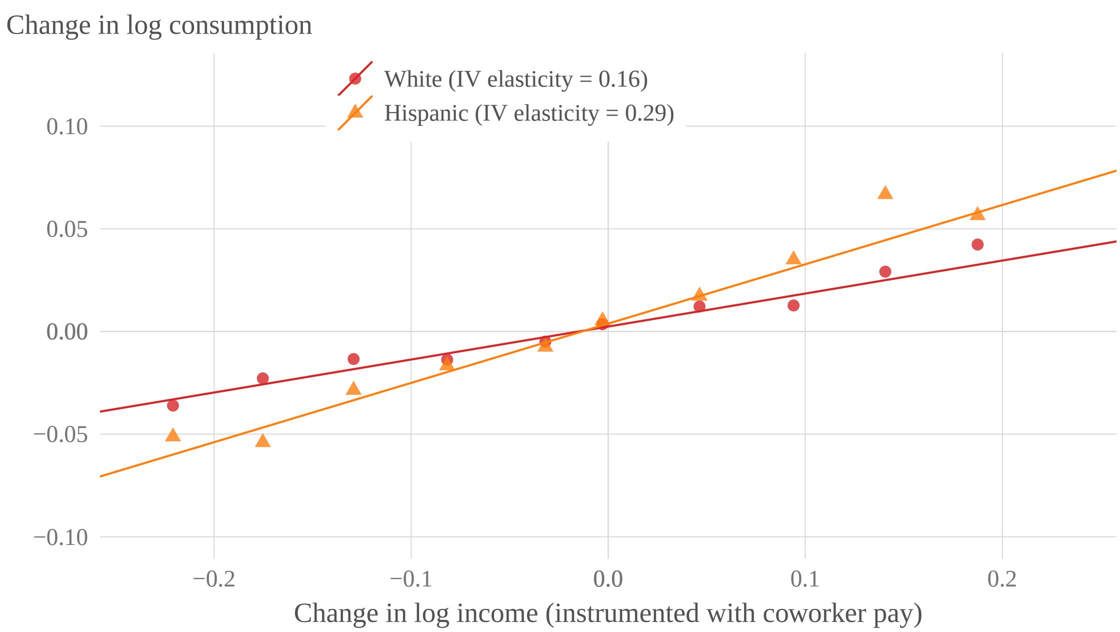

Note: This figure shows the effect of changes in labor income on nondurable consumption by race in the matched Chase-voter data. The x-axis stratifies the sample into 10 equally spaced bins of the predicted change in log own pay, where the omitted instrument is the change in pay of all other workers at the same firm. The y-axis shows conditional means of the change in log nondurable consumption by race. This is the second stage plot from Figure 1c, shown separately by race. See Section 3.5 for details.

Figure 3: Consumption Sensitivity by Asset Buffer Decile

(a) Lowest Decile

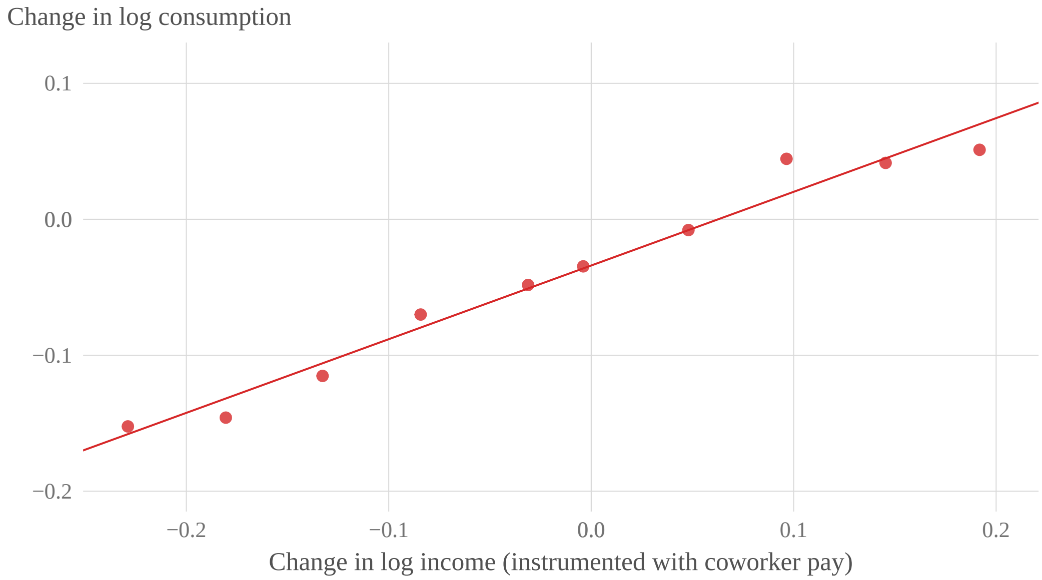

(b) Highest Decile

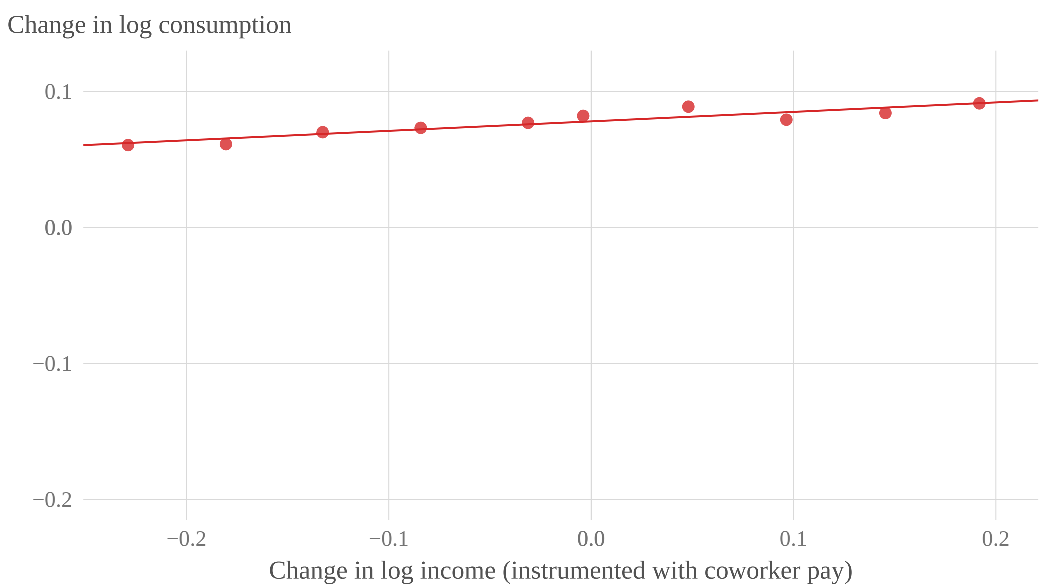

Note: This figure shows the effect of changes in labor income on nondurable consumption by asset buffer in the matched Chase-voter data. The x-axis stratifies the sample into 10 equally spaced bins of the predicted change in log own pay, where the omitted instrument is the change in pay of all other workers at the same firm. The y-axis shows conditional means of the change in log nondurable consumption separately by asset buffer decile. This is the second stage plot from Figure 1c, shown separately by asset buffer decile. We construct asset buffer deciles as the ratio of lagged assets to lagged nondurable consumption. Thus, households that are classified as high assets are a mix of households with permanently high assets and households with temporarily low consumption in prior months. The inclusion of this second group means that Figure 3b shows increasing consumption on average, across the distribution of pay shocks. The same pattern is also apparent in Figure 3a where some households are classified as having low assets because they had temporarily high consumption in prior months and therefore the group as a whole has falling consumption on average. See Section 3.6 for details.

Figure 4: Marginal Propensity to Consume by Asset Buffer

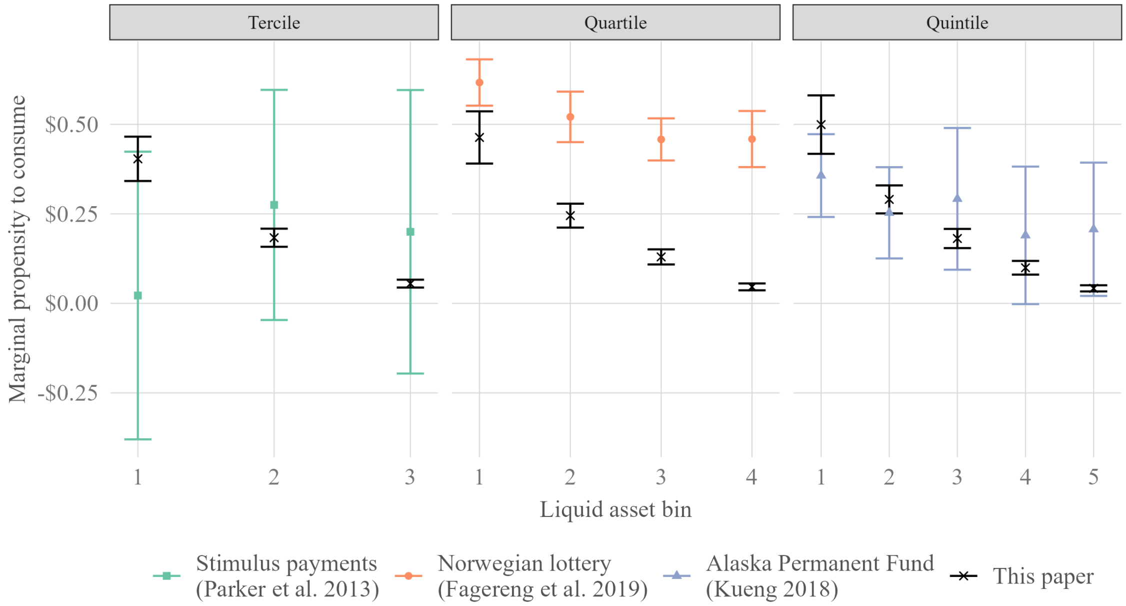

Note: This figure compares the estimates of heterogeneity by assets in the passthrough of income shocks to consumption. Parker et al. (2013), Fagereng, Holm and Natvik (2021), and Kueng (2018) use terciles, quartiles, and quintiles respectively. To enable comparability with these prior papers, we calculate the marginal propensity to consume (instead of the elasticity of consumption to income) using their respective bin cutoffs. Our paper, Parker et al. (2013), and Kueng (2018) measure the MPC on nondurables. Fagereng, Holm and Natvik (2021) measures the MPC on total consumption. See Section 3.6 for details.

Figure 5: Racial Inequality in Consumption Smoothing and Role of Assets

(a) Firm Pay Shocks

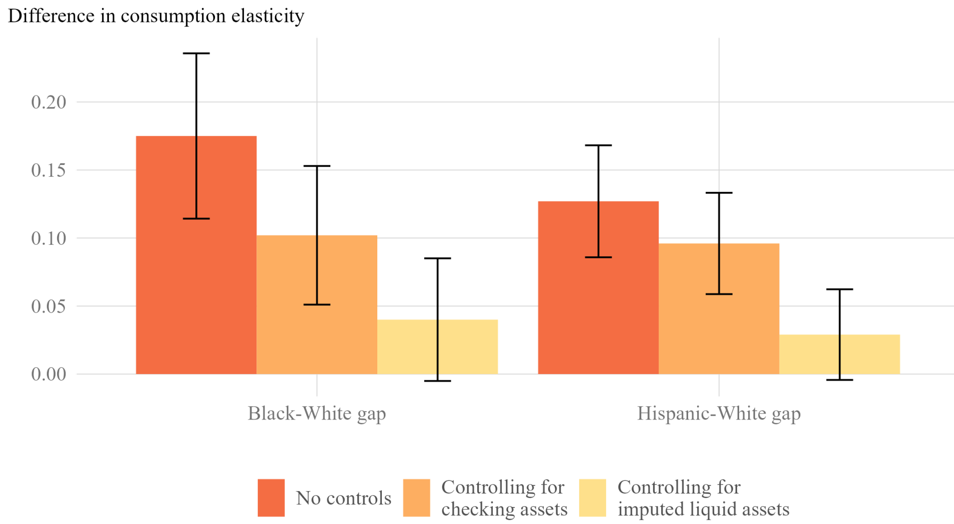

(b) Unemployment

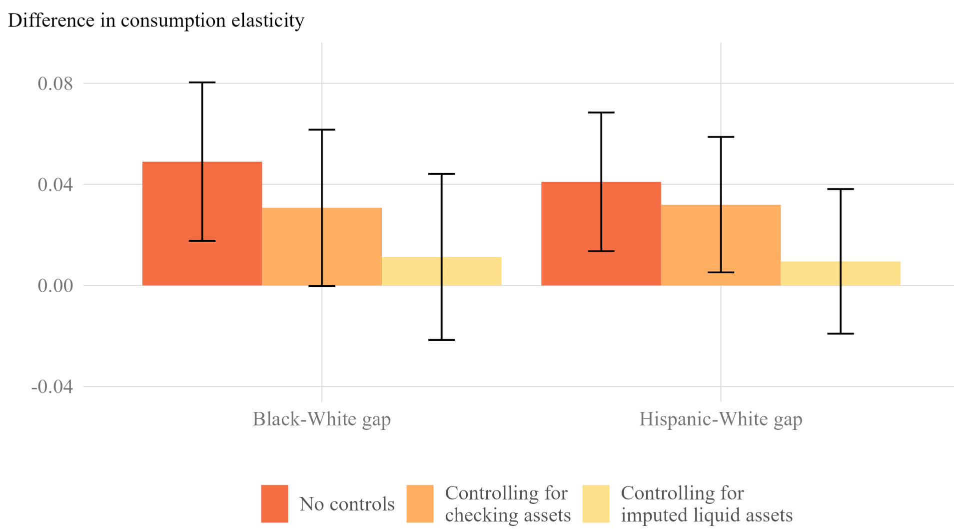

Note: This figure illustrates regression coefficients from Table 1 and 2. Within each race, the figure shows three bars. The left-most bar shows the difference in the elasticity between Black or Hispanic households and White households. The middle bar shows the difference when controlling for observed checking account balance. The right-most bar shows the difference when controlling for liquid assets, which are imputed using checking account balance and race. This imputation addresses the problem that for a given level of checking assets, Black and Hispanic households have fewer liquid assets than White households. See Section 3.7 for details on firm pay fluctuations and Section 3.8 for details on unemployment.

Figure 6: Effect of Unemployment on Income and Consumption by Race

(a) Labor Income + Unemployment Insurance

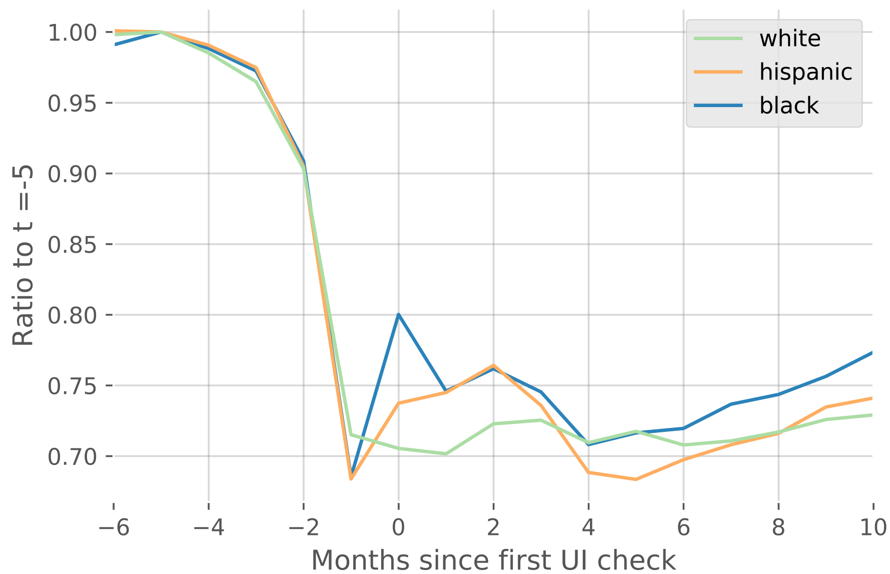

(b) Nondurable Spending

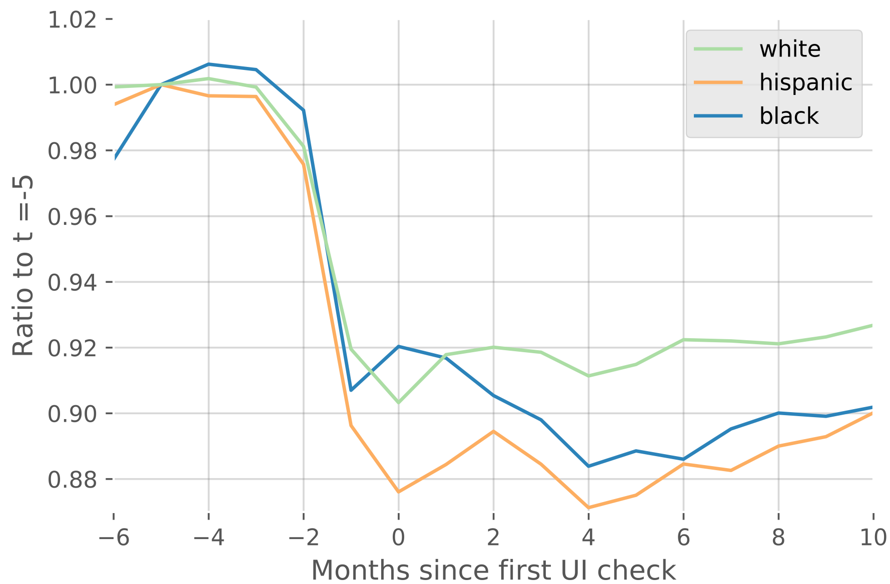

Note: This figure shows the evolution of income and consumption around unemployment in the matched Chase-voter data. Unemployment is measured using the household's first receipt of an unemployment insurance (UI) check. See Section 3.8 for details.

Table 1: Impact of Income on Consumption by Assets and Race

	Dependent Variable: Δ Log Non Durable Consumption							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δ Log Income	0.075	0.057	0.214	0.163	0.258	0.220	0.268	0.252
	(0.004)	(0.003)	(0.017)	(0.011)	(0.019)	(0.013)	(0.020)	(0.015)
(Δ Log Income) x Black		0.048		0.175		0.102		0.040
		(0.003)		(0.031)		(0.026)		(0.023)
(Δ Log Income) x Hispanic		0.023		0.127		0.096		0.029
		(0.003)		(0.021)		(0.019)		(0.017)
(Δ Log Income) x Checking					−0.440	−0.404		
					(0.039)	(0.032)		
(Δ Log Income) x Liquid(Imputed)							−0.476	−0.448
							(0.046)	(0.036)
OLS/IV	OLS	OLS	IV	IV	IV	IV	IV	IV
Black and Hispanic Dummies		Yes		Yes		Yes		Yes
Asset Rank Control					Yes	Yes	Yes	Yes
Observations	25,774,028	25,774,028	20,095,473	20,095,473	20,095,473	20,095,473	20,095,473	20,095,473
Adjusted R^2	0.004	0.004	−0.001	−0.004	−0.006	−0.008	−0.008	−0.008

Note: This table shows estimates of the elasticity of consumption with respect to income ($\hat{\beta}$). Columns (1) and (2) show OLS estimates of the effect of income on consumption using equations (1) and (2) respectively. Columns (3)-(8) show IV estimates using equations (7) and (6). Standard errors are clustered at the firm level. Columns (5) and (6) control for a narrow measure of assets: checking account balance. Asset variables are parameterized as $AssetRank/N - 0.5$, so the variable is scaled from -0.5 for the lowest asset household to 0.5 for the highest asset household. Columns (7) and (8) control for liquid assets. Liquid assets are imputed using checking account balance and race. The IV specifications control for five lags of the change in coworker pay. See Section 3.7 for details.

Table 2: Impact of Unemployment on Consumption by Assets and Race

	Dependent Variable: Δ Log Non Durable Consumption					
	(1)	(2)	(3)	(4)	(5)	(6)
Δ Log Income	0.176	0.153	0.177	0.160	0.178	0.171
	(0.006)	(0.008)	(0.006)	(0.009)	(0.006)	(0.009)
(Δ Log Income) x Black		0.049		0.031		0.011
		(0.016)		(0.016)		(0.017)
(Δ Log Income) x Hispanic		0.041		0.032		0.010
		(0.014)		(0.014)		(0.015)
(Δ Log Income) x Checking			−0.111	−0.106		
			(0.020)	(0.021)		
(Δ Log Income) x Liquid(Imputed)					−0.113	−0.116
					(0.021)	(0.023)
OLS/IV	IV	IV	IV	IV	IV	IV
Black and Hispanic Dummies		Yes		Yes		Yes
Asset Rank Control			Yes	Yes	Yes	Yes
Observations	65,811	65,811	65,811	65,811	65,811	65,811
Adjusted R^2	0.170	0.179	0.184	0.198	0.177	0.198

Note: This table shows IV estimates of the elasticity of consumption with respect to income using the unemployment design in equations (10) and (11). Columns (3) and (4) control for a narrow measure of assets: checking account balance. Asset variables are parameterized as $AssetRank/N - 0.5$, so the variable is scaled from -0.5 for the lowest asset household to 0.5 for the highest asset household. Columns (5) and (6) control for liquid assets. Liquid assets are imputed using checking account balance and race. See Section 3.8 for details.

Table 3: Welfare Gain of Eliminating Transitory Shocks

	Consumption elasticity of transitory shocks: β	transitory variance: σ_θ^2	Welfare gain: $\lambda \approxeq \left[\beta + \beta^2 (\gamma - 1)\right] \frac{\sigma_\theta^2}{2}$ Coefficient of relative risk aversion $\gamma = 1$	$\gamma = 2$	$\gamma = 3$	$\gamma = 4$
Monthly						
Black	0.338	0.10	1.75%	2.33%	2.92%	3.51%
Hispanic	0.290	0.10	1.44%	1.85%	2.27%	2.68%
White	0.163	0.09	0.72%	0.83%	0.95%	1.07%
Overall	0.214	0.09	1.01%	1.22%	1.44%	1.65%
Quarterly						
Black	0.396	0.21	4.20%	5.86%	7.52%	9.18%
Hispanic	0.343	0.22	3.82%	5.14%	6.45%	7.76%
White	0.224	0.19	2.07%	2.54%	3.00%	3.46%
Overall	0.294	0.20	2.95%	3.82%	4.69%	5.56%

Note: This table conducts an exercise in the spirit of Lucas (1987) by calculating the welfare gain from eliminating transitory income shocks. The welfare gain is calculated using a Taylor approximation in equation (15) and is expressed as a fraction of lifetime consumption. The transitory elasticities β are from columns (3) and (4) of Table1 (monthly) and columns (1) and (2) of Table A-12 (quarterly). The variance of transitory shocks σ_θ^2 are calculated using the method of Carroll and Samwick (1997). See Section 4 for details.

Wealth, Race, and Consumption Smoothing of Typical Income Shocks – Online Appendix

Peter Ganong, Damon Jones, Pascal Noel, Diana Farrell, Fiona Greig, and Chris Wheat

Contents

A Appendix Figures and Tables **1**

B Data Appendix **27**

B.1 Matching algorithm 27

B.2 Quality of matches 27

B.3 IV regression details 28

B.4 Orthogonality conditions 28

C Notes on welfare calculations **30**

C.1 Derivation of formula for statistical welfare gain λ 30

C.2 Structural Model 31

A Appendix Figures and Tables

Figure A-1: Assets and Checking Account Balances

(a) Financial Assets

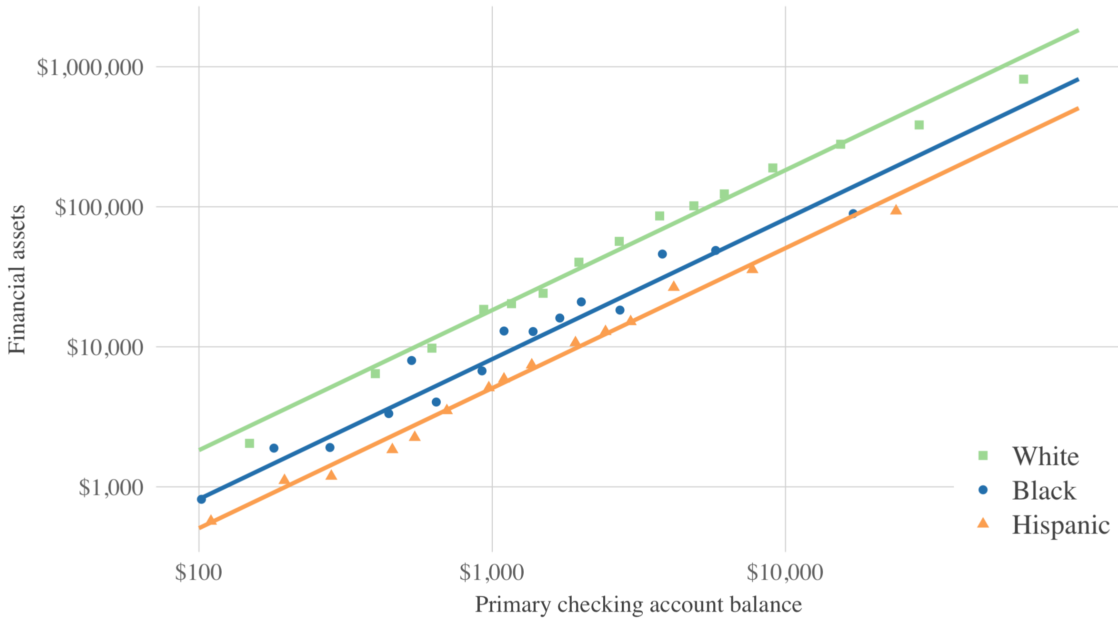

(b) Liquid Assets

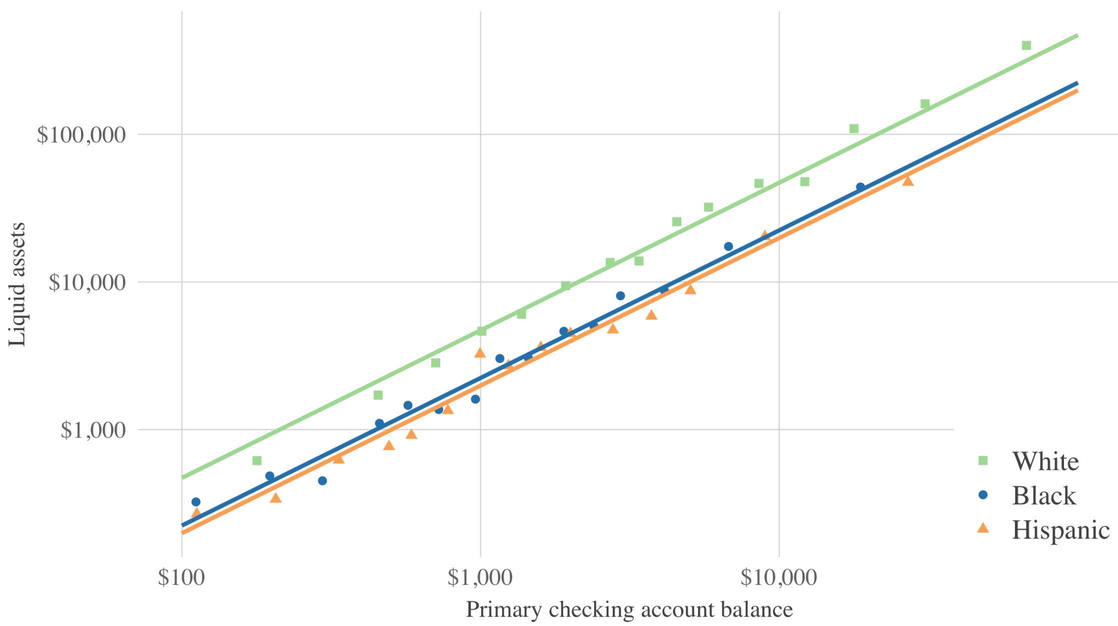

Note: We define primary checking account balance as the sum of balances at the institution that respondents "use the most" for checking. We stratify the sample into 20 equally-sized groups by race and show mean assets for each bin.
Source: 2010, 2013 and 2016 Survey of Consumer Finances

Figure A-2: Availability of Public Race Data and Chase Bank Presence

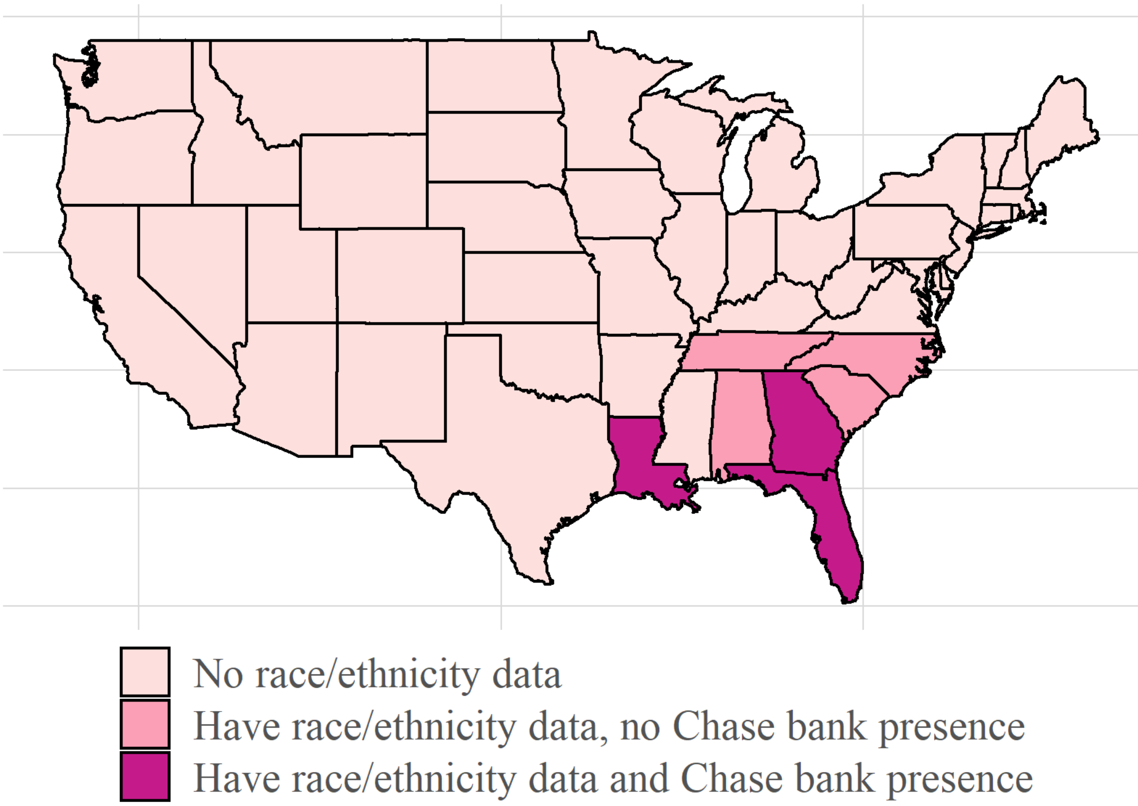

Source: Bank presence is measured using the bank's 2017 footprint (Chase Consumer Bank 2019). The set of eight states with race and Hispanicity data is from Hersh (2015). Pennsylvania's voter registration form also asks for race, but the variable is rarely populated in the voter file.

Figure A-3: Share of Voters Matched

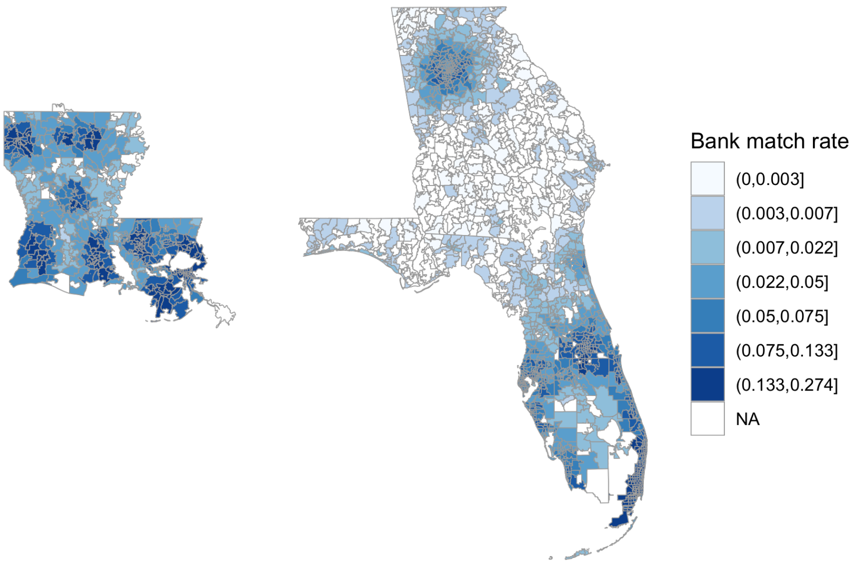

Note: This figure shows the share of registered voters in each zip code that have been matched to Chase bank accounts.

Figure A-4: Local Race and Hispanicity Distributions

(a) Share Black

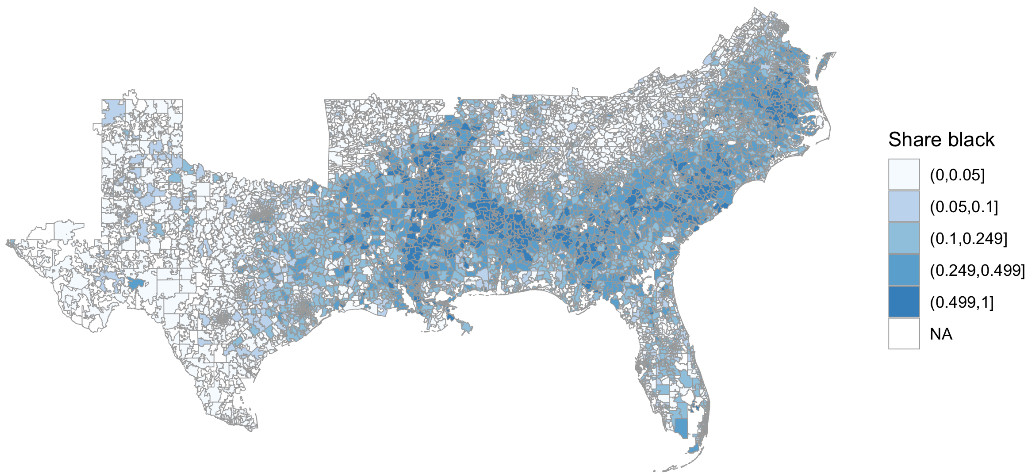

(b) Share Hispanic

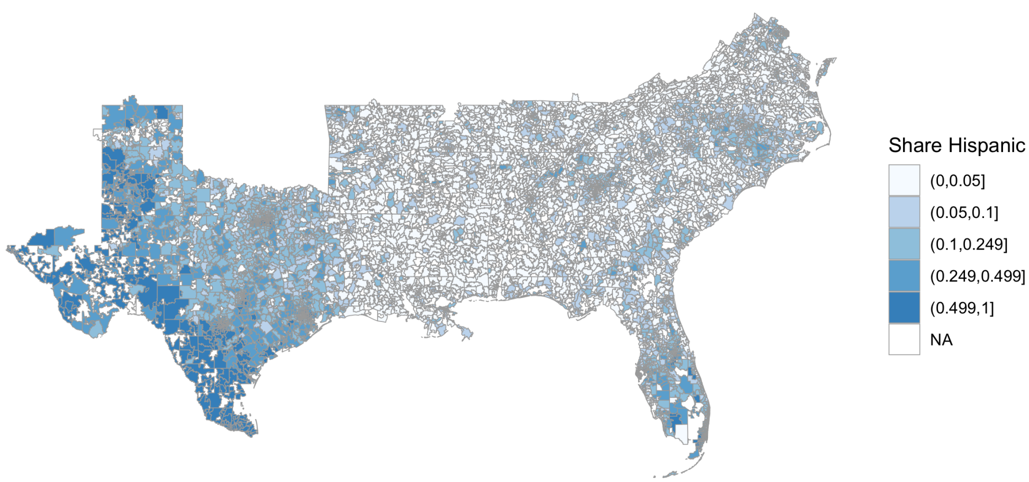

Source: 2010 U.S. Census

Figure A-5: Economic and Demographic Variables by Race

(a) Income and Asset Gaps

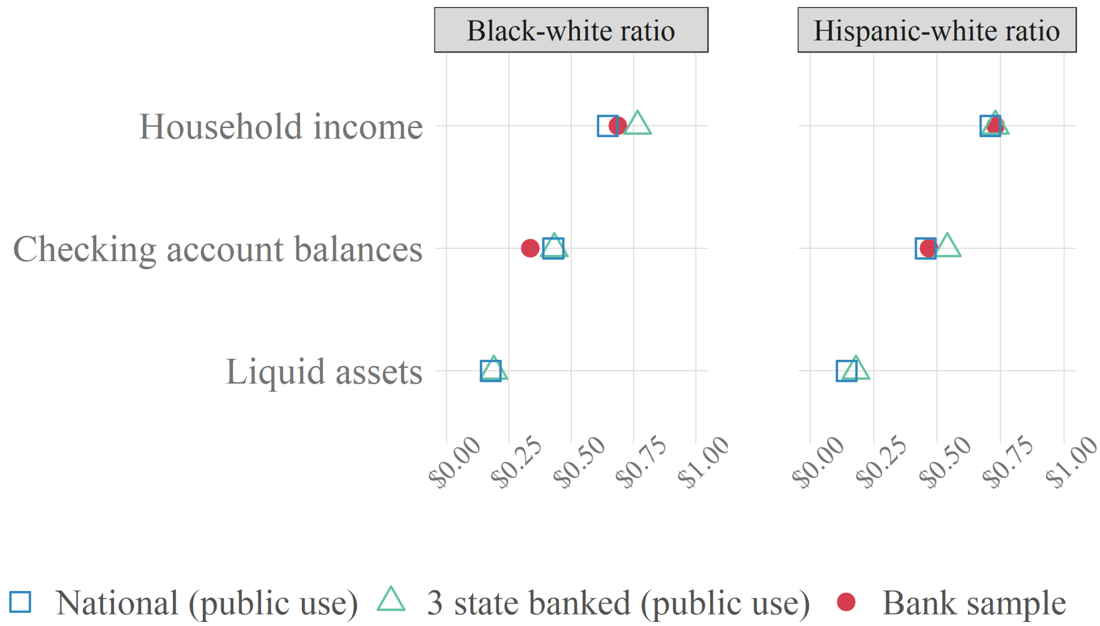

(b) Demographics

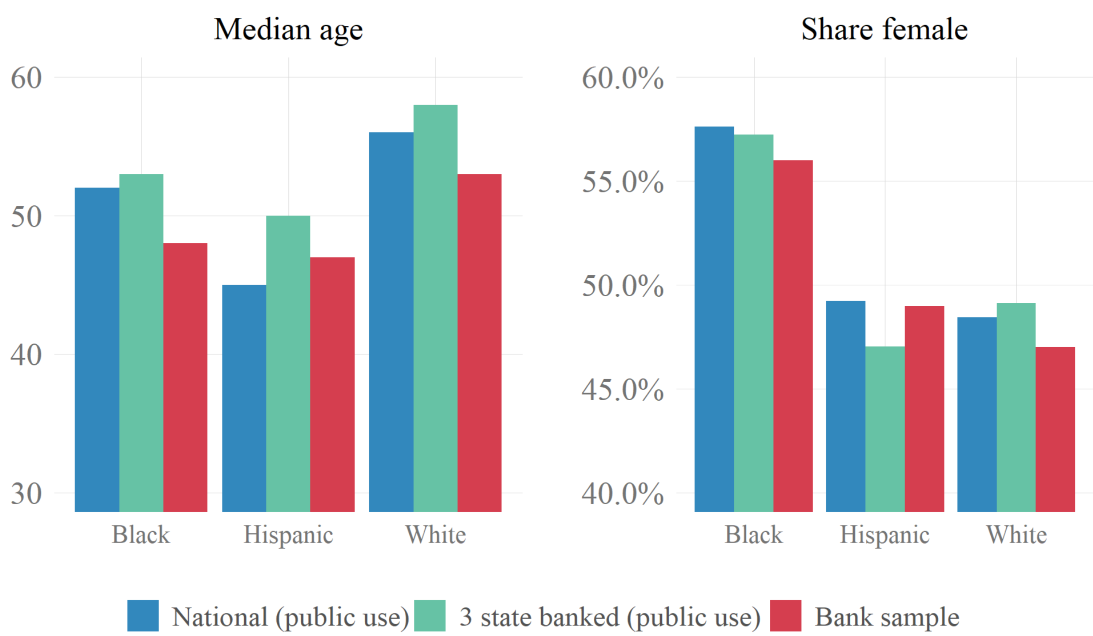

Note: Household income, checking account balances, and liquid assets are measured using medians.
Data sources: Public use household income, age and gender are from the 2015-2017 Current Population Survey. Public use checking account balances and liquid assets are from the Survey of Consumer Finances for 2016 and the Health and Retirement Study. See footnote 16 for details. Bank sample is the matched Chase-voter data described in the text.

Figure A-6: Distribution by Race – National versus Sample Frame versus Bank Sample

(a) Income Distribution

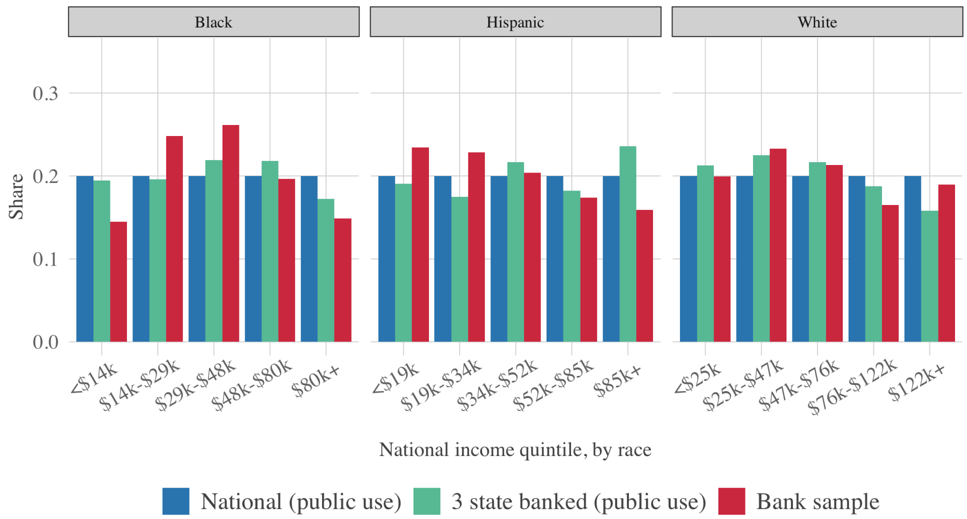

(b) Age Distribution

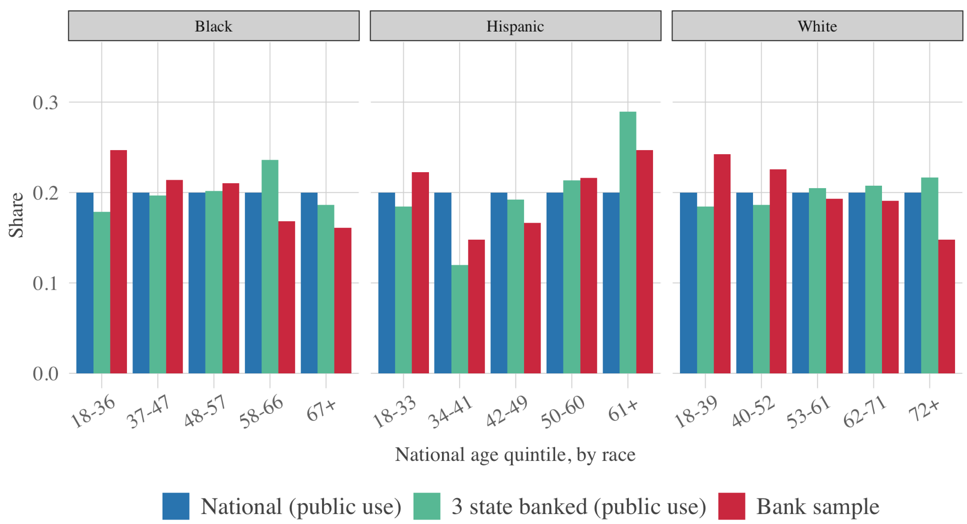

Note: This figure shows the distribution of income and age in the U.S. as a whole (blue), in a representative sample of banked households in the 3 states (green), and in the bank sample (red). The cutoff points for the bins are chosen to capture the quintiles of the national, race-specific distributions. Thus, by construction, the blue bars are all equal to 20 percent. These bars are included to ease comparison to the green and red bars. Source: Current Population Survey (CPS) for 2012, 2014, and 2016. Banked status is measured using the CPS Unbanked/Underbanked supplement. Household income is taken from the 2013, 2015, and 2017 CPS Annual Social and Economic Supplement where it is measured for the prior calendar year (2012, 2014, 2016). Income is inflated to 2016 dollars using CPI-U.

Figure A-7: Hispanic Origin in Three State Sample Frame Compared to U.S.

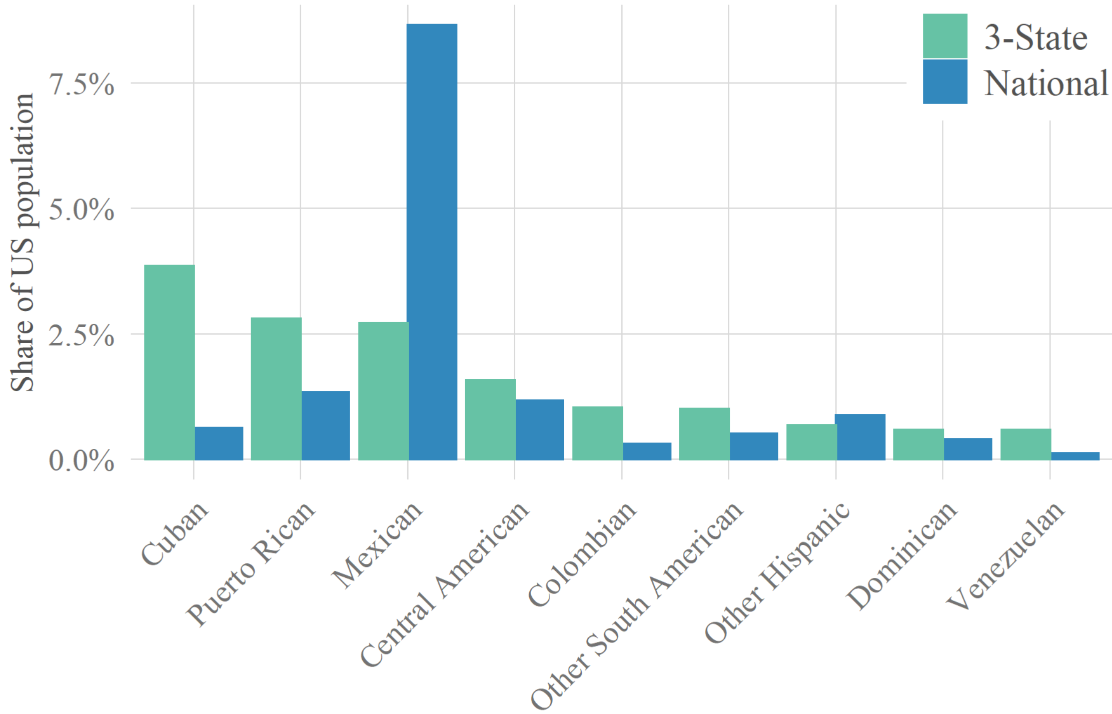

Source: 2018 American Community Survey 1-year estimates, Table B03001.

Figure A-8: Financial Vulnerability – National and in Sample Frame

(a) Share Ever Behind on Bills in Last Year

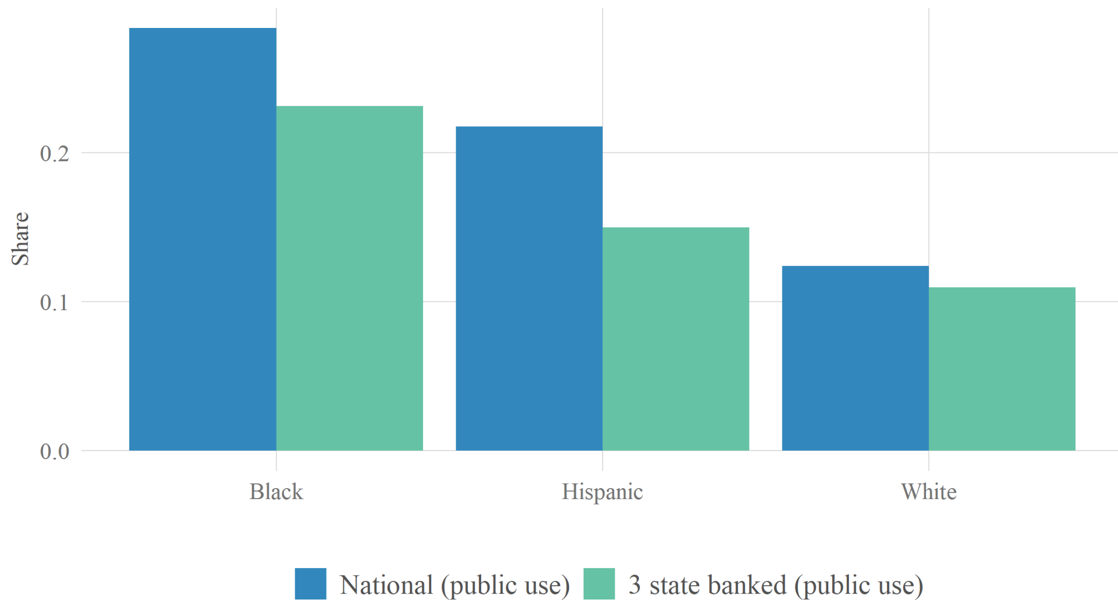

(b) Share with No Emergency Savings

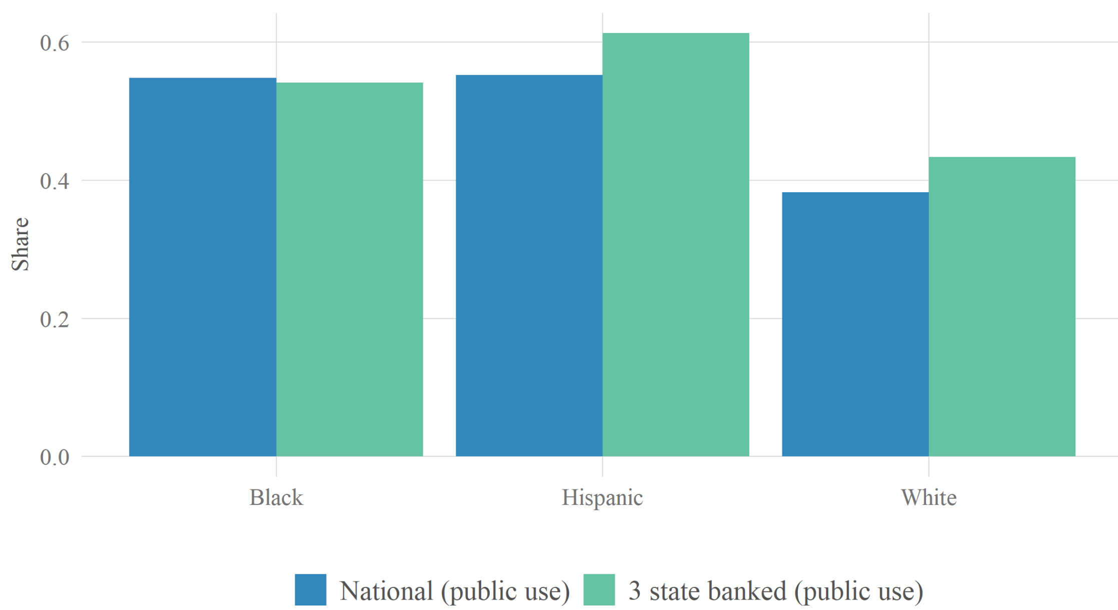

Note: The questions are "Often times, households find that they are not able to keep up with their bills. Over the last 12 months, was there a time when (you/ or someone else in your household) fell behind on bill payments?" and "Did (you/ or anyone in your household) set aside any money in the past 12 months that could be used for unexpected expenses or emergencies? I'm only asking about funds that could be easily spent if necessary, and am not asking about retirement or other long-term savings."
Source: Current Population Survey (CPS), June 2015 and June 2017 Unbanked/Underbanked CPS Supplements.

Figure A-9: Leads and Lags of Firm Pay Shocks

(a) Pay Per Paycheck with No Lags

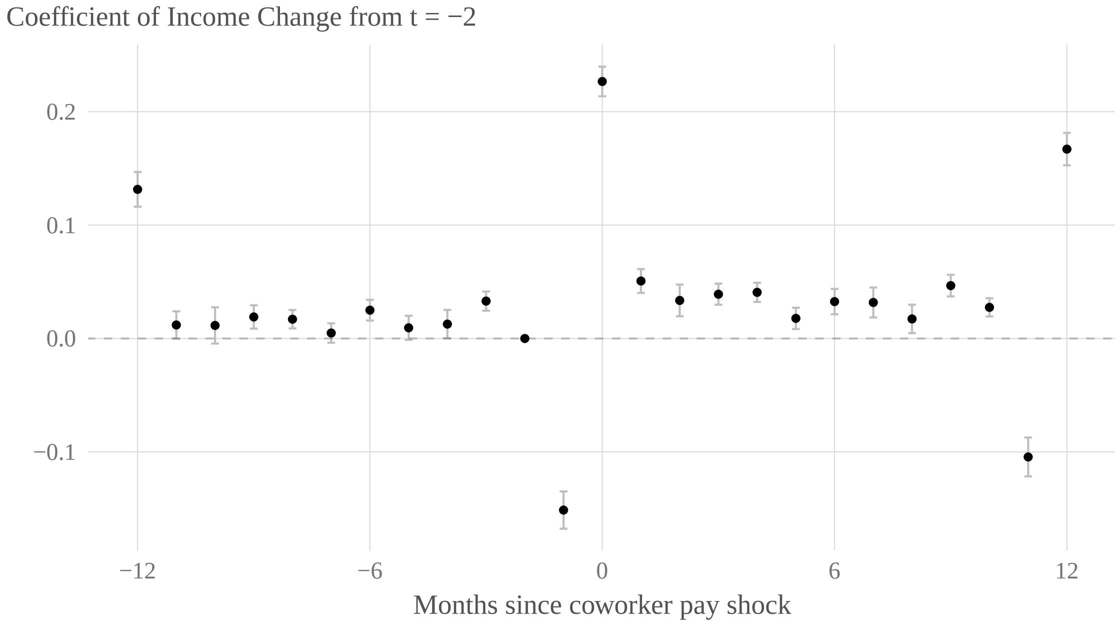

(b) Pay Per Paycheck with Five Lags

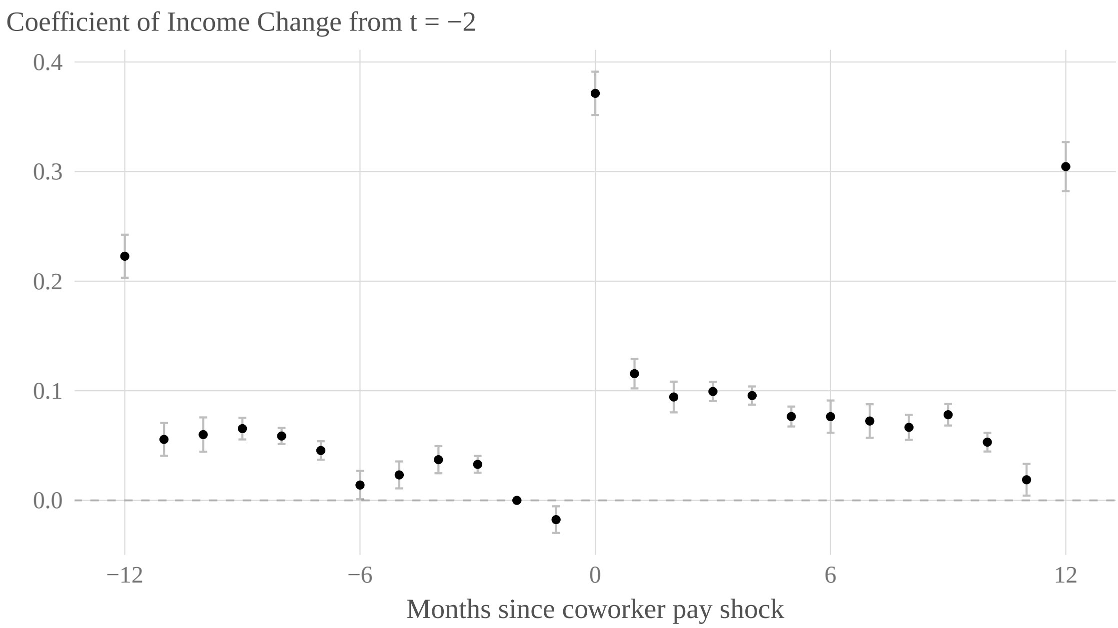

Note: This figure shows the relationship between the instrument at date $t = 0$ and changes in individual income between date t and date $t = -2$ for equation (4) in the top panel and equation (7) in the bottom panel. The specification shown in the bottom panel is our preferred specification.

Figure A-10: First Stage by Race and Income Quartile

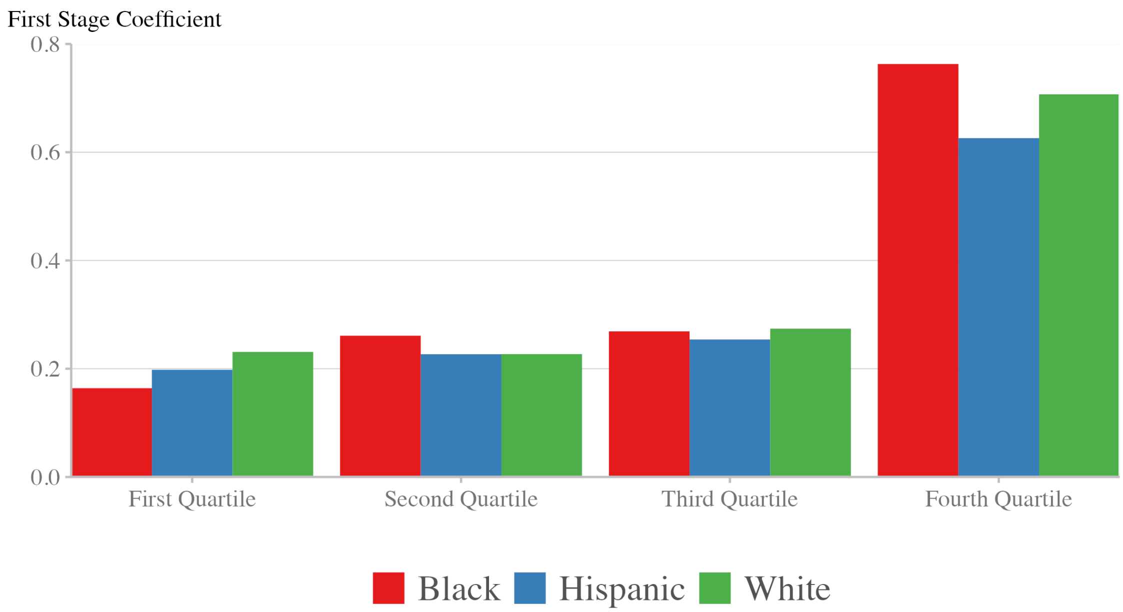

Note: We re-estimate the first stage equation (6) separately by race and by income quartile in the month prior to the shock; the coefficients are shown in this figure. The first stage coefficient is much larger in the top income quartile, perhaps consistent with a common timing for annual bonus payments together with top quartile workers being more likely to receive bonuses (Lemieux, MacLeod and Parent 2009). The first stage coefficient is similar across races within income quartile.

Figure A-11: Leads and Lags of Firm Pay Shocks by Race

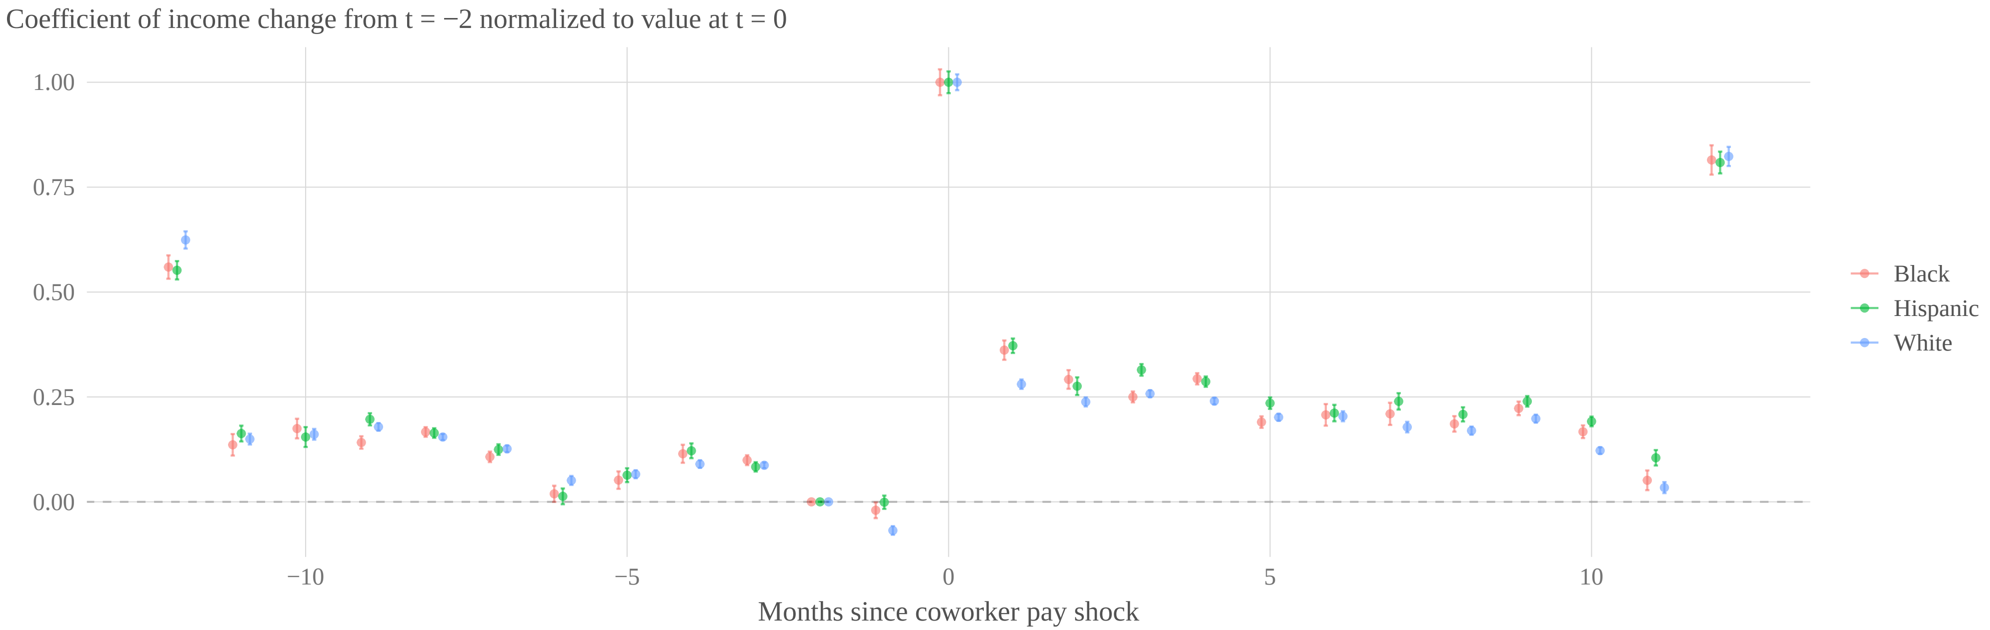

Note: This figure shows the relationship between the instrument at date $t = 0$ and changes in individual income between date t and date $t = -2$ for various racial groups. The values for each group are normalized to 1 at date $t = 0$.

Table A-1: Relation of Checking Balances and Liquid Assets

	Dependent variable:				
	Liquid (broad)			Liquid (narrow)	Financial
	(1)	(2)	(3)	(4)	(5)
Black	0.017	−0.685***	−0.743***	−0.359***	−0.805***
	(0.122)	(0.023)	(0.022)	(0.014)	(0.023)
Hispanic	−0.162	−0.810***	−0.862***	−0.470***	−1.283***
	(0.134)	(0.025)	(0.024)	(0.015)	(0.024)
White x Log(checking)	1.103***				
	(0.006)				
Black x Log(checking)	1.005***				
	(0.016)				
Hispanic x Log(checking)	1.013***				
	(0.017)				
Log(checking)		1.083***			
		(0.005)			
Constant	0.755***	0.912***	1.548***	1.194***	2.905***
	(0.048)	(0.043)	(0.009)	(0.006)	(0.009)
Observations	46,637	46,637	46,637	56,028	56,028
R^2	0.502	0.501	0.043	0.025	0.060

Notes: This table shows the relationship of log checking account balance and log total liquid and financial assets by race. See Section 3.6 for details. Sample is 2010, 2013, and 2016 Survey of Consumer Finances.

Table A-2: Race & Hispanicity Question in Various Datasets

(a) Florida Voter File

Race/Ethnicity
American Indian/Alaskan Native
Asian/Pacific Islander
Black, not of Hispanic Origin
Hispanic
White, not of Hispanic Origin
Multi-racial
Other

(b) Georgia Voter File

Race/Ethnicity
Black
White
Hispanic/Latino
Asian/Pacific Islander
American Indian/Alaskan Native
Other

(c) Louisiana Voter File

Race
White
Black
Asian
Hispanic
American Indian
Other

(d) Survey of Consumer Finances

"Which of these categories do you feel best describe you: White, Black or African-American, Hispanic or Latino, Asian, American Indian or Alaska Native, Hawaiian Native or other Pacific Islander, or another race?"

Source: Voter registration forms in 2020 and survey text from 2016 SCF.

Table A-3: Distribution of Race and Ethnicity in State Voter Files

	Florida	Georgia	Louisiana	Sample
White	0.6348	0.5412	0.5979	0.6036
Black	0.1345	0.2986	0.2754	0.1987
Hispanic	0.1631	0.0276	0.0120	0.1058
Asian	0.0190	0.0206	0.0078	0.0181
American Indian	0.0032	0.0011	0.0041	0.0027
Multi-racial	0.0066	NA	NA	0.0039
Other	0.0161	0.0124	0.0217	0.0157
Missing	0.0226	0.0986	0.0811	0.0515

Notes: Georgia asks if voter is "Hispanic/Latino". Florida asks if voter is "White, not of Hispanic Origin", "Black, not of Hispanic Origin", and "American Indian/Alaskan Native". Georgia and Florida ask if voter is "Asian/Pacific Islander". Exact question wording is in FigureA-2.

Table A-4: Location of Matched Sample

MSA	Observations	Percentage	Cumulative percentage
Miami	546887	32.8%	32.8%
Atlanta	207341	12.4%	45.2%
Orlando	176195	10.6%	55.8%
Tampa	129856	7.8%	63.6%
New Orleans	120678	7.2%	70.8%
Baton Rouge	66617	4.0%	74.8%
Sarasota	36908	2.2%	77.1%
Palm Bay	36786	2.2%	79.3%
Port St. Lucie	30440	1.8%	81.1%
Shreveport	29426	1.8%	82.9%
Rest of Florida	145061	8.7%	91.6%
Rest of Louisiana	128344	7.7%	99.3%
Rest of Georgia	12309	0.7%	100.0%

Notes: This table shows the number of observations in the matched sample by metro area.
Source: Matched JPMCI data and voter registration records.

Table A-5: Distribution of Race and Ethnicity

Race	Voter	Voter (Reweighted)	Matched
Black	20.3%	21.5%	25.7%
Hispanic	10.9%	18.0%	22.9%
White	61.9%	52.1%	45.7%

Notes: This table describes the racial composition of registered voters in FL, GA, and LA, both in the voter files and in the matched bank data. The first column ("Voter") shows the racial composition in the voter files. The second column ("Voter (Reweighted)") adjusts for the fact that bank customers are more likely to live in certain ZIP codes. Specifically, we compute the racial composition of the voter file, weighting each ZIP code by the percentage of registered voters successfully matched to the bank data. The third column ("Matched") shows the racial composition of the matched bank-voter data. The rows do not sum to 100 percent because some people report a race other than White, Black or Hispanic, or report no race at all.

Table A-6: Summary Statistics for Analysis Sample

Variable	Race	Mean	Median	Std. Dev	N Households	N Household-Months
Labor Income	White	3,792.5	2,864.1	7,881.5	431,800	13,722,548
	Black	2,550.4	2,059.2	3,183.6	230,862	6,870,680
	Hispanic	2,995.5	2,407.9	3,700.9	221,204	6,673,632
Coworker Pay	White	3,331.0	2,906.4	3,504.9		
	Black	2,994.4	2,675.3	3,199.3		
	Hispanic	3,150.1	2,844.5	3,195.9		
Non Durable Consumption	White	2,524.3	1,827.8	9,121.4		
	Black	2,286.3	1,750.8	3,213.6		
	Hispanic	2,176.2	1,634.5	3,923.9		
Checking Assets	White	8,924.4	2,721.5	26,568.2		
	Black	3,587.3	1,142.6	11,120.5		
	Hispanic	5,069.5	1,579.7	15,118.4		
Buffer Ratio	White	15.2	1.3	7,811.3		
	Black	3.3	0.6	75.2		
	Hispanic	5.8	0.9	150.4		

Note: This table shows summary statistics for the analysis sample in our coworker pay design. See Section 3 for details. Buffer ratio is the ratio of checking assets to nondurable consumption.

Table A-7: First Stage for Various Specifications

	Dependent Variable: Δ Log Income		
	(1)	(2)	(3)
Δ Log Coworker Pay	0.721	0.383	0.197
	(0.008)	(0.013)	(0.017)
IV	Total Pay	Pay per Paycheck	Pay per Paycheck
Firm-by-Calendar-Month FEs	No	No	Yes
Observations	20,223,107	20,223,107	14,680,597
Adjusted R^2	0.207	0.021	0.155

Note: This table shows the first stage estimate of $\hat{\rho}$ from equation (7) for various specifications. Column (1) defines the instrument as the change in log mean coworker total pay. Columns (2) and (3) defines the instrument as the change in log mean coworker pay per paycheck. Column (3) adds firm-by-calendar-month fixed effects.

Table A-8: First Stage for Multiple Datasets

		First stage	
Periodicity	Data Source	Coef	(SE)
Paycheck	Bank account	0.38	(0.01)
	Time clock	0.40	(0.01)
Monthly	Bank account	0.72	(0.01)
	Time clock	0.80	(0.01)
Quarterly	Bank account	0.64	(0.02)
	Time clock	0.81	(0.01)
	Tax, WA, Lachowska et al. (2022)	0.56 to 0.65	(0.01 to 0.02)
	Tax, 7 states, CWBH	0.42 to 0.61	(0.01 to 0.03)

Notes: This table shows the results of the first stage regression on multiple datasets and with multiple levels of time aggregation. The rows labeled bank account correspond to the Chase data and regress change in log worker pay on change in log coworker pay. The rows labeled time clock correspond to the Homebase data and regress change in log worker hours on change in log coworker hours. The row labeled CWBH corresponds to the Continuous Wage and Benefit History data and regresses change in log worker pay on change in log coworker pay. Standard errors are clustered by firm. The paycheck-level and monthly-level regressions that we calculates ourselves in the bank account and time clock data control for five lags of the change in log coworker pay. The quarterly regressions do not control for lags, to align with the specification in Lachowska et al. (2022). That specification is reported in Table C.1 of Lachowska et al. (2022) and uses earnings per hour data.

Table A-9: Reasons for Income Volatility

Response	**Rate**
Irregular work schedule	41.6%
Bonuses	12.4%
Commissions	6.9%
Seasonal employment	12%
Periods of unemployment	14.7%
Investment income	9.2%
Other	23.3%
Refused	4.3%

Source: Survey of Household Economics and Decisionmaking, Table C.86

Table A-10: Impact of Income on Consumption: Comparison to BPP Specification

	Dependent Variable: Δ Log Non Durable Consumption			
	(1)	(2)	(3)	(4)
Δ Log Income	0.104	0.262	0.116	0.177
	(0.006)	(0.031)	(0.004)	(0.007)
IV	Own Pay per Paycheck (t+1)	Coworker Pay per Paycheck (t+1)	Own Pay (t+1)	Coworker Pay (t+1)
Time Aggregation	Month	Month	Month	Month
Observations	23,404,316	23,426,550	23,404,316	23,426,550
Adjusted R^2	0.004	−0.016	0.003	−0.002

Notes: This table reports the elasticity of consumption with respect to income combining the panel data methods in Blundell, Pistaferri and Preston (2008, BPP) with the firm pay shock instrument in this paper. Following BPP, columns (1) and (3) isolate the effect of temporary shocks by instrumenting for contemporaneous income changes using the income change for the same individual one period forward. Columns (2) and (4) use one period forward co-worker pay changes as the instrument instead.

Table A-11: Impact of Income on Non-Work-Related Expenditure

Category	Pooled	White	Black	Hispanic
All nondurables	0.216	0.163	0.338	0.295
	(0.017)	(0.011)	(0.042)	(0.032)
Non-work-related	0.199	0.149	0.295	0.279
	(0.021)	(0.014)	(0.052)	(0.040)

Notes: This table reports the elasticity of consumption with respect to income separately for all nondurable spending (as in our baseline specification) and for the subset of non-work-related nondurable spending. We define a spending category as work-related if it exhibits a larger-than-median drop at retirement (Aguiar and Hurst 2013). We define retirement as the first receipt of a Social Security check as in Ganong and Noel (2019).

Table A-12: Impact of Income on Consumption at a Quarterly Frequency

	Dependent Variable: Δ Log Non Durable Consumption	
	(1)	(2)
Δ Log Income	0.294	0.224
	(0.017)	(0.012)
(Δ Log Income) x Black		0.172
		(0.025)
(Δ Log Income) x Hispanic		0.119
		(0.017)
OLS/IV	IV	IV
Black and Hispanic Dummies		Yes
Observations	19,512,278	19,512,278
Adjusted R^2	0.087	0.072

Note: This table replicates columns (3) and (4) of Table 1, but compares the three-month consumption change to the three-month income change (rather than the one-month changes in Table 1).

Table A-13: Impact of Income on Consumption by Assets and Race with Total Pay as Instrument

	Dependent Variable: Δ Log Non Durable Consumption							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δ Log Income	0.075	0.057	0.179	0.155	0.182	0.168	0.183	0.182
	(0.004)	(0.003)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)
(Δ Log Income) x Black		0.048		0.068		0.039		0.013
		(0.003)		(0.004)		(0.004)		(0.004)
(Δ Log Income) x Hispanic		0.023		0.036		0.023		−0.007
		(0.003)		(0.004)		(0.004)		(0.004)
(Δ Log Income) x Checking					−0.185	−0.171		
					(0.006)	(0.006)		
(Δ Log Income) x Liquid(Imputed)							−0.191	−0.188
							(0.006)	(0.006)
OLS/IV	OLS	OLS	IV	IV	IV	IV	IV	IV
Black and Hispanic Dummies		Yes		Yes		Yes		Yes
Asset Rank Control					Yes	Yes	Yes	Yes
Observations	25,774,028	25,774,028	20,095,473	20,095,473	20,095,473	20,095,473	20,095,473	20,095,473
Adjusted R^2	0.004	0.004	0.002	0.002	0.005	0.005	0.005	0.005

Note: This table replicates Table 1 using average coworker total monthly pay as the instrument rather than average coworker pay-per-paycheck.

Table A-14: Impact of Income on Consumption by Assets and Race with Firm-by-Calendar-Month FEs

	Dependent Variable: Δ Log Non Durable Consumption							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δ Log Income	0.076	0.056	0.171	0.134	0.187	0.161	0.192	0.184
	(0.001)	(0.001)	(0.016)	(0.014)	(0.018)	(0.016)	(0.019)	(0.017)
(Δ Log Income) x Black		0.053		0.115		0.066		0.022
		(0.001)		(0.017)		(0.016)		(0.015)
(Δ Log Income) x Hispanic		0.028		0.073		0.053		0.006
		(0.001)		(0.014)		(0.014)		(0.014)
(Δ Log Income) x Checking					−0.302	−0.281		
					(0.025)	(0.023)		
(Δ Log Income) x Liquid(Imputed)							−0.325	−0.312
							(0.028)	(0.026)
OLS/IV	OLS	OLS	IV	IV	IV	IV	IV	IV
Black and Hispanic Dummies		Yes		Yes		Yes		Yes
Asset Rank Control					Yes	Yes	Yes	Yes
Observations	18,404,294	18,404,294	14,590,652	14,590,652	14,590,652	14,590,652	14,590,652	14,590,652
Adjusted R^2	0.025	0.025	0.022	0.022	0.024	0.024	0.023	0.024

Note: This table replicates Table 1 adding firm-by-calendar-month fixed effects.

Table A-15: First Stage by Race

	Dependent Variable: Δ Log Income		
	(1)	(2)	(3)
Δ Log Coworker Pay	0.711	0.435	0.247
	(0.007)	(0.012)	(0.015)
(Δ Log Coworker Pay) x Black	0.033	-0.137	-0.136
	(0.007)	(0.013)	(0.011)
(Δ Log Coworker Pay) x Hispanic	0.011	-0.111	-0.100
	(0.007)	(0.014)	(0.016)
IV	Total Pay	Pay per Paycheck	Pay per Paycheck
Firm-by-Calendar-Month FEs	No	No	Yes
Observations	20,223,107	20,223,107	14,680,597
Adjusted R^2	0.207	0.022	0.156

Note: This table replicates Table A-7 adding interactions by race.

Table A-16: Second Stage by Asset Quartile

	Dependent Variable: Δ Log Non Durable Consumption			
	Asset Quartile 1	Asset Quartile 2	Asset Quartile 3	Asset Quartile 4
Δ Log Income	0.442	0.289	0.188	0.104
	(0.035)	(0.020)	(0.016)	(0.011)
Consumption	2,567	2,576	2,434	1,937
Income	2,390	2,983	3,463	4,293
MPC	0.475	0.250	0.132	0.047
	(0.038)	(0.017)	(0.011)	(0.005)
Observations	4,821,070	5,031,689	5,104,845	5,137,869
Adjusted R^2	-0.021	-0.012	-0.004	-0.001

Note: This table reports second stage estimates of the elasticity of consumption with respect to income from column (3) of Table 1 separately by quartile of liquid asset buffer. We construct asset buffer as the ratio of lagged assets to lagged nondurable consumption. Thus, households that are classified as high assets are a mix of households with permanently high assets and households with temporarily low consumption in prior months. This table also reports the scaling factors used to translate each group's elasticity into an MPC.

Table A-17: Unemployment Experiences by Race

	White	Black	Hispanic
3 state (public use)			
Individual UI recipiency rate	0.2%	0.6%	0.3%
Unemployment rate	2.9%	5.7%	3.4%
Mean unemployment duration (weeks)	24	25	21
Bank sample			
Household UI recipiency rate	1.3%	1.2%	1.5%
Mean completed UI duration (weeks)	15	15	15

Source: UI recipiency rate is from the Department of Labor. Unemployment statistics are for banked households in the 2017 Current Population Survey.

Table A-18: Variance of Transitory Labor Income Shocks in Chase and SIPP Data

	Monthly Chase	Monthly SIPP	Quarterly Chase	Quarterly SIPP
Overall	0.09	0.13	0.20	0.17
Black	0.10	0.13	0.21	0.16
Hispanic	0.10	0.08	0.22	0.12
White	0.09	0.14	0.18	0.18

Note: This table shows the variance of transitory labor income shocks in two different data sets and at two levels of time aggregation. We calculate this variance following the method in Appendix 2 of Carroll and Samwick (1997). Because we do not require estimates of this variance for each household, we simplify this method by calculating population-level regressions rather than household-by-household regressions. Following Carr, Moffitt and Wiemers (2020) we trim the earnings distribution at the 1st and 99th percentiles before calculating monthly or quarterly income changes.
Source: 2004 SIPP and JPMCI.

Table A-19: Structural Model Parameters

Parameter					Source
A. Common Parameters	All Households				
Coefficient of relative risk aversion: γ	1.0				N/A
Monthly discount factor: δ	0.937				Kaplan and Violante (2022)
Implied annual ratio of transitory to permanent income shock variance: $\sigma^2_{\theta,annual} \big/ \sigma^2_{\theta,annual}$	2.03				Carroll and Samwick (1997)
Monthly coefficient of variation: $\sigma^2_y \big/ \bar{y}$	0.38				Farrell, Greig and Yu (2019)
Monthly variance of permanent shocks (when working): σ^2_θ	0.0004				Calibrated to Carroll and Samwick (1997)
Monthly variance of transitory shocks (when working): σ^2_ζ	0.1296				& Farrell, Greig and Yu (2019)
Unemployment replacement rate: α_{UI}	0.575				Mitchell and Phillips (2006)
Social Security replacement rate:	0.45				Mitchell and Phillips (2006)
Retirement age: T_R	62				Biggs and Springstead (2008)
Maximum age: T	85				N/A
Group-Specific Parameters	All Households	Black	Hispanic Households	White Households	
Unemployment probability: π_U	0.04	0.06	0.04	0.03	CPS (2015-2018)
Lifecycle income growth: Γ_t	-		-	-	CPS (2014-2016)
Probablity of death: S_t	-		-	-	CDC (1999-2016)

Table A-20: Structural Model Results

	(1)		(2)		(2)		(3)	
	All Households		Black Households		Hispanic Households		White Households	
A. Target moments								
	Data	Model	Data	Model	Data	Model	Data	Model
Average elasticity: β	0.21	0.21	0.35	0.33	0.29	0.29	0.16	0.16
Total Asset to Income Ratio (mean)	4.3	4.2	2.1	2.0	2.5	2.6	4.8	4.9
Total Asset to Income Ratio (median)	1.7	3.9	0.5	1.8	0.8	2.2	2.3	5.0
Liquid Asset to Income Ratio (mean)	0.6	1.6	0.2	1.9	0.2	1.8	0.7	1.3
Liquid Asset to Income Ratio (median)	0.1	1.5	0.02	1.8	0.02	1.7	0.1	1.1
B. Estimated parameters								
Low liquid interest rate: r^1	-3.0%		-6%		-4%		0.0%	
Medium liquid interest rate: r^2	0.0%		-6%		-4%		0.0%	
High liquid interest rate: r^3	0.0%		-3%		-2%		1.0%	
Low illiquid interest rate: r_n^1	4.0%		2%		2%		8%	
Medium illiquid interest rate: r_n^2	9.0%		2%		4%		8%	
High illiquid interest rate: r_n^3	10.0%		7%		8%		11%	
C. Welfare gain from eliminating temporary income volatility								
Average gain (consumption equivalent):	0.8%		1.6%		1.3%		0.5%	

Note: This table shows estimates from the model in Section 4. Panel A shows the empirical moments from the data, which we use to estimate the parameters of the model. The definition of liquid assets follows Kaplan and Violante (2014), who subtract credit card debt from assets held in checking and savings accounts. Panel B shows the calibrated model parameters. Interest rates are annualized. Panel C shows the results from eliminating temporary income volatility and calculating the welfare gain in lifetime consumption equivalents. All numbers are rounded to two significant digits. See Section C.2 for details.

B Data Appendix

B.1 Matching algorithm

We draw on standard fuzzy matching algorithms to match bank customers to voter records. The match uses name, address, and age. Exact date of birth is available for matching in the bank data. In the voter records, exact date of birth is available in Florida and Georgia, but in Louisiana age is only available in years. Our algorithm first blocks on street name, street number, and ZIP code. Then, within each block, the algorithm uses a fuzzy match. We measure name similarity and address similarity using a Jaro-Winkler distance metric. The matching algorithm is implemented by a third-party, so we as the researchers never see personally identifiable information.

The exact algorithm used for matching is the following:

1. Starting with customer data, construct a list of candidate matches by left-joining voter data using house number, street name, and ZIP.
   (a) Choose the "best" voter record. How? Choose the record with the highest address score. If address scores are a tie, then choose the record with the highest name score.
   (b) Accept a candidate match if either of two conditions are met
      i. Address similarity > 0.97 and full name similarity > 0.97
      ii. Address similarity$= 1$, full name similarity < 0.97, first name similarity > 0.97, and birth year differs by no more than one year[2]
2. With customers who were not matched in step 1, join to voter data using house number, street name, and city. Repeat steps (a) and (b).
3. With customers who were not matched in step 1 or step 2, join to voter data using ZIP and birth year. Repeat steps (a) and (b).

This procedure yields about 6.5 million customers across all product types with accepted matches to the voter data. The analysis sample is drawn from the 1.8 million accepted matches that have checking accounts with Chase.

B.2 Quality of matches

We validate the quality of accepted matches using data from mortgage applications. The Home Mortgage Disclosure Act requires lenders to ask two race & ethnicity questions: one is about race (American Indian or Alaska Native, Asian, Black or African American, Native Hawaiian or Other Pacific Islander, or White) and one about Hispanicity (Hispanic or Latino, or Not Hispanic or Latino). We classify a voter record and a mortgage application record as in agreement if the household reports: (i) Black in the voter record and race as Black in the mortgage application, (ii) White in the voter record and race as White in the mortgage application, or (iii) Hispanic in the voter record and Hispanicity as Hispanic or Latino in the mortgage application. Among households that identify as Black, White, or Hispanic in the voter record, the agreement rate is 98.9 percent. The agreement rate is 93.7 percent for

[2] If a customer's birth year is not available in the bank records, then this condition is address similarity $=$ 1 and first name similarity > 0.97.

voters that identify as Black, 99.5 percent for voters that identify as White, and 90.4 percent for voters that identify as Hispanic. The 98.9 percent overall agreement rate is better than the agreement rate we would achieve if we assigned all customers as White (in which case the agreement rate would be 89.5 percent), Hispanic (which would yield an agreement rate of 10.7 percent), or Black (which would yield an agreement rate of 10.0 percent).

B.3 IV regression details

The two stage least squares equations we estimate are

$$\Delta c_{it} = \alpha + \beta \Delta y_{it} + X_{it} + \varepsilon_{it}$$
$$\Delta y_{it} = \phi + \rho \Delta y_{j(-i,t),t} + X_{it} + \nu_{it}.$$

Two small technical challenges are how to handle households with more than one job and workers who have newly joined a firm.

To address the first challenge, if a household i is employed by more than one firm j because there are two workers in the household or because one worker has two jobs then the household enters the regression twice, with standard errors clustered by firm j.

The second challenge arises because the first stage equation (equation 7) in the baseline specification includes the change in coworker pay ($\Delta y_{j(-i,t),t}$) and X_{it} as five lags of the change in coworker pay. Let all the workers at firm j at time t be captured by j_t. For an individual i who works at firm j at time t, we construct the leave out average firm income as

$$y_{j(-i,t),t} = \frac{\sum_{i' \in j_t} y_{i't} - y_{it}}{\sum_i \mathbf{1}(i \in j_t) - 1} \tag{16}$$

and then the regressor is the change in log coworker income: $\Delta y_{j(-i,t),t} = \log y_{j(-i,t),t} - \log y_{j(-i,t),t-1}$. The control for the k^{th} lag of the change in coworker pay is $\Delta y_{j(-i,t),t-k} = \log y_{j(-i,t),t-k} - \log y_{j(-i,t),t-k-1}$. Finally, in the case where the worker was not employed in period $t-k$ at that firm, equation (16) simplifies to $y_{j(-i,t),t-k} = \frac{\sum_{i' \in j_t} y_{i't-k}}{\sum_{i'} \mathbf{1}(i' \in j_{t-k})}$.

In practice, the way that we calculate these statistics is to first construct an employer panel with total income ($\sum_{i' \in j_t} y_{i't}$) and the total number of workers ($\sum_{i'} \mathbf{1}(i' \in j_t)$). We then left join individual income onto this employer panel and construct leave-out means. Had we instead started with individual income and left joined on the employer panel, this would inadvertently impose a requirement that a worker be employed at the same firm for K periods before they enter the analysis sample. We stress that our analysis, which starts from the employer panel, does not impose this requirement.

B.4 Orthogonality conditions

A more demanding version of assumption FO is

$$E(\varepsilon_{it} | \{\psi_{j(i,t),s}\}_{s=1}^{T}) = 0. \tag{17}$$

When this condition holds, the $\hat{\beta}$ from equation (6) identifies the consumption response to an unanticipated income shock lasting exactly one period. The condition in equation (17) can best be understood as adding two additional exclusion restrictions over Assumption FO.

First, it requires a pre-shock exclusion restriction for $s < t$: consumption does not respond ahead of time to future firm-wide pay changes. This assumption would fail if agents know that they will receive additional income in the future because of an increase in firm-wide pay ($\psi_{j(i,t),s} > 0$) and increase their consumption in advance ($\Delta c_{i,t-1} > 0$ leads to lower ε_{it}). This no-anticipation assumption is common in the literature and is made in, e.g. Blundell, Pistaferri and Preston (2008). It is useful to note that this assumption does not require that households are fully ignorant of future firm-wide pay changes. It only requires that they not adjust their consumption before the pay change occurs.

Second, it requires a post-shock exclusion restriction for $s > t$: once a shock hits, consumption responds only to the income change received in that period rather than any potential future changes in income. This assumption would fail if $\psi_{j(i,t),s}$ contains news about the path of income after time t which then affects consumption at time t. This assumption is satisfied if agents are unaware of the news content (which the econometrician is aware of based on Figure A-9b) or do not increase their consumption in response to the news.

C Notes on welfare calculations

C.1 Derivation of formula for statistical welfare gain λ

We begin with equation (14) from Section 4:

$$\mathbb{E}\left[\sum \delta^t \frac{((1+\lambda)\, C_t)^{1-\gamma}}{1-\gamma}\right] \equiv \mathbb{E}\left[\sum \delta^t \frac{\tilde{C}_t^{1-\gamma}}{1-\gamma}\right].$$

Passing the expectations operator through the summation operator, and using the definitions of C_{it} and $\tilde{C}_{it}$ above, we have:

$$\sum \delta^t \frac{(1+\lambda)^{1-\gamma}\, \mathbb{E}\left[\tilde{C}_t^{\,1-\gamma}\right] \mathbb{E}\left[e^{\beta\theta_t(1-\gamma)}\right]}{1-\gamma} = \sum \delta^t \frac{\mathbb{E}\left[\tilde{C}_t^{\,1-\gamma}\right]}{1-\gamma}$$

Subtracting the left-hand side from the right, we have:

$$\begin{aligned} 0 &= \sum \left\{ \delta^t \frac{\mathbb{E}\left[\tilde{C}_t^{\,1-\gamma}\right]}{1-\gamma} \left[1-(1+\lambda)^{1-\gamma}\, \mathbb{E}\left[e^{\beta\theta_t(1-\gamma)}\right]\right] \right\} \\ &= \left[1-(1+\lambda)^{1-\gamma}\, \mathbb{E}\left[e^{\beta\theta_t(1-\gamma)}\right]\right] \sum \left\{ \delta^t \frac{\mathbb{E}\left[\tilde{C}_t^{\,1-\gamma}\right]}{1-\gamma} \right\}, \end{aligned}$$

which leaves us with:

$$\begin{aligned} 1-(1+\lambda)^{1-\gamma}\, \mathbb{E}\left[e^{\beta\theta_t(1-\gamma)}\right] &= 0 \\ \Rightarrow (1+\lambda)^{\gamma-1} &= \mathbb{E}\left[e^{\beta\theta_t(1-\gamma)}\right] \end{aligned} \tag{18}$$

For the term on the right, recall that we have assumed that $\mathbb{E}\left[e^{\theta}\right] = 1$, which means that $\mathbb{E}[\theta] = -\sigma_\theta^2/2$. Also, recall that $var[\theta] = \sigma_\theta^2$. It follows that $\mathbb{E}[(1-\gamma)\,\beta\theta] = -(1-\gamma)\,\beta\sigma_\theta^2/2$ and $var[(1-\gamma)\,\beta\theta] = (1-\gamma)^2\,\beta^2\sigma_\theta^2$. Thus, we have:

$$\begin{aligned} \mathbb{E}\left[e^{\beta\theta_t(1-\gamma)}\right] &= e^{-(1-\gamma)\beta\sigma_\theta^2/2+(1-\gamma)^2\beta^2\frac{\sigma_\theta^2}{2}} \\ &= e^{(\gamma-1)\left[\beta+\beta^2(\gamma-1)\right]\frac{\sigma_\theta^2}{2}} \end{aligned}$$

Taking logs of both sides of (18), we have

$$\begin{aligned} (\gamma-1)\log(1+\lambda) &= (\gamma-1)\left[\beta+\beta^2(\gamma-1)\right]\frac{\sigma_\theta^2}{2} \\ \Rightarrow \log(1+\lambda) &= \left[\beta+\beta^2(\gamma-1)\right]\frac{\sigma_\theta^2}{2} \end{aligned}$$

Finally, if we use the approximation $\log(1+\lambda) \approxeq \lambda$, we have:

$$\lambda \approxeq \left[\beta+\beta^2(\gamma-1)\right]\frac{\sigma_\theta^2}{2} \tag{19}$$

Note, when $\beta = 1$, the expression collapses to the Lucas (1987) result:

$$\lambda \approxeq \gamma\frac{\sigma_\theta^2}{2}$$

As we can see from (19) λ is increasing in σ_θ^2. We can also rewrite equation (19) as follows:

$$\lambda \approxeq \left[\gamma + \frac{(1-\beta)}{\beta}\right]\beta^2\frac{\sigma_\theta^2}{2},$$

which makes it clear that λ is also increasing in γ. Finally, we have:

$$\frac{\partial\lambda}{\partial\beta} = \left[1 + 2\beta\left(\gamma - 1\right)\right]\frac{\sigma_\theta^2}{2}$$

When $\gamma \geq 1$, λ is unambiguously increasing in β. Otherwise, when $\gamma < 1$, $\partial\lambda/\partial\beta > 0$ whenever $\beta < 1/\left(2\left(1-\gamma\right)\right)$.

C.2 Structural Model

The advantage of our statistical model in Section 4 is that it only requires estimates of differences in consumption elasticities and transitory income volatilities to capture heterogeneity in welfare costs. We now turn to estimating a structural, lifecycle model of consumption and saving, which, among other things, allows households to fully reoptimize under different volatility scenarios and incorporates wealth into the consumption decision. This has several advantages, including capturing a broader measure of welfare and tying more directly into our discussion of liquid assets and consumption smoothing. To demonstrate the welfare implications of volatility we build a model that can match the elasticities we observe in our data. Within this model, we can then shut off transitory income volatility and calculate the expected lifetime welfare improvement for each agent, demonstrating that the type of volatility we observe in our data is economically significant.

As discussed in Section 4.3, we consider a two-asset lifecycle model, in the spirit of Kaplan and Violante (2014), which is able to match the relatively low level of liquid asset accumulation observed in the data. In each period, a household chooses how much to consume, c_t, in order to maximize expected lifetime utility. The household may save liquid assets in each period, a_t, which earn a gross return R at the beginning of the next period. Alternatively, they can make a deposit or withdrawal, d_t, using an illiquid asset, which requires an adjustment cost of k dollars. Each period, a gross return of R_n is earned on the balance of the illiquid asset, n_t. As in Kaplan and Violante (2022), we assume that households cannot borrow. Labor income, Y_t, follows the same process as in (12) above, when one is employed. In addition, there is a probability π_{UI} of entering into unemployment for one period, with a household-level unemployment replacement rate of α_{UI}. Upon reaching the retirement age, T_R, households receive a stream of Social Security payments, which replace a fixed percentage of income prior to retirement.

The household lives for T periods and there is a time-varying probability of surviving at the beginning of each period, S_{it}. By the end of the last period, T, household assets must be non-negative. Households maximize a discounted stream of expected utility:

$$\max_{\{c_t\}} E\left[\sum_{j=0}^{T-t}\delta^j u\left(c_{t+j}\right)\right]$$

where δ is the per-period discount factor and utility takes a CRRA form: $u\left(c\right) = \frac{c^{1-\gamma}}{1-\gamma}$. The parameter γ captures risk aversion and intertemporal substitution preferences. The expectation is taken over permanent and transitory income shocks, and unemployment and survival probabilities. Assets evolve according to the following transition rules and constraints:

$$
\begin{aligned}
a_t &= m_t - c_t - d_t - \kappa \cdot 1\{d_t \neq 0\} \geq 0 \\
m_t &= R a_{t-1} + y_t \\
n_t &= R_n \left(n_{t-1} + d_{t-1}\right) \\
c_t &> 0
\end{aligned}
$$

where y_t is log income and m_t is amount of liquid resources available once liquid savings are also considered. We solve the problem by adapting the Nested Endogenous Grid Method (NEGM) (Druedahl and Jørgensen 2017). We estimate the model at a monthly frequency to mirror our empirical estimates in Section 3.

We calibrate several parameters based on our administrative bank sample or representative survey data. Unemployment rates and lifecycle income growth patterns by race are taken from the CPS, while wealth to income ratios are estimated using the SCF among households who are over 18 and banked. The variances of monthly transitory and permanent income shocks while working are calibrated, conditional on unemployment rates and UI replacement rates, to match a monthly coefficient of variation of 0.38, estimated previously using JPMCI data (Farrell, Greig and Yu 2019). We set the Social Security replacement rate to 45 percent, which matches average replacement rates (Mitchell and Phillips 2006), and household income falls to 57.5 percent during unemployment (Zedlewski and Nichols 2012). The retirement age is set to 62, which reflects the median retirement age of initial benefit receipt (Biggs and Springstead 2008).

Our remaining parameters are set based on prior literature. Monthly discount factors are set to match an annualized discount rate of 0.937 (Kaplan and Violante, 2022). Because the monthly coefficient of variation is not enough to separately calibrate monthly permanent and transitory shock variances, we rely on prior estimates based on annual income patterns to set their relative magnitudes. Specifically, we find values for the monthly variances that, conditional on the unemployment probability and unemployment replacement rate, generate annual income patterns consistent with Carroll and Samwick (1997), whose estimates imply a ratio of transitory to permanent shock variances of 2. Each agent lives to at most age 85, with the first period in the model representing the beginning of one's working life, i.e. age 25. We set the CRRA parameter to 1, i.e. a log utility functional form.[3] See Table A-19 for a complete list of calibrated parameters and their sources.

In order to match our empirical elasticity estimates and distribution of assets, we set a discrete, joint distribution of interest rates on liquid assets and illiquid assets with three mass points: $\left(R^j, R_n^j\right)$, $j \in \{1, 2, 3\}$. This generates a low, middle, and high asset type, who have, respectively, decreasing sensitivity of consumption to income volatility. The variation in interest rates within- and across groups is meant to capture differences in returns to assets, access to credit, investment opportunities, tax preferences, and other financial frictions that may affect the accumulation of liquid and illiquid wealth. As is the case in Kaplan and Violante (2022), our lowest types will generally face negative nominal interest rates on liquid assets, and all types will face a nontrivial spread between the liquid and illiquid return. An alternative approach to generating asset heterogeneity is to vary the discount factors, but since interest rates and discount factors effectively perform the same role in the economic model, it is not surprising that we can achieve similar results by varying interest rates.

[3]We could have also chosen a higher CRRA parameter. In our model, this results in the need for even more negative interest rates in order to continue to match the distribution of liquid assets.

Without any strong reason to choose a model with discount factor heterogeneity over one with interest rate heterogeneity, we proceed with the latter. To better capture the difference in asset levels between each type of household, we similarly assign differences in income levels, with the lowest type earning 1/8 of the middle type, and the highest type earning 5 times as much as the middle type. These ratios are taken from IRS Statistics of Income household income tables. The lowest type of households comprise the lowest third of households, the middle type spans the middle 42 percent, and the highest type represent the top 25 percent.

A summary of our model fit and calibrated parameters is presented in Table A-20. A primary goal is to find a set of parameters that is consistent with our key moment, the consumption elasticity. Using a mixture of three types of agents, we are able to do so. Although we do not calibrate the model to match a marginal propensity to consume (MPC), when we simulate an MPC from a one-time \$500 cash transfer, our simulated MPC of 0.21 is within the range of empirical estimates in prior literature, 0.15 to 0.25 (Kaplan and Violante 2022). We additionally target moments of the wealth distribution. Following Kaplan and Violante (2022) we target a mean total wealth to income ratio of 4.3 and median ratio of 1.7. We are able to achieve a mean ratio of 4.2, but we are less able to fit the right-skew of the data. Our median total asset ratio of 3.9 is significantly higher than our target ratio. This inability to generate enough skew using a two-asset model is also common in the literature (see Kaplan and Violante 2022). Likewise, the ratio of liquid assets to income, both at the mean and the median, are significantly higher than their empirical counterparts.

Our calibrated interest rates for liquid assets include a negative value for the lowest group. The observed level of liquid assets, relative to income, tend to be lower than a standard buffer-stock model would predict. Our model therefore uses negative interest rates to fit the liquid asset distribution and to generate a type that has a high consumption elasticity. We interpret this as capturing general frictions in financial markets. More concretely, negative interest rates can arise in cases where financial costs, such as ATM fees, account maintenance fees, and overdraft fees, effectively impose negative real returns on cash balances. Likewise, a negative interest rate may capture cases where liquid assets are at risk of being borrowed by friends or family members who face negative shocks, i.e. kinship taxation (Chiteji and Hamilton, 2002; Squires, 2016). The negative interest rate of -3% is within the range of rates considered in prior literature (Kaplan and Violante 2022).

In order to induce savings in the illiquid account, we require a relatively high return on those assets. For the two types who do use the account—the middle and high-asset households—interest rates are set at 9% and 11%, respectively. The average spread between the liquid and illiquid return in our case, 7.6%, is comparable to other common, two-asset applications (Kaplan and Violante 2014; Kaplan and Violante 2022). A common extension in these cases is to include a flow of utility from holding illiquid assets, which captures the value of durable assets such as owner-occupied housing, which tends to relax the requirement for such a high spread between the liquid and illiquid interest rates.

With our model parameters in hand, we can now construct measures of the welfare benefits of consumption smoothing. Our measure of welfare gain is derived by calculating the expected lifetime utility across all types, and the associated certainty equivalent, that is, the stream of constant consumption that would yield the same level of lifetime utility. Next, we set the standard deviation of monthly transitory shocks and the unemployment probability to zero. We then re-solve the model with agents that are fully aware of the elimination of uncertainty. Relative to our statistical model above, we allow the new consumption path to endogenously adapt to the new economic environment. We once again calculate the certainty equivalent, and compare to the previous case. We report the welfare gain of turning off income volatility as the percent rise in certainty equivalents. We estimate a welfare gain of 0.8% when agents no longer face transitory income shocks or unemployment, which is consistent with our estimate using our statistical model.

When we estimate separate models to match the consumption elasticity and wealth moments for each racial group, we likewise see differences in welfare gains that larger for those with fewer liquid assets. Table A-20 shows the results for each group. The ordering and relative magnitude of welfare gains from our structural model are very similar our estimates in Table 3, where we use a statistical model and a common CRRA parameter of 1. In order to match the relatively high consumption elasticities among Black and Hispanic households, we must assume particularly negative interest rates, ranging between -2% and -6%. As was the case in our baseline model, we are also able to closely match the ratio of total assets to mean income, but do less well in matching the median ratio and liquid asset moments.