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

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

June 10, 2020

Abstract

We study the consumption response to typical labor income shocks and investigate how these vary by wealth and race. First, we estimate the elasticity of consumption with respect to income using an instrument based on firm-wide changes in monthly pay. While much of the consumption-smoothing literature uses variation in *unusual* windfall income, this instrument captures the temporary income variation that households *typically* experience. In addition, because it can be constructed for every worker in every month, it allows for more precision than most previous estimates. We implement this approach in administrative bank account data and find an average elasticity of 0.23, with a standard error of 0.01. This increased precision also allows us to address an open question about the extent of heterogeneity by wealth in the elasticity. We find a much lower consumption response for high-liquidity households, which may help discipline structural consumption models.

Second, we use this instrument to study how wealth shapes racial inequality. An extensive body of work documents a substantial racial and ethnic wealth gap. However, less is known about how this gap translates into differences in welfare on a month-to-month basis. We combine our instrument for typical income volatility with a new dataset linking bank account data with race and Hispanicity. We find that black (Hispanic) households cut their consumption 50 (20) percent more than white households when faced with a similarly-sized income shock. Nearly all of this differential pass-through of income to consumption is explained in a statistical sense by differences in liquid wealth. Combining our empirical estimates with a model, we show that temporary income volatility has a substantial welfare cost for all groups. Because of racial disparities in consumption smoothing, the cost is at least 50 percent higher for black households and 20 percent higher for Hispanic households than it is for white households.

1 Introduction

Forty-two percent of Americans report in the Current Population Survey that they do not have money set aside that could be used for unexpected expenses or emergencies. This statistic raises concerns that it may be difficult for such households to smooth consumption in the face of even modest amounts of income and expense volatility. Yet, perhaps surprisingly, there is little empirical evidence on how typical month-to-month fluctuations in income affect consumption, and even less evidence about how this consumption smoothing varies with wealth.

Concerns about the ability to smooth consumption are particularly pronounced for black households. Fifty-five percent of black households say that they do not have savings for unexpected shocks, as compared to thirty-eight percent of white households. This black-white disparity is one manifestation of the racial wealth gap in the US. 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 similarly high share of Hispanic households do not have savings for unexpected shocks, and similar gaps exist in liquid wealth between Hispanic and white households. Although the racial wealth gap has been extensively documented, there is little quantitative research about whether and how this gap affects households' lives from one month to the next.

The goal of this paper is to make progress on these questions by constructing precise estimates of the consumption response to typical labor income shocks and investigating how these vary by wealth and race. This goal has been difficult to achieve in prior work due to data and methodological challenges. First, existing datasets suitable for studying the pass-through of income to consumption either have small sample sizes or do not include information on both wealth and race. This has made it difficult to obtain precise estimates of the average sensitivity of consumption to income, and especially difficult to uncover how this varies by both race and wealth. Second, most feasible quasi-experimental research designs for convincingly identifying this pass-through have relied on *unusual* income variation such as lotteries and tax refunds. Researchers may have therefore been hesitant to extrapolate from these estimates to make conclusions about the welfare cost of *typical* income shocks.

To overcome the data challenge, we use a novel, large, de-identified dataset with information on income, consumption, liquid wealth, and race. The first three variables come from Chase bank account records. Self-reported race and Hispanicity information come from voter registration records.[1] These two datasets were matched using personal identifiers and then

---

[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" as mutually exclusive categories. Going forward, we use the shorthand "race"

delivered to the research team stripped of these personal identifiers. This is the first time we are aware of that administrative bank account data have been linked with race.[2] Relative to similar data, the two key strengths of this dataset are a large sample size of nearly two million households and employer identifiers which enable our main empirical strategy. 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 US as a whole.

To overcome the methodological challenge, we instrument for changes in individual labor income using employer-wide changes in monthly pay. The three most common reasons for income volatility reported by employed households are, in descending order, irregular work schedules, bonuses and commissions, and seasonal employment. Employer-wide pay shocks capture all these sources of variation and can be identified using a generalization of the classic model of firm effects by Abowd, Kramarz and Margolis (1999). Our research design therefore captures the type of income variation that households typically experience. Furthermore, because the instrument can be constructed for all workers in every month, it yields more precise estimates than were usually feasible in prior work. This design builds on prior strategies to examine the consumption response to firm-level shocks by Shea (1995), Baker (2018) and Koustas (2018).

We implement this research design in our new dataset and present four reduced-form empirical results. We then examine their welfare implications using a model.

First, we find that typical labor income fluctuations meaningfully affect consumption for the average household. Our central estimate is that the elasticity of consumption of nondurable goods and services with respect to income is 0.23. This estimate is precise, with a standard error of 0.01. We show that our instrument captures temporary fluctuations in pay, and so we interpret this elasticity as capturing the consumption response to *temporary* income shocks. We analyze the impact of *persistent* shocks in a secondary research design discussed below.

Second, we find that the consumption of black and Hispanic households is 20 to 50 percent more sensitive to income shocks than the consumption of white households. Specifically, we estimate elasticities of 0.29 for black households, 0.24 for Hispanic households, and 0.20 for white households. As is conventional in the literature on racial inequality, these estimates capture heterogeneity by race in a causal estimate; they do not capture the causal effect of race.

---

to describe responses to the question. We therefore use "Hispanic" to refer to respondents who have selected that identity 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 US over time.

[2]See Section 2.1 for more detail on the matching process and steps taken to protect privacy. See Akee, Jones and Porter (2019), Chen et al. (2019) and Chetty et al. (2018) for examples of other papers that analyze racial inequality using administrative data.

Third, we find that the extent of consumption smoothing also varies dramatically with liquid assets. We find that the elasticity declines monotonically as liquid assets increase, ranging from 0.39 for the lowest-asset households to 0.12 for the highest-asset households.

Fourth, we find that racial inequality in wealth can 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, the black-white difference in elasticities falls to 0.01; this eliminates 89 percent of the raw black-white elasticity gap of 0.09. Similarly, after controlling for liquid assets, the Hispanic households actually have a smaller elasticity than white households.

One way to interpret this 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 or cancel each other out. These 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.

However, this finding should not be interpreted to mean that race is irrelevant for inequality in consumption smoothing. Rather, it indicates 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.[3] 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 income fluctuations than white households.

These findings also hold in a second research design using income declines arising from unemployment. This design is a useful complement for two reasons. First, while our primary design captures the typical sources of income volatility for *employed* households, transitions *between employment states* are another meaningful source of volatility. Second, because of downward nominal pay rigidity, people with stable employment rarely have large and persistent negative shocks to earnings. Identifying such shocks requires job separations, which we measure through the receipt of unemployment insurance benefits. White, black, and Hispanic households experience similar income drops after involuntary job loss. Nevertheless, as in the primary research design, we find substantially 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.

In the final section of the paper we consider the welfare implications of our results. We combine our reduced-form empirical estimates with a model to quantify the welfare cost of idiosyncratic temporary income volatility, and how it differs by race. We consider a statistical model in the spirit of Lucas (1987). Three sufficient statistics are required to calculate the welfare cost: the elasticity of consumption with respect to temporary income shocks, the

---

[3]We discuss prior literature on the determinants of the racial wealth gap in Section 3.6.

coefficient of relative risk aversion, and the variance of temporary income shocks.

We find that income volatility has a substantial welfare cost for all groups, but that the cost is substantially larger for black and Hispanic households than it is for white households. Specifically, we calculate the welfare gain from eliminating idiosyncratic income volatility in terms of lifetime consumption equivalents. Evaluated at our empirical estimates, we find a gain of 1.6 percent for white households, 1.9 percent for Hispanic households, and 2.3 percent for black households. We find qualitatively similar patterns in relative welfare gains using a structural model, albeit with even higher baseline values. Relative to the benchmark view in Lucas (1987) that a cost of 0.5 percent of lifetime consumption is "large", these estimates constitute a substantial welfare cost for all groups, and an especially large cost for Hispanic and black households.

Methodologically, our study seeks to capture the strengths of two distinct strands of the consumption smoothing literature. One strand, pioneered by Blundell, Pistaferri and Preston (2008), uses *all* variation in income and defines income shocks by modeling the earnings process.[4] However, this methodology 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.

A second strand uses exogenous but *unusual* variation arising from windfalls. Examples of such windfalls include life insurance dividends, restitution payments, tax rebates, veterans' bonuses, lottery winnings, and Alaska Permanent Fund payments.[5] However, the marginal propensity to consume (MPC) out of windfall income may differ from the MPC out of other income if recipients use a mental accounting heuristic. In fact, anomalously high spending out of windfalls is used to motivate the canonical mental accounting papers (Thaler 1985; Thaler 1990). The same concern may apply to modern MPC studies; 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). Recent evidence shows that labels can have dramatic effects on how funds are spent (Hastings and Shapiro 2018; Beatty et al. 2014). Because research designs that use *typical* variation in labor income are not subject to this labeling concern, they are particularly useful for understanding the consequences of income volatility.

An added benefit of our methodology is statistical precision, which enables us to address two types of ambiguity surrounding estimates in the prior consumption smoothing literature. First, as noted in Kaplan and Violante (2010), reliable measures of the *average* consumption response to income shocks are central for macroeconomic models. However, a recent meta-

---

[4]See also 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).

[5]See Bodkin (1959), Kreinin (1961), Johnson, Parker and Souleles (2006), Hausman (2016), Fagereng, Holm and Natvik (2018), and Kueng (2018), respectively.

analysis by Havranek and Sokolova (2020) finds that published MPC estimates dramatically overstate the true MPC after correcting for publication bias. Such concerns may be muted for our estimates because the standard error is less than one cent, and therefore even economically insignificant consumption responses will be statistically significant at conventional levels.

Second, looking beyond the average MPC, there is even more ambiguity in the consumption-smoothing literature about *heterogeneity* in MPCs by liquid assets. A central implication of many consumption models is that there is a "*tight negative correlation* between the size of the consumption response and the ratio of holdings of liquid wealth to income" (Kaplan and Violante 2014, emphasis added).[6] However, it has been difficult to statistically distinguish between high and low asset MPCs in prior work. Furthermore, recent literature using variation from infrequent windfalls finds substantial consumption responses for high-asset households that are nearly as high as those for low-asset households (Kueng 2018; Fagereng, Holm and Natvik 2018). These results have spurred the development of alternative models designed to explain excess sensitivity for high-liquidity households (Lian 2019). In contrast, we find that the MPC out of typical income fluctuations is four times larger for low-asset households than for high-asset households. Our estimates therefore provide renewed empirical support for this fundamental prediction of standard models.

Our welfare analysis revisits 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). Historically, there has been good reason for this assumption, since the standard one-asset consumption model predicts that most agents accumulate sufficient assets to easily smooth temporary income fluctuations. Yet empirically, low liquid asset holding is prevalent among a substantial share of Americans. Recent work using two-asset models has sought to explain the *causes* of this low liquid asset holding. In contrast, our analysis documents a *consequence* of low liquid asset holding which is intuitive, yet has not been documented in prior work to the best of our knowledge: temporary income volatility has a significant welfare cost. This is because even temporary income fluctuations meaningfully affect consumption when households hold few liquid assets. Furthermore, this cost is highest for groups with low average liquid assets such as black and Hispanic households.

Finally, ours is the first study we know of to measure consumption smoothing by race at a month-to-month frequency. 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 homeownership (Charles and Hurst 2002; Aaronson, Hartley and Mazumder 2017), marriage (Addo and Lichter 2013), and inter-household trans-

---

[6]This is true in both modern two-asset models (Laibson, Repetto and Tobacman 2007; Kaplan and Violante 2014) and also in the canonical buffer stock model (Aiyagari 1994, Deaton 1991, Carroll 1997)

fers (Chiteji and Hamilton 2002); as well as potential policy solutions (Hamilton and Darity 2010). Most studies focus on differences in *levels* of wealth and longer term trends in these gaps.[7] However, less is known about how wealth gaps translate into differences in month-to-month economic welfare. We contribute to this literature by examining the relationship between the racial wealth gap—and by extension gaps in liquid assets—and the transmission of income *volatility* to consumption, thus linking the literatures on racial inequality and consumption smoothing.

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 model, and 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 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 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 20 million households each month with a Chase checking account. The unit of observation is household-by-month, from October 2012 through April 2018.

Our analysis centers on three variables: 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

---

[7]However, see Hardy et al. (2018) for an exception in the case of income volatility.

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, haircuts, 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 US 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 financial assets increase one-for-one with average checking account balances, as shown in online Appendix Figure 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 online Appendix Table 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 US 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 states subject to additional scrutiny

under the Act as well as the set that had white-only primaries in the Jim Crow era. Online Appendix Figure 2 shows a map of these states. The voter files also contain name, address, and date of birth, which we use for matching.[8]

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 (e.g. 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 online Appendix Table 2 and the distribution of responses is shown in online Appendix Table 3.

2.2 Matched Dataset with Bank Account Data and Voter Registration

A JPMorganChase 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.[9] We discuss details of the algorithm in online 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.[10] Thus, the number of observations in the anal-

---

[8]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.

[9]This 56 percent estimate is the ratio of 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 2018, 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.

[10]Federal law requires that mortgage lenders collect data on race. Fuster et al. (2018), Bartlett et al. (2019), and Bhutta and Hizmo (2019) 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. In related work, Dobbie et al. (2018) documents discrimination against

ysis 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.[11]

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 second about their Hispanicity. There are 194,000 de-identified Chase mortgage applications that JPMorgan Chase is able to match to the voter files.[12]

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. Online 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 US 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.[13]

---

immigrants in the UK.

[11]Online Appendix Table 4 and online Appendix Figure 3 demonstrate this concentration in urban areas. 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. Online Appendix Table 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.

[12]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.

[13]To align with the mutually exclusive concept of race and Hispanicity from the voter files, we allocate

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 US 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.

In practice, we find that black-white income inequality in the sample frame is quite similar to national patterns. Figure 1a 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 higher: 77 cents. This is 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 online Appendix Figure 4a). 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 US is that black household heads are more likely to be female than male; Figure 1b shows that this also holds in the sample frame. In addition, Figure 1b 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 US 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. Online Appendix Figure 6 compares the countries of origin for Hispanics across the US 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

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. (2018), Card and Rothstein (2007), and Federal Reserve Board (2017).

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 1a shows that the Hispanic-white income ratio is similar in the sample frame to the nation and online Appendix Figure 4a shows that the distribution of income is similar. Figure 1b 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 online Appendix Figure 4b).

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 1a 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 for black households.[14]

One striking fact visible in Figure 1a 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.[15] 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. As we discuss in the introduction, 42 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

---

[14]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 1a, 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 analyze people 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.

[15]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.

to 38 percent of white households, for a black-white gap of 17 percentage points (see Online Appendix Figure 7a). 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 Online Appendix Figure 7b). 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 1a shows that racial gaps in income and assets are similar. Figure 1b 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. Online Appendix Figure 4 compares the distribution of income and age in the sample frame and in the analysis sample.

To summarize, we use a novel, large de-identified dataset with information on income, consumption, liquid wealth, and race. We accomplish this by matching bank account records and voter files in Florida, Georgia, and Louisiana. The matching procedure is accurate. The well-known and large racial economic gaps in the US 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

To assess whether and how much consumption smoothing varies by race and assets, we first provide descriptive analysis of the matched dataset described above. After describing our primary identification strategy, we estimate the causal impact of typical labor income variation on consumption. We report results for the pooled sample, separately by race, separately by assets, and by race and asset together. Finally, we report results from a second identification strategy that uses unemployment.

3.1   Ordinary Least Squares

For our descriptive analysis, we use a first-difference ordinary least squares (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 (BPP 2008). In Section 4 we show that this elasticity is one of the sufficient statistics needed

for measuring the welfare cost of income volatility. Table 1 provides summary statistics for our analysis sample.

This equation identifies β if

Assumption EO (Earnings Orthogonality)

$$E(\varepsilon_{it}|\{\Delta y_{it}\}_{t=1}^{T}) = 0$$

holds. This assumption says that income changes in every month are strictly exogenous with respect to the unobserved error ε_{it}, where the data span T months. 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 2 column (1) reports an estimate for $\hat{\beta}$ of 0.12, which means that a 10 percent change in income is associated with a change in consumption of 1.2 percent. The relationship is present for both positive and negative income changes as well as small and large income changes (see online Appendix Figure 8a).

This correlation between income and consumption is about twice as strong for black households and fifty 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 2 column (2) reports estimates for $\hat{\beta}_1$ of 0.09, $\hat{\beta}_2$ of 0.08, and for $\hat{\beta}_3$ of 0.04. Online Appendix Figures 8b, 8c, and 8d show non-parametric binscatter plots of the same relationships.

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. In particular, labor supply changes are problematic for Assumption EO because of reverse causality. $\hat{\beta}$ will be biased upward if a worker increases her labor supply to finance a specific consumption need. $\hat{\beta}$ will be biased downward if leisure and market consumption are substitutes (working additional hours lowers desired contemporaneous market consumption) or if there is classical measurement error in income (Bound et al. 1994). The overall sign of the bias from these different forces is ambiguous.

The endogeneity of income and measurement error has motivated researchers to find instruments for income that are plausibly orthogonal to ε_{it}. However, many such papers rely on sources of income that arrive infrequently and are labeled for a specific purpose. This may limit the applicability of these estimates for measuring the consumption smoothing response to typical income shocks.

3.2 Firm Pay Shock as Instrument for Labor Income

To measure the causal effect of typical income variation on consumption, we use variation in monthly firm-level pay. This strategy avoids contamination from the worker-specific factors discussed above. It is inspired by the labor economics model of firm effects pioneered by Abowd, Kramarz and Margolis (henceforth AKM, 1999). Consider the data-generating process

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

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. (2020) in generalizing AKM to allow for *time-varying* firm effects $(\psi_{j(i,t),t})$, while the most of the AKM literature assumes *time-invariant* firm effects $(\psi_{j(i,t)})$.[16]

What are the important sources of variation in monthly firm pay $(\psi_{j(i,t),t})$? External data suggest that hours volatility is the leading source of month-to-month variation in firm pay. The Survey of Household Economics and Decisionmaking (SHED) asks households about whether they face income volatility and specific reasons for volatility. By far, the most commonly-cited reason for income volatility is "irregular work schedule".

A second source of volatility is from asymmetric changes in compensation such as bonuses and commissions. Grigsby, Hurst and Yildirmaz (2019) and Lemieux, MacLeod and Parent (2009) document 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 tabulate all the responses to this question in online Appendix Table 6. We also note that workers who are paid a higher hourly rate for additional hours ("overtime") will generate asymmetry in the income shocks. Online Appendix Figure 9 documents this asymmetry in the JPMCI data: positive deviations of monthly pay from the firm median are larger than negative ones.

Variation in firm pay arises from two types of economic fundamentals: changes in labor demand and firm rent-sharing. A broad and active literature regresses changes in annual earnings or wages on *observable* firm-level shocks such as value added or revenue (Guiso, Pistaferri and Schivardi 2005; Kline et al. 2019; Lamadon, Mogstad and Setzler 2019; Friedrich et al. 2019; Juhn et al. 2018a; Juhn et al. 2018b). Some of these papers interpret the regressor as isolating firm rents, while others interpret the regressor as capturing both rents and changes in labor demand. Our empirical strategy captures both of these observable channels. It also captures changes in labor demand which occur for reasons not captured by such regressors

---

[16]Lachowska et al. (2020) 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. They conclude that firm effects are quite stable from year to year. One way to jointly interpret our results and this prior work is that the month-to-month fluctuations in firm pay captured by our research design do not reflect the type of permanent changes that would show up in a specification with time-varying firm effects estimated using data aggregated to the annual level.

(e.g. when labor demand today responds to future expectations about product demand).

We estimate firm pay shocks using a first-difference estimator that regresses changes in a worker's own pay on the change in the average pay of their coworkers. We estimate

$$\Delta y_{it} = \phi + \rho \Delta y_{j(-i,t),t} + \sum_{s=t-L}^{t-1} \omega_{s-t} \Delta y_{j(-i,t),s} + \nu_{it} \tag{4}$$

where $\Delta y_{j(-i,t),t} \equiv \frac{1}{\sum_i \mathbf{1}(i \in j)-1} \sum_{i' \in j, i' \neq i} y_{i't} - y_{i',t-1}$ is a leave-out mean that uses the coworkers at firm j of worker i and L is the number of lags. To improve precision, coworker pay is estimated using everyone with a Chase bank account, not just the workers for whom we have matched race.

This specification controls for several lags of the change in coworker pay. Without lags, innovations in $\psi_{j(i,t),t}$ would affect the regressor $\Delta y_{j(-i,t),t}$ twice. For example, if $\psi_{j(i,t),t} = 0.1$ at time t and $\psi_{j(i,t),t}$ is zero in every other period ($\psi_{j(i,t),t'} = 0 \quad \forall t' \neq t$), 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.[17] The inclusion of lags enables us to interpret our results as capturing innovations in income, and not simply the reversion of innovations from the prior period. The estimates reported in this section use five lags ($L = 5$). However, the paper's key quantitative results continue to hold with other choices of lags. In Online Appendix Table 7 columns (1)-(3), as a robustness check, we report the paper's central estimates without controlling for any lags. In unreported results, we find that our estimates are very similar in specifications that control for between one and twelve lags.

Firm-level pay shocks estimated using coworkers have strong predictive power for individual pay, as shown in Figure 2a. Nearly every one of the binscatter points falls on the linear regression line, indicating that the relationship in means is almost exactly linear. Consistent with the asymmetry in pay fluctuations discussed above, Figure 2a shows an asymmetric distribution in the instrument: positive shocks are larger than negative ones. We estimate $\hat{\gamma}$ of 0.70, which means that a worker's own pay changes by 7 percent when their coworkers' pay changes by 10 percent. Furthermore, the firm pay shock instrument accounts for a sizable fraction of the total volatility of household income. We calculate that firm effects account for between one-fourth and one-fifth of volatility in monthly take-home earnings.

The firm pay shocks appear to be transitory innovations. We demonstrate this by examining the change in income for months around the shock. Online Appendix Figures 10 and 11 replicate Figure 2a, changing the dependent variable to lags and leads of the change in individual pay. Here, and throughout the analysis, we require that the worker is in the sample in the ten months prior to the coworker shock.[18] We find little change in income in

---

[17]BPP solve this problem in an alternative way: by assuming an earnings process with temporary and permanent shocks and implementing an estimator that, in the context of this assumption, captures only the innovations to earnings.

[18]We also require that the worker be employed at the same firm in the ten months prior for computational

the months leading up to the shock, suggesting that our specification does in fact capture innovations in income. We also find that income has fully reverted to its pre-shock level by one month after the shock.

Although most consumption smoothing research uses all variation in income or variation in windfall income, there are at least three prior papers that instrument for individual income using variation in firm pay. The closest antecedent of the research design described above is the design in Koustas (2018), who also uses coworker pay as an instrument. He studies people who use an online financial aggregator and drive for Uber or Lyft (as well as their coworkers), and uses a panel specification with firm fixed effects to instrument for variation in individual pay. We build on Koustas (2018) by studying a broader sample, providing formal orthogonality conditions, reporting nonparametric and parametric estimates of the first stage of the instrument, documenting the asymmetric distribution of the instrument, and reporting pre-trends and post-trends for income and consumption for several periods. In related work, Shea (1995) and Baker (2018) instrument for individual pay using explicitly *observable* firm-level changes in union contracts and stock prices respectively.

We measure the effect of transitory income shocks on consumption using a linear two-stage least squares estimator that uses a modified equation (1) together with equation (4):

$$\Delta c_{it} = \alpha + \beta \Delta y_{it} + \sum_{s=t-5}^{t-1} \tau_{s-t} \Delta y_{j(-i,t),s} + \varepsilon_{it} \tag{5}$$

$$\Delta y_{it} = \phi + \rho \Delta y_{j(-i,t),t} + \sum_{s=t-5}^{t-1} \omega_{s-t} \Delta y_{j(-i,t),s} + \nu_{it}. \tag{6}$$

These equations identify β if $\rho \neq 0$ (as Figure 2a shows to be the case) and the condition

Assumption FO (Firm-pay Orthogonality)

$$E(\varepsilon_{it} | \{\psi_{j(i,t),t}\}_{t=1}^{T}) = 0 \tag{7}$$

holds. Assumption FO would fail if the firm is large enough that its payroll changes have a noticeable effect on the regional economy through changes in prices (e.g. housing prices or local prices of non-traded goods). It would also fail if households' consumption is responding to news about the firm's future payroll. However, the fact that in a statistical sense the firm pay shock has essentially no impact on the subsequent months suggests that there is little or no news in the shock.[19]

---

convenience. We will drop this requirement in a future draft.

[19]A third way that this assumption can fail is if a change in the average pay of all workers at a firm has affects consumptions through peer effects. De Giorgi, Frederiksen and Pistaferri (2020) document that idiosyncratic 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 Jones" 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 same income shock.

Assumption FO is weaker than Assumption EO, which identifies the OLS model discussed in the previous section. 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 therefore the pair of assumptions

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

together imply Assumption EO. Thus, Assumption EO nests both Assumption FO and a second assumption $(E(\varepsilon_{it}|\{\eta_{it}\}_{t=1}^{T}))$. 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 Causal Impact of Income on Consumption

We use the firm pay pay shock as an instrument to measure the causal effect of income on consumption using an instrumental variables (IV) specification. Figure 2b shows a binscatter of the reduced form and 2c shows the second stage. The relationship is precisely estimated. One striking feature of this plot is that even large changes in income (such as a $500 decrease or increase) cause changes in consumption. This pattern suggests that when income dips, households either do not have or choose not to access illiquid assets to smooth consumption. This motivates a focus on heterogeneity by liquid assets in our subsequent analysis.

To provide a quantitative interpretation of the evidence in Figure 2, Table 2 column (3) reports an IV estimate of an elasticity of consumption to income of 0.23 ($\hat{\beta} = 0.23$). This estimate is larger than the OLS estimate. We examine the implications of this IV estimate for the welfare cost of income volatility in Section 4.

This elasticity is precisely estimated, with a standard error of 0.009. The comparable MPC estimate is 21.9 cents with a standard error of 0.9 cents. In a recent meta-analysis, Havranek and Sokolova (henceforth HS, 2020) reviews 246 estimates of the MPC for non-durables. The mean (median) MPC in these prior papers is 29 (25) cents with a mean (median) standard error of 19 (13) cents. Online Appendix Figure 12 plots the distribution of standard errors for MPC estimates in the HS literature review. None of the papers reviewed by HS have an MPC less than one cent in their preferred specification.[20] The standard error of our estimates is sufficiently small that even modest consumption responses that are economically insignificant will be statistically significant at conventional levels, thereby

---

[20]The most precise estimates in prior work, which have standard errors close to one cent in their primary specification, come from papers that estimate daily or weekly MPCs using within-month variation in income. Olafsson and Pagel (2018) and Gelman et al. (2014) examine the timing of spending relative to payday. Baker and Yannelis (2017) and Gelman et al. (2018) examine the effects of a delayed paycheck due to the federal government shutdown. Thus, our estimates for the monthly MPC have similar degree of statistical precision to the literature on the daily or weekly MPC and a substantial gain in precision relative to the prior literature on monthly MPCs.

allaying the publication bias concern raised by HS. Indeed, given the argument in HS, it is not surprising that our point estimate is smaller than the average in prior studies.

Our elasticity estimate of 0.23 implies less consumption smoothing of temporary income shocks than the estimates in BPP, which estimates an elasticity of consumption to temporary income shocks of 0.05. Our estimates could differ for three reasons: (1) 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), (2) BPP uses leads of the change in income to isolate temporary income shocks, and/or (3) the elasticity of consumption with respect to (exogenous) firm pay shocks differs from the elasticity with respect to (potentially endogenous) household income. We run two alternative regression specifications to test between these potential explanations. When we implement BPP's estimator in our data using all income shocks, we estimate passthrough of 0.17. When we implement BPP's estimator using firm pay shocks, we estimate a passthrough of 0.24; thus, BPP's estimator applied to firm pay shocks yields the same conclusion as our preferred specification. Our estimate of higher passthrough is therefore driven mostly by explanation 1 (having a measure of income from administrative data) and to a lesser degree by explanation 3 (using exogenous variation from firm pay shocks instead of using all variation in income).

Just as the impact of the firm pay shock on income is temporary, so too is its impact on consumption. Online Appendix Figure 13 and 14 replicate Figure 2b, changing the dependent variable to lags and leads of the change in household consumption. This is useful for evaluating Assumption FO. One way that this assumption could be violated is if households modify their consumption in anticipation of the firm pay shock. The absence of pre-trends in consumption suggests that there is not such a violation. We also find little to no consumption response in the months following the firm shock. This suggests that that whatever fraction of the income deviation that is not spent in the same month as the shock is saved.

One interesting question is to what extent the firm pay shocks are predictable from the perspective of households. We do not have a good way to answer this question with our data and this is an important question for future research. However, we note that the predictability of the shock does not necessarily imply a violation of Assumption FO. Johnson, Parker and Souleles (2006) and Kueng (2018) are examples of papers with predictable income shocks and an assumption similar to Assumption FO.

Although we do not require that the firm pay shocks be entirely unpredictable, we are nevertheless interested in isolating the unpredictable component of those shocks. We therefore construct an alternative version of the instrument, where we residualize $\Delta y_{j(-i,t)}$ to be the deviation of coworker pay from its calendar month average (e.g. for each firm, compare this February to the same firm's other Februaries). This requires firm panels of at least 24 months. This specification seeks to eliminate firm-specific seasonal variation in payroll and annual bonuses as a source of identifying variation. As a result, this specification amplifies the importance of hours volatility, which is the remaining source of identification. In Online

Appendix Table 7, we report the paper's central estimates using the specification with this alternative instrument and find similar results to our preferred specification.

3.4 Consumption Smoothing by Race

Figures 3 and 4 provide non-parametric evidence that when a firm raises or lowers its pay for all workers, 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 (5) and (6), interacting Δy_{it} and $\Delta y_{j(-i,t),t}$ with $Black_{it}$ and $Hispanic_{it}$ and adding categorical dummies by race. Two features of this specification are worth noting. First, we re-estimate the first stage equation separately by race, but the estimates are very similar and so differences in Figures 3 and 4 almost entirely reflect differences in the reduced form relationship between the instrument and consumption. Second, as is typical in economic studies of racial inequality, this specification captures the causal impact of a variable (in our case, income) for different racial groups, not the causal impact of race itself.

We find that the causal effect of an income shock on consumption is about 50 percent larger for black households than it is for white ones, and about 20 percent larger for Hispanic households than for white ones. Table 2 column (4) estimates $\hat{\beta} \times Black$ of 0.09 and $\hat{\beta} \times Hispanic$ of 0.04, where $\hat{\beta}$ is 0.20, which captures the elasticity for omitted race category of white. To give a sense of the difference in magnitudes in dollar terms, the median absolute deviation in labor income is \$1,178. Our estimates imply that a decline in income of this size will lead to a decline in consumption of \$236 for white households, \$286 for Hispanic households, and \$344 for black households.

3.5 Consumption Smoothing by Assets

As we discuss in the introduction, there are several mechanisms that might give rise to differences in elasticities by race. In this paper, we focus on just one candidate mechanism: differences in liquid assets. We focus on this variable because MPC heterogeneity 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 3.6, 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 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 some households with high assets also have high consumption, such that they have little financial buffer to smooth consumption in response to income shocks.

Figure 5 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 2c for the bottom and top deciles. The slope is quite steep for the bottom decile (0.39 in Figure 5a) and shallow for the top decile (0.12 in Figure 5b). The consumption changes for low-asset households are symmetric. They cut spending when income drops, presumably because they lack the assets needed to smooth consumption. They also increase spending when income rises, perhaps because they have more deferred consumption needs when a positive income shock arrives.

One interesting pattern is that high-asset households have an asymmetric consumption response to firm pay shocks. Visual inspection of Figure 5b shows that high-asset households with positive income shocks increase consumption, but do *not* appear to decrease consumption when income falls. This finding is consistent with evidence in Baugh et al. (2018).

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.[21]

The extent of MPC heterogeneity by liquid assets has been difficult to precisely quantify in prior work. Figure 6 illustrates the statistical uncertainty of estimates in Parker et al. (2013), Fagereng, Holm and Natvik (2018), and Kueng (2018).[22] 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 MPCs decline steeply as liquid wealth rises.

---

[21] An alternative strategy would be to estimate equations (5) and (6) in levels instead of logs. We report such estimates in online Appendix Table 8, and also 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. However, these are not our preferred point estimates because a dollar-based specification upweights the consumption response of high-wage households, 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. Similarly, our main specification in logs is not susceptible to this concern since pay shocks are similar in percentage terms across the income distribution, and hence our preferred strategy is to scale these elasticity measures into MPCs.

[22] 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 (2018). Misra and Surico (2014) document heterogeneity in the MPC among people with high mortgage debt. BPP and Fisher et al. (2019) 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.

One lesson from our estimates is that the nondurables MPC out of typical income variation for high liquidity households is substantially smaller than the MPC of low liquidity households. This finding contrasts with prior work studying the response of high-asset households to infrequent windfalls. Kueng (2018) emphasizes the high 21-cent MPC out of Alaska Permanent Fund checks among top quintile households, and the lack of a steep gradient across quintiles, as inconsistent with standard consumption models. Fagereng, Holm and Natvik (2018) estimates a total consumption MPC of 46 cents out of lottery winnings for the top quartile households (as compared to 62 cents for the bottom quartile households). For comparison, we estimate a 10.6 cent MPC for top quartile households that is only one-fourth as large as the response for bottom-quartile households. One possible reason for these differences is that our estimates rely on typical income variation rather than income windfalls. If households use mental accounting heuristics when deciding how much of a windfall to spend, their MPC may differ from the MPC out of regular income. An interesting path for future work would be to measure the response to both types of income shocks for the same population, in the same dataset, using the same consumption measures.

Our results are consistent with a literature documenting MPC heterogeneity in covariates that are likely to be correlated with liquidity constraints. Gross, Notowidigdo and Wang (2020) documents higher MPCs during the Great Recession and interpret their findings as consistent with liquidity constraints binding for a larger share of households during this time period. Patterson (2019) shows heterogeneity in MPCs by several covariates that are correlated with liquidity. We discuss the relationship of our estimates to hers in Section 3.7.

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 5 and 6. 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). How important is each channel? On one hand, some prior work studying other outcomes such as health has found that the causal effect of wealth is much smaller than the cross-sectional relationship with wealth (Cesarini et al. 2016). On the other hand, there are strong theoretical reasons (e.g. Aiyagari 1994) to believe that households who have liquid assets will use them to smooth consumption, while those who do not have assets will be unable to smooth consumption. 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.

3.6 Consumption Smoothing by Race and Assets

This subsection explores the link between our first two results: can racial inequality in wealth account for racial inequality in consumption smoothing? As discussed in Section 2 and shown in Figure 1a, 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 3.5, it is possible that racial wealth gaps contribute to the racial gaps in consumption smoothing that we document in Section 3.4.

We use a simple specification to summarize the role of assets in consumption smoothing in columns (5)-(8) of Table 2. We re-estimate columns (3) and (4), adding a linear control for the rank of the household's asset buffer ($Buffer = AssetRank/N - 0.5$). This variable is scaled from -0.5 for the lowest asset household to 0.5 for the highest asset household. Thus, the coefficient on $\Delta y \times Buffer$, captures the change in consumption smoothing from being the lowest-ranked household in terms of assets to being the highest-ranked household. As in the prior columns, the coefficient on Δy continues to capture for the median household in terms of assets. Thus, column (5) implies that the consumption elasticity for the median asset household is 0.23 and that this elasticity ranges from 0.36 for the lowest-asset households to 0.11 for the highest-asset households.

Controlling for observed asset buffers reduces racial gaps in consumption smoothing by a factor of two. Column (6) controls for checking account balance as a measure of asset buffer. The coefficient on $\Delta y \times Black$ falls to 0.05 (from 0.09 in column (4)) and the coefficient on $\Delta y \times Hispanic$ falls to 0.02 (from 0.04 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 affect 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 imputation addresses a problem of differential measurement error: for a given level of checking assets, black and Hispanic households have fewer assets than white households. This is analogous to the procedure in Rothstein and Wozny (2013), which seek to measure the black-white test score gap while controlling for permanent income, but lacks a single dataset with both variables; they 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 all of the racial gaps in consumption smoothing are explained by controlling for differences in liquid assets. These results can be seen in columns (7) and (8) of Table 2. Once we control for liquid assets, the black-white gap in consumption sensitivity to income is largely eliminated and the Hispanic-white gap is reversed. This finding is depicted graphically in Figure 7a. A negative Hispanic-white gap means that, after controlling for liquid assets, Hispanic households do more consumption smoothing than white households. This may reflect Hispanic households receiving more transfers than white households from family and friends to smooth consumption.

We explore robustness to alternative imputation strategies. We re-estimate the specification in Table 2 column (8) using two alternative definitions of assets. We consider financial assets, which are substantially broader than the Kaplan and Violante (2014) definition of liquid assets. We also consider the definition of liquid assets used by SCF summary file, which is narrower than the Kaplan and Violante (2014) definition because the SCF summary file excludes credit card debt. The racial wealth gap is largest for the broadest definition of assets and smallest for the narrow liquid asset definition. Online Appendix Table 7 shows that racial gaps in consumption smoothing are smaller when we use a broader definition of assets and larger when we use a narrower definition. Understanding which non-checking account assets people actually use to smooth consumption is an important topic for future work.

There are two related interpretations of these results. First, this finding suggests that when faced with the same financial circumstances (i.e., 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. 2018; Perry, Rothwell and Harshbarger 2018; Keys, Mahoney and Yang 2020; Stein and Yannelis 2019), 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 frame 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 race.

3.7 Alternative Research Design: Consumption Sensitivity to Unemployment

One limitation of the preceding analysis is that it captures *temporary* income fluctuations that last on average for only one month. Because of downward nominal pay rigidity, people with stable employment rarely have large and persistent negative shocks to earnings. One might imagine that households are better able to smooth larger and more persistent negative income shocks, for example if they are able to tap into an illiquid asset. This motivates a second research design studying unemployment. This is the same event studied in Ganong and Noel (2019). We demonstrate that the same racial gaps exist in consumption smoothing for this larger negative income shock.

Figure 8a 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, which is similar to the average shock below the 5th percentile of the firm pay shock distribution shown in Figure 2a. While the firm pay shock lasts for one *month* on average, Figure 8a illustrates that income remains depressed in the *year* following UI receipt. (In unreported results, we find that incomes recover substantially from month 12 to 24 after UI receipt.)

The elasticity of consumption with respect to income is similar for unemployment than it is for typical income variation. We quantify this using a linear two-stage least squares specification:

$$c_{it} = \alpha + \beta y_{it} + \varepsilon_{it} \tag{8}$$

$$y_{it} = \phi + \rho Post_{it} + \nu_{it} \tag{9}$$

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 ($\hat{\beta} = 0.18$), which is a bit smaller than the elasticity estimate from the firm pay shock design.

Figure 8a also shows that black, Hispanic, and white households experience a similar

drop in income in the year after the onset of unemployment. This might be surprising given the well-known fact that black workers have longer average unemployment durations in the US as well as the results of Fryer, Pager and Spenkuch (2013), who document that black jobseekers in the Great Recession received lower wage offers than white jobseekers. However, these patterns do not hold in our sample. Online Appendix Table 9 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 8a 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 forward—the effect of insured job loss on household income in the year after job separation is similar by race.

Faced with a similar-sized income shock, Figure 8b shows that black and Hispanic households cut spending more than white households. We re-estimate equations (8) and (9), interacting y_{it} and $Post_{it}$ with $Race_i$ and adding a dummy $Race_i$. Table 3 column (2) shows that the elasticity is about 30 percent higher for black and Hispanic households. This racial gap is smaller than the black-white gap and larger than the Hispanic-white gap estimated in our prior research design. It is smaller than estimates of the racial gap in Patterson (2019) using the PSID.[23]

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

Just as with firm pay shocks, it appears that liquid assets are the key channel for understanding racial differences in consumption smoothing during unemployment. Table 3 column (3) shows that assets are strongly correlated with consumption smoothing. Controlling for assets reduces or eliminates the racial gap in consumption smoothing. Table 3 column (4) controls for checking assets and column (5) controls for imputed liquid assets. In the latter column, the racial gap in consumption smoothing is economically small and statistically insignificant. Figure 7b provides a visualization of this finding.

In a companion paper (Farrell et al. 2020), we provide additional details on this analysis of spending around unemployment. The companion paper also repeats approximately the same analysis using a third source of income variation: tax refunds.

Across three research designs, we find substantially less consumption smoothing by black

---

[23]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). She 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 differences in gender composition. 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 1b shows that more than half of black households in the US and in the JPMCI data are women. To ease comparison of our estimates to Patterson's, we will disaggregate our estimates by race and gender in a future draft.

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 identified statistically significant differences in consumption elasticities across racial and ethnic groups, but are these differences economically meaningful? The classic view in macroeconomics is that temporary income shocks are irrelevant for welfare. Indeed, a standard approach in the literature is simply to assume that these welfare costs are zero (see e.g. Heathcote, Storesletten and Violante 2008). This assumption is rooted in the fact that the standard model predicts that most agents accumulate sufficient assets to smooth temporary income fluctuations. Yet empirically, low liquid asset holding (and high credit card debt) is prevalent among a substantial share of Americans and is particularly large for black and Hispanic households. A new generation of two-asset consumption models (Kaplan and Violante 2014; Laibson, Repetto and Tobacman 2007) have been used to explain the pattern of people simultaneously holding substantial illiquid assets and relatively few liquid assets. Our analysis documents a consequence of low liquid asset holding which is intuitive, yet has not been fully quantified in prior work to the best of our knowledge: idiosyncratic temporary income volatility has a significant welfare cost for households with few liquid assets. This cost is highest for groups with low average liquid assets, e.g. black and Hispanic households.

To get a better handle on how quantitatively important these differences may be, 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 deterministic component (κ_t), a permanent random component (z_t), and a transitory random component (θ_t):

$$y_t = \kappa_t + z_t + \theta_t$$
$$z_t = z_{t-1} + \zeta_t, \tag{10}$$

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. 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:

$$c_t = \beta\theta_t + \tilde{c}_t(\kappa_t, z_t, \xi_t), \tag{11}$$

with an error term, ξ_t, orthogonal to (κ, z, θ). Blundell, Pistaferri and Preston (2008) and 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. This specification allows consumption to respond differently to predictable changes in income, permanent shocks, and transitory ones.

We now derive a formula, similar to that of Lucas (1987), to express the welfare gain from shutting off the effect of transitory shocks on consumption. From (11) 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:

$$\mathbb{E}\left[\sum_{t=1}^{T} \delta^t \frac{\left((1+\lambda)C_t\right)^{1-\gamma}}{1-\gamma}\right] \equiv \mathbb{E}\left[\sum_{t=1}^{T} \delta^t \frac{\tilde{C}_t^{1-\gamma}}{1-\gamma}\right]. \tag{12}$$

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

$$\lambda \cong \left[\beta + \beta^2 (\gamma - 1)\right] \frac{\sigma_\theta^2}{2}. \tag{13}$$

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

In order to apply this formula in our empirical setting, we need values for β, γ, and σ_θ^2. We use our firm pay shock estimates of β by race from column (4) of Table 2 and vary the values of γ between 1 and 4. Finally, we calibrate σ_θ^2 using prior estimates of monthly income volatility within the bank data, in combination with external estimates of the relative variances of permanent and transitory shocks at an annual level. That is, if we simulate a monthly earnings process with a value of $\sigma_\theta^2 = 0.16$ and a value of $\sigma_\zeta^2 = 0.0005$, we match the monthly coefficient of variation of 0.38 from Farrell, Greig and Yu (2019), which also uses the administrative bank data. These estimates are similar in magnitude to survey-based estimates by Hannagan and Morduch (2015). Next, if we aggregate the data to an annual frequency and apply a standard estimator, we arrive at a ratio of $\sigma_{\theta,annual}^2 \big/ \sigma_{\zeta,annual}^2 = 0.012/0.006 = 2$, which matches the ratio of estimates from Carroll and Samwick (1997).[24]

Table 4 shows that the welfare gain from eliminating transitory income volatility is sub-

---

[24]This ratio of transitory to permanent shock variances is in line with other estimates using the PSID (Blundell, Pistaferri and Preston 2008), but lower than that of Campbell and Cocco (2015); Laibson, Repetto and Tobacman (2007); and Guvenen et al. (2015). The estimates in all of these prior studies imply a larger monthly variance of transitory shocks than the one based on our administrative bank data. The welfare gains we estimate would be larger were we match the variance of transitory shocks in these studies, or if we were to allow for serial correlation in transitory shocks.

stantial: at least 1.5 percent for every racial group. Although Lucas (1987) conjectures that the welfare costs of transitory income volatility might be large, we are unaware of any prior estimate of these costs.[25] Some prior papers assume the cost is zero (Heathcote, Storesletten and Violante 2008) and the assumption that these costs are trivial is in line with previously-available evidence using annual data. For example, using equation (13) with the estimates of the passthrough of temporary income shocks to consumption in BPP, implies an estimate of $\lambda \approx [0.05 + 0.05^2(\gamma - 1)]\frac{0.012}{2}$, which is equal to 0.11 percent when γ is 1. This is over an order of magnitude less than our estimate.

Two forces contribute to our much higher estimate of the welfare impact of temporary income volatility. First, we find a higher elasticity of consumption to income than BPP, as discussed in Section 3.3. Second, while a welfare calculation based on evidence at an annual frequency assumes that income is constant from month-to-month, our analysis incorporates within-year income volatility.

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 black welfare gain is 46 percent larger than the white one. The Hispanic welfare gain is 22 percent larger. These differences grow modestly as we increase the value of γ. Since the welfare gain is linear in the monthly transitory variance, σ_θ^2, adjusting for differences in income volatility across racial groups would likely cause these gaps to widen even more.

In Online Appendix C.2 we complement this statistical model using a structural lifecycle model of consumption and saving. 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 and borrowing constraints that allow us to match the elasticities estimated in Section 3 and race-specific distributions of liquid assets, 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 50 percent larger for black households and 32 percent larger for Hispanic households than it is for white households.

5 Conclusion

In this paper we set out to estimate the consumption response to typical month-to-month changes in income. To overcome issues of endogeneity, we leverage firm-wide changes in monthly pay, instrumenting for a household's pay using the pay of all other co-workers.

---

[25]Hai, Krueger and Postlewaite (2013) and De Nardi, Fella and Paz-Pardo (2019) document a significant welfare gain from eliminating *all* income volatility (both temporary shocks and permanent shocks).

This type of income variation captures the most common sources of income variation for households, such as changes in hours, bonuses and commissions, and seasonal employment. This approach contrasts with standard identification strategies using unusual windfalls. We estimate an average consumption elasticity of 0.23.

Our relatively large data set allows us to precisely explore heterogeneity in the consumption responses across household race and liquid wealth on hand. We find sizable differences in consumption elasticities across racial groups. Black and Hispanic households respectively have elasticities that are 50 and 20 percent larger than those of white households. We likewise find a substantial role for liquid assets in driving heterogeneity in consumption responses. Households with the lowest asset levels exhibit sizable consumption elasticities, while those with the highest asset levels show little response to income variation. We then explore the interplay between these two patterns. We again compare consumption responses across racial groups, while controlling for the level of liquid assets on hand. We find that conditional on liquid assets, the differences in consumption elasticities across racial groups decrease by fifty to one hundred percent, depending on the specification.

To better interpret the differences in consumption smoothing that we find, we quantify the welfare implications of our results. We use a statistical model of consumption and income to calculate the welfare gain of fully insulating a household from transitory income volatility. The welfare gain from eliminating transitory shocks to income for black households are 46 to 56 percent higher than those of white households, while the gains for Hispanic households are 22 to 32 percent higher than those of white households. Our modeling exercises therefore mirror our empirical estimates: larger gaps in consumption elasticities translate into larger welfare differentials.

Our findings complement prior research in a number of ways. First, we introduce a novel method of identifying individual-level race in administrative data sets by linking administrative bank data to publicly available voter registration data. Second, ours is one of a few studies to leverage firm-wide changes in monthly pay as an instrument for income volatility, thus identifying consumption responses to typical sources of income variation. We are then able to identify month-to-month differences in consumption smoothing across racial groups. Likewise, our data and identification strategy allows us to estimate differences in consumption responses across liquid asset quantiles with a significantly high level of precision relative to previous studies. This precision allows us, for example, to identify a substantial role for liquid assets in statistically explaining differences in consumption elasticities across racial groups. Our results therefore highlight a potential focus area when considering policy approaches to mitigating the impact of income volatility on household welfare and to reducing inequality between racial groups.

References

Aaronson, Daniel, Daniel A. Hartley, and Bhashkar Mazumder. 2017. "The Effects of the 1930s HOLC "Redlining" Maps." Social Science Research Network SSRN Scholarly Paper ID 3038733, Rochester, NY.

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.

Aiyagari, S. Rao. 1994. "Uninsured Idiosyncratic Risk and Aggregate Saving." *The Quarterly Journal of Economics*, 109(3): 659–84.

Akee, Randall, Maggie R. Jones, and Sonya R. Porter. 2019. "Race Matters: Income Shares, Income Inequality, and Income Mobility for All U.S. Races." *Demography*, 56(3): 999–1021.

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.

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.

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., and Constantine Yannelis. 2017. "Income Changes and Consumption: Evidence from the 2013 Federal Government Shutdown." *Review of Economic Dynamics*, 23: 99–124.

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. 2019. "Consumer-Lending Discrimination in the FinTech Era." National Bureau of Economic Research Working Paper 25943. Series: Working Paper Series.

Baugh, Brian, Itzhak Ben-David, Hoonsuk Park, and Jonathan A Parker. 2018. "Asymmetric Consumption Smoothing." National Bureau of Economic Research Working Paper 25086. Series: Working Paper Series.

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. 2019. "Do Minorities Pay More for Mortgages?" Social Science Research Network SSRN Scholarly Paper ID 3352876, Rochester, NY.

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.

Campbell, John Y., and João F. Cocco. 2015. "A Model of Mortgage Default." *The Journal of Finance*, 70(4): 1495–1554.

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

Carroll, Chris, Nathan Palmer, Matthew N. White, Jacqueline Kazil, and David Low. 2018. "econ-ark/HARK: 0.8.0 (Version pre)."

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

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."

Cesarini, David, Erik Lindqvist, Robert Östling, and Björn Wallace. 2016. "Wealth, Health, and Child Development: Evidence from Administrative Data on Swedish Lottery Players." *The Quarterly Journal of Economics*, 131(2): 687–738. Publisher: Oxford Academic.

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.

Chen, M. Keith, Kareem Haggag, Devin G Pope, and Ryne Rohla. 2019. "Racial Disparities in Voting Wait Times: Evidence from Smartphone Data." National Bureau of Economic Research Working Paper 26487. Series: Working Paper Series.

Chetty, Raj, Nathaniel Hendren, Maggie R Jones, and Sonya R Porter. 2018. "Race and Economic Opportunity in the United States: An Intergenerational Perspective." National Bureau of Economic Research Working Paper 24441.

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.

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

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

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.*

Dobbie, Will, Andres Liberman, Daniel Paravisini, and Vikram Pathania. 2018. "Measuring Bias in Consumer Lending." National Bureau of Economic Research Working Paper 24953. Series: Working Paper Series.

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.

Du Bois, W. E. B. 1868-1963. 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.

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, Martin B. Holm, and Gisle J. Natvik. 2018. "MPC Heterogeneity and Household Balance Sheets." CESifo Group Munich 7134.

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, David Johnson, Timothy M. Smeeding, and Jeffrey P. Thompson. 2019. "Estimating the Marginal Propensity to Consume Using the Distributions of Income, Consumption and Wealth." Social Science Research Network SSRN Scholarly Paper ID 3374467, Rochester, NY.

Friedrich, Benjamin, Lisa Laun, Costas Meghir, and Luigi Pistaferri. 2019. "Earnings Dynamics and Firm-Level Shocks." National Bureau of Economic Research Working Paper 25786.

Fryer, Roland, Devah Pager, and Jörg L. Spenkuch. 2013. "Racial Disparities in Job Finding and Offered Wages." *Journal of Law and Economics*, 56(3): 633–689. Publisher: University of Chicago Press.

Fuster, Andreas, Paul Goldsmith-Pinkham, T. Ramadorai, and Ansgar Walther. 2018. "Predictably Unequal? The Effects of Machine Learning on Credit Markets." Social Science Research Network SSRN Scholarly Paper ID 3072038, Rochester, NY.

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

Gelman, Michael, Shachar Kariv, Matthew D. Shapiro, Dan Silverman, and Steven Tadelis. 2014. "Harnessing naturally occurring data to measure the response of spending to income." *Science*, 345(6193): 212–215.

Gelman, Michael, Shachar Kariv, Matthew D. Shapiro, Dan Silverman, and Steven Tadelis. 2018. "How individuals respond to a liquidity shock: Evidence from the 2013 government shutdown." *Journal of Public Economics*, 103917.

Grigsby, John, Erik Hurst, and Ahu Yildirmaz. 2019. "Aggregate Nominal Wage Adjustments: New Evidence from Administrative Payroll Data." National Bureau of Economic Research Working Paper 25628.

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*.

Guiso, Luigi, Luigi Pistaferri, and Fabiano Schivardi. 2005. "Insurance within the Firm." *Journal of Political Economy*, 113(5): 1054–1087.

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.

Guvenen, Fatih, Fatih Karahan, Serdar Ozkan, and Jae Song. 2015. "What Do Data on Millions of U.S. Workers Reveal about Life-Cycle Earnings Risk?" National Bureau of Economic Research Working Paper 20913. Series: Working Paper Series.

Hai, Rong, Dirk Krueger, and Andrew Postlewaite. 2013. "On the Welfare Cost of Consumption Fluctuations in the Presence of Memorable Goods." National Bureau of Economic Research Working Paper 19386. Series: Working Paper Series.

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.

Havranek, Tomas, and Anna Sokolova. 2020. "Do consumers really follow a rule of thumb? Three thousand estimates from 144 studies say "probably not"." *Review of Economic Dynamics*, 35: 97–122.

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.

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.

Juhn, Chinhui, Kristin McCue, Holly Monti, and Brooks Pierce. 2018*a*. "Firm-level Shocks and the Volatility of Worker Earnings within Megafirms." Working Paper.

Juhn, Chinhui, Kristin McCue, Holly Monti, and Brooks Pierce. 2018*b*. "Firm Performance and the Volatility of Worker Earnings." *Journal of Labor Economics*, 36(S1): S99–S131.

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.

Kempf, Elisabeth, and Margarita Tsoutsoura. 2018. "Partisan Professionals: Evidence from Credit Rating Analysts." National Bureau of Economic Research Working Paper 25292. Series: Working Paper Series.

Keys, Benjamin J, Neale Mahoney, and Hanbin Yang. 2020. "What Determines Consumer Financial Distress? Place- and Person-Based Factors." National Bureau of Economic Research Working Paper 26808. Series: Working Paper Series.

Kijakazi, Kilolo, Karen E. Smith, and Charmaine Runes. 2019. "African American Economic Security and the Role of Social Security." Library Catalog: www.urban.org.
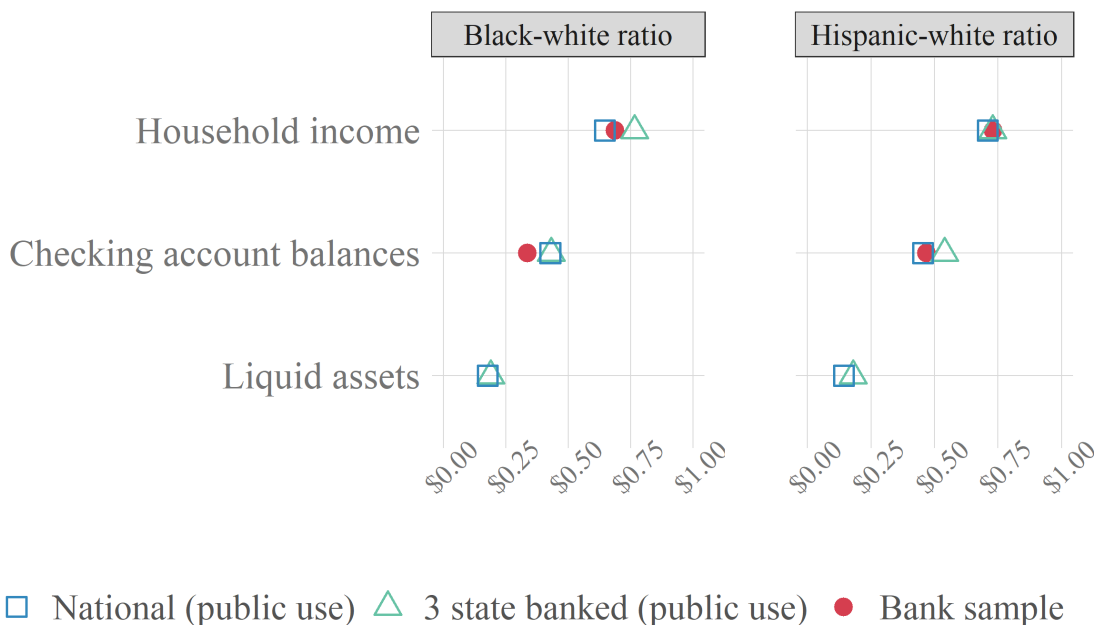
Kline, Patrick, Neviana Petkova, Heidi Williams, and Owen Zidar. 2019. "Who Profits from Patents? Rent-Sharing at Innovative Firms." *The Quarterly Journal of Economics*, 134(3): 1343–1404.

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, Raffaele D Saggio, and Stephen A Woodbury. 2020. "Do Firm Effects Drift? Evidence from Washington Administrative Data." National Bureau of Economic Research Working Paper 26653.

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.

Lamadon, Thibaut, Magne Mogstad, and Bradley Setzler. 2019. "Imperfect Competition, Compensating Differentials and Rent Sharing in the U.S. Labor Market." National Bureau of Economic Research Working Paper 25954.

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

Lian, Chen. 2019. "Consumption with Imperfect Perception of Wealth." *working paper*.

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.

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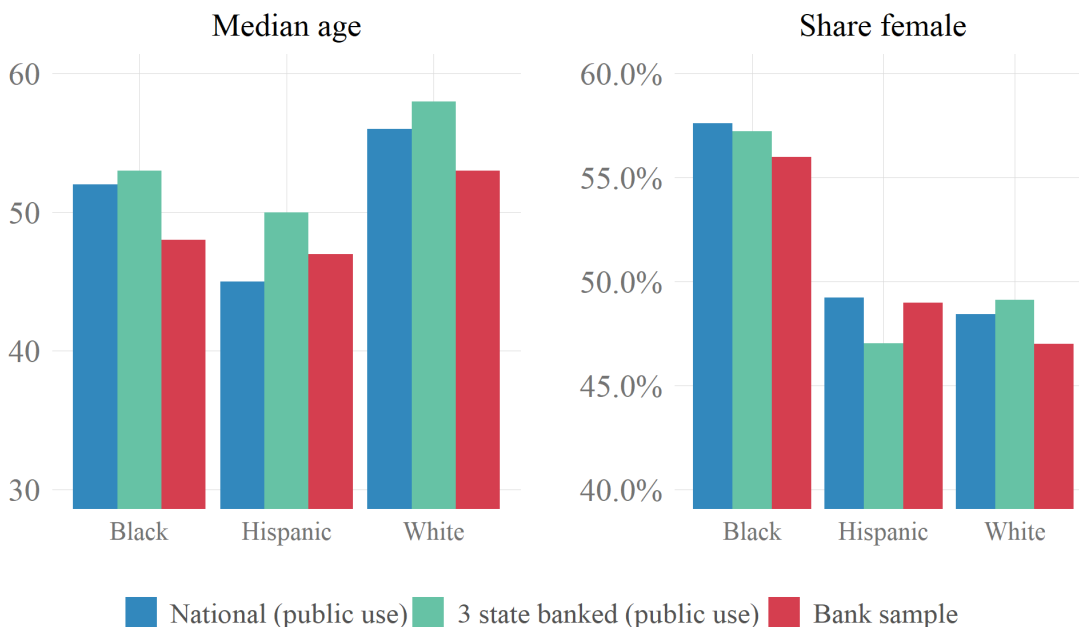
Squires, Munir. 2016. "Kinship taxation as a constraint on microenterprise growth." PhD diss. The London School of Economics and Political Science (LSE).

Stein, Luke C. D., and Constantine Yannelis. 2019. "Financial Inclusion, Human Capital, and Wealth Accumulation: Evidence from the Freedman's Savings Bank." Social Science Research Network SSRN Scholarly Paper ID 3302996, Rochester, NY.

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.

Thompson, Jeffrey P., and Gustavo Suarez. 2015. "Updating the Racial Wealth Gap." Social Science Research Network SSRN Scholarly Paper ID 3072923, Rochester, NY.

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

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.

Figure 1: Economic and Demographic Variables by Race

(a) Income and Asset Gaps

□ National (public use)   △ 3 state banked (public use)   ● Bank sample

(b) Demographics

■ National (public use)   ■ 3 state banked (public use)   ■ Bank sample

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 14 for details. Bank sample is the matched Chase-voter data described in the text.
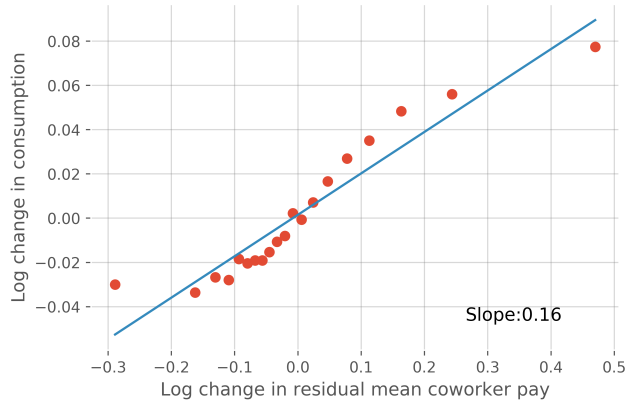
Figure 2: Consumption Sensitivity to Payroll Fluctuations

(a) First Stage

Slope:0.70

(b) Reduced Form


Slope:0.16

### (c) Second Stage

IV Elasticity Estimate = 0.23

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 of 20 vingtiles of the change in log 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 (5). See Section 3.3 for details.

40

Figure 3: Consumption Sensitivity by Race: Black and White
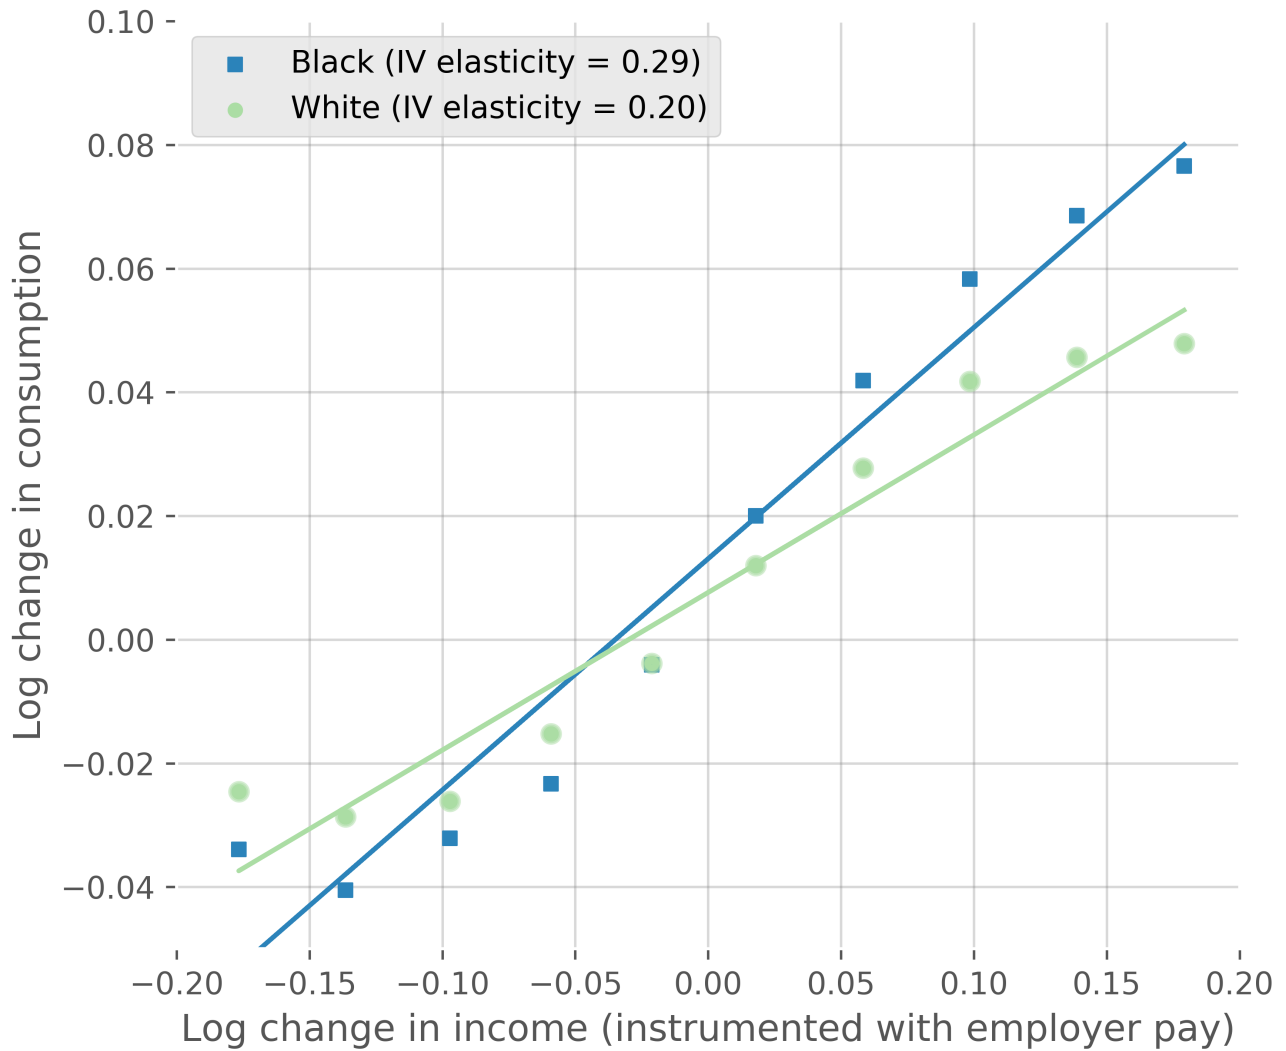
Note: This figure shows the effect of changes in labor income on nondurable consumption by race in the matched Chase-voter data. The x-axis stratifies the sample into 20 vingtiles of the predicted change in log own pay, where the omitted instrument is the change in pay of all other workers at the same firm. The y-axis shows conditional means of the change in log nondurable consumption for black and white households. This is the second stage plot from Figure 2c, shown separately by race. See Section 3.4 for details.

Figure 4: Consumption Sensitivity by Race: Hispanic and Non-Hispanic White
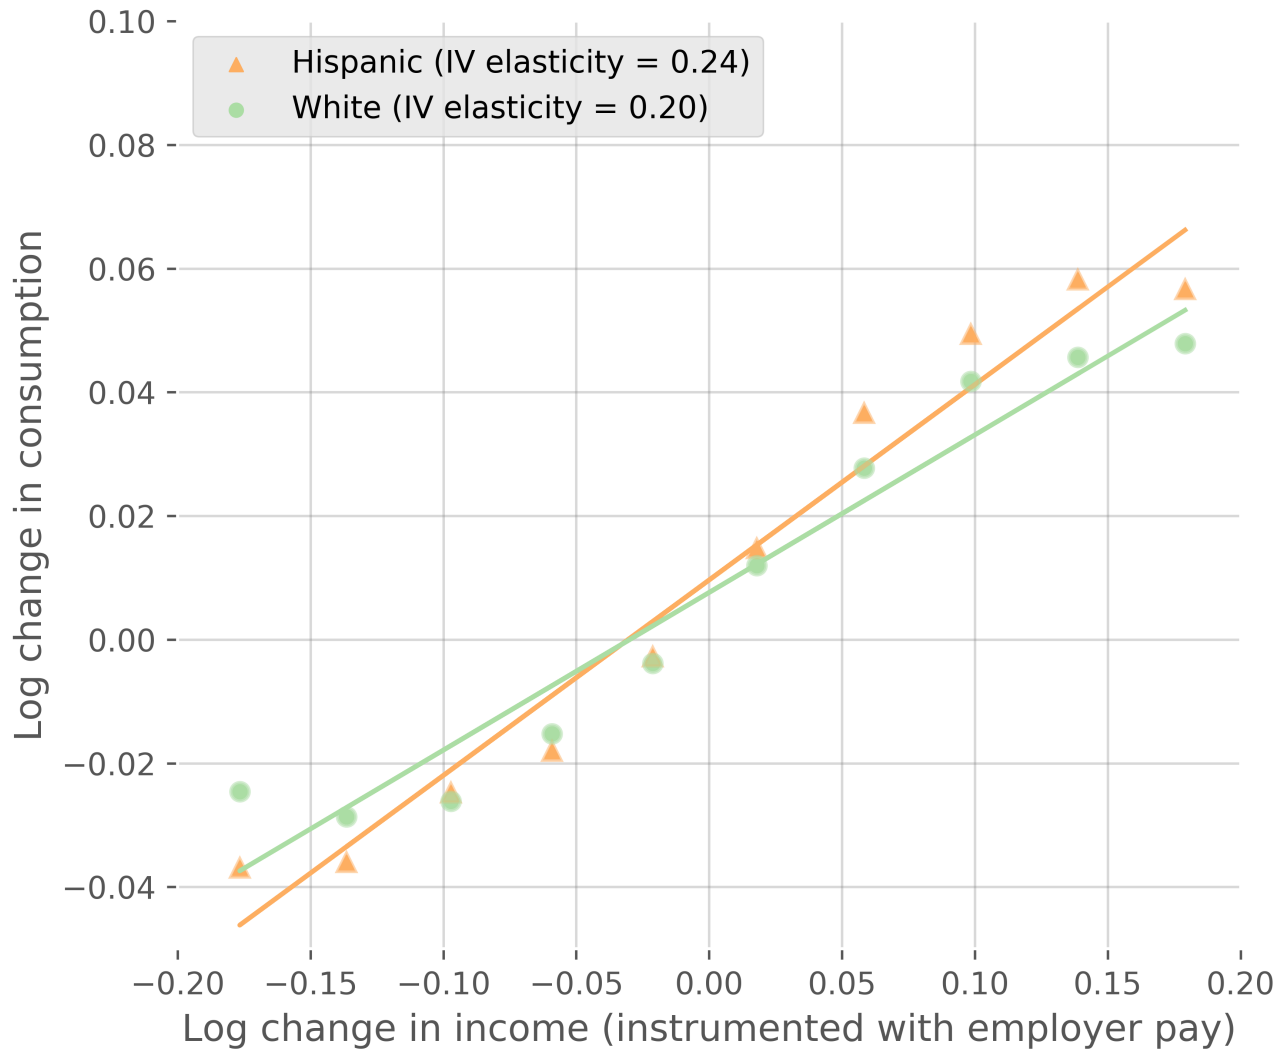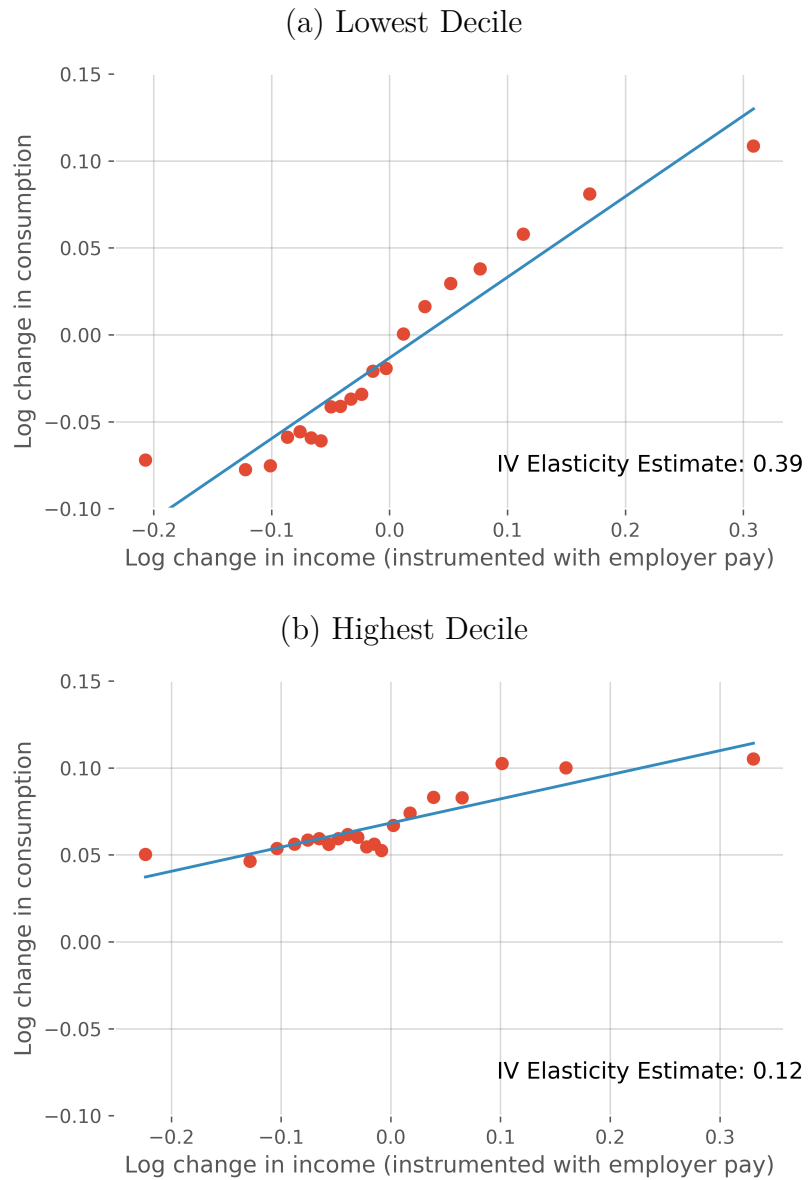

Note: This figure shows the effect of changes in labor income on nondurable consumption by race in the matched Chase-voter data. The x-axis stratifies the sample into 20 vingtiles of the predicted change in log own pay, where the omitted instrument is the change in pay of all other workers at the same firm. The y-axis shows conditional means of the change in log nondurable consumption for Hispanic and non-Hispanic white households. This is the second stage plot from Figure 2c, shown separately by Hispanicity. See Section 3.4 for details.

Figure 5: Consumption Sensitivity by Asset Buffer Decile

(a) Lowest Decile

IV Elasticity Estimate: 0.39

(b) Highest Decile


IV Elasticity Estimate: 0.12

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 20 vingtiles 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 2c, 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 5b shows increasing consumption on average, across the distribution of pay shocks. See Section 3.5 for details.

Figure 6: Marginal Propensity to Consume by Asset Buffer
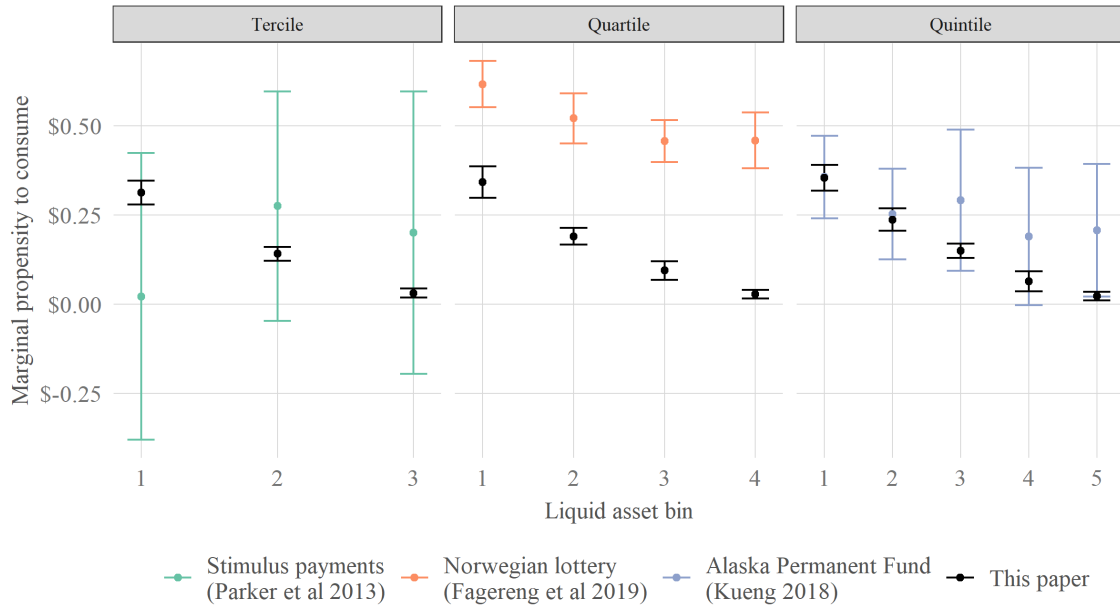
Note: This figure compares the estimates of heterogeneity by assets in the passthrough of income shocks to consumption. Parker et al. (2013), Fagereng, Holm and Natvik (2018) and Kueng (2018) use terciles, quartiles, and quintiles respectively. To enable comparability with these prior papers, we calculate 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 (2018) measures the MPC on total consumption. See Section 3.5 for details.

Figure 7: Racial Inequality in Consumption Smoothing and Role of Assets

(a) Firm Pay Shocks

(b) Unemployment
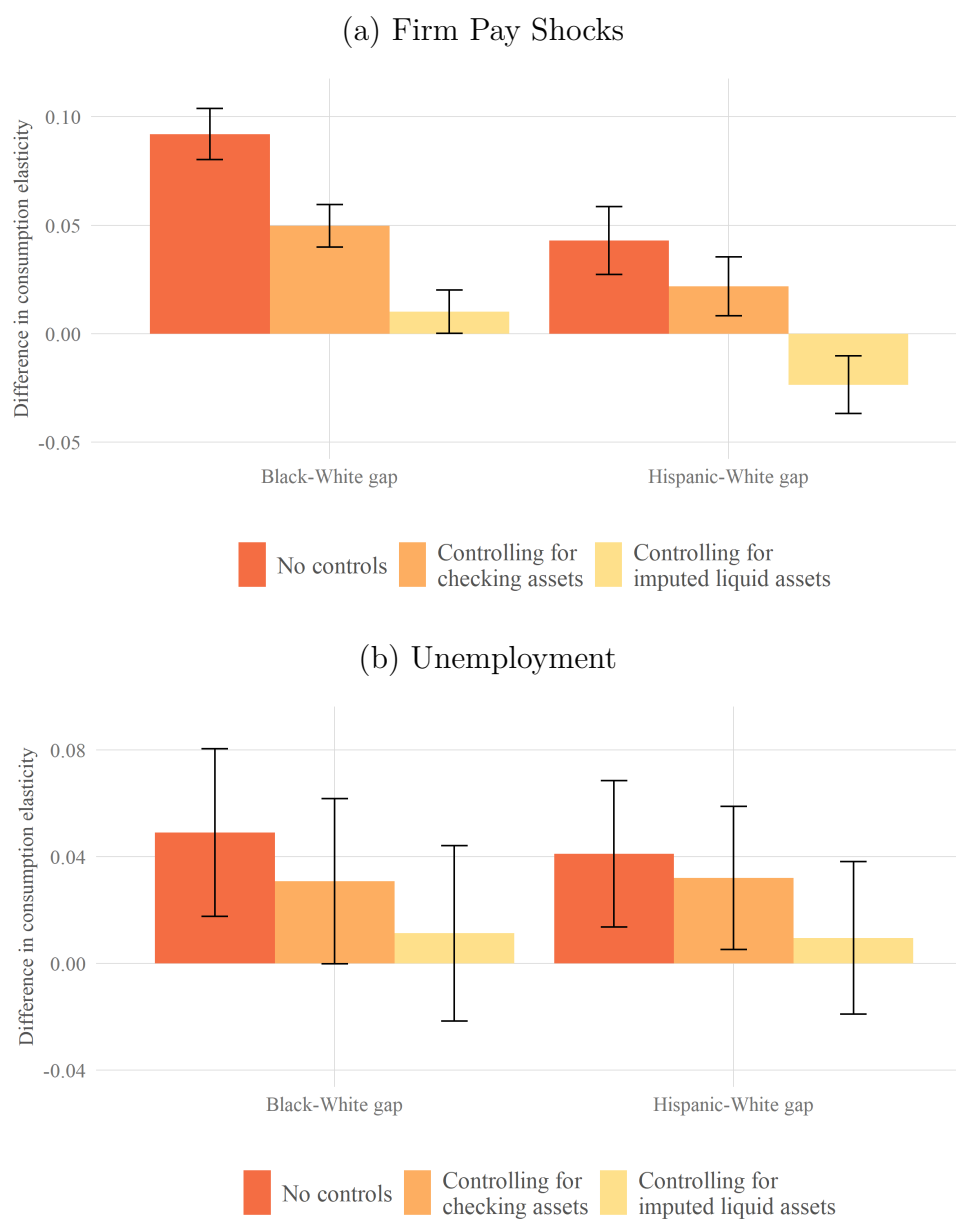
Note: This figure illustrates regression coefficients from Table 2 and 3. Within each race, the figure shows three bars. The left-most dark bar shows the difference in the elasticity of between black or Hispanic households and white households. The middle lighter bar shows the difference when controlling for observed checking account balance. The right-most light 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 3.6 for details on firm pay fluctuations and Section 3.7 for details on unemployment.

Figure 8: Effect of Unemployment on Consumption by Race

(a) Labor Income + Unemployment Insurance
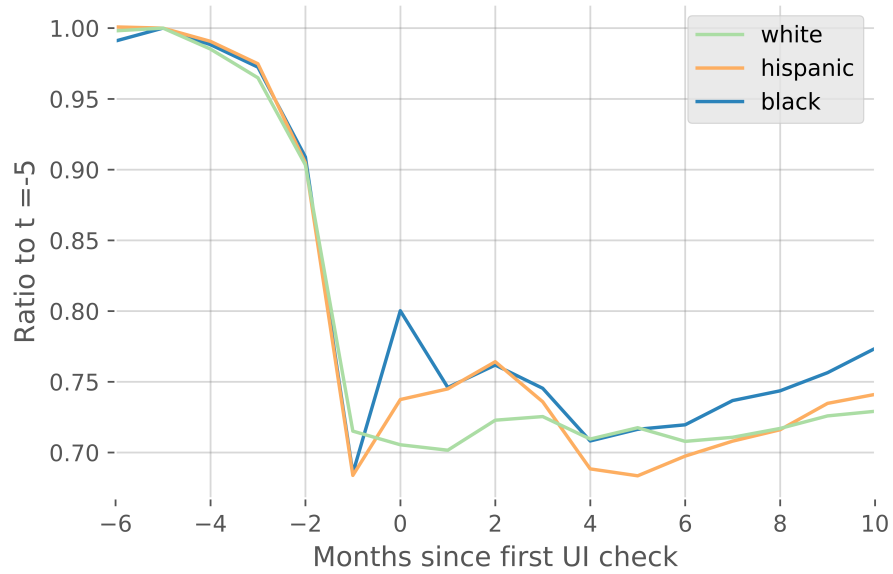
(b) Nondurable Spending
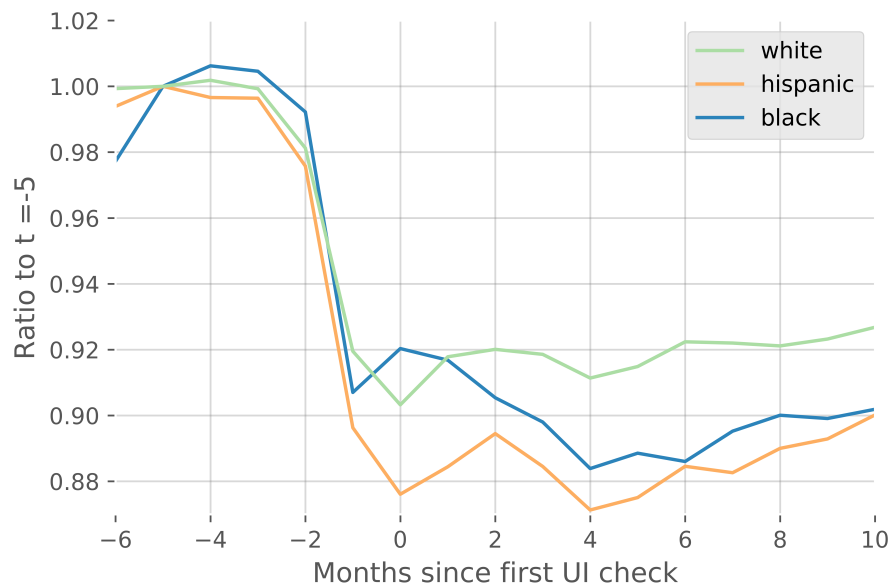
Note: This figure shows the evolution of income around unemployment in the matched Chase-voter data. Unemployment is measured using the household's first receipt of an unemployment insurance (UI) check. See Section 3.7 for details.

Table 1: Summary Statistics for Analysis Sample

Variable	Race	Mean	Median	Std. Dev	N Households	N Household-Months
Labor Income	White	3,531.2	2,700.6	8,884.1	351,208	10,266,939
	Black	2,242.0	1,864.3	2,489.5	183,601	4,949,641
	Hispanic	2,674.5	2,178.3	3,359.6	162,301	4,280,128
Coworker Pay	White	2,986.6	2,746.3	1,813.6		
	Black	2,700.0	2,544.4	1,456.9		
	Hispanic	2,828.3	2,665.1	1,804.0		
Non Durable Consumption	White	3,229.0	2,453.8	6,692.2		
	Black	2,377.0	1,909.4	2,394.9		
	Hispanic	2,534.9	1,981.2	3,012.6		
Checking Assets	White	7,276.1	2,271.9	25,671.6		
	Black	2,509.8	837.0	8,734.4		
	Hispanic	3,459.5	1,204.7	10,604.9		
Buffer Ratio	White	15.7	0.8	3,921.4		
	Black	4.1	0.4	1,057.4		
	Hispanic	9.4	0.6	2,444.1		

Note: This table shows summary statistics for the analysis sample in our coworker pay design. See Section 3 for details. This sample, which only includes observations with at least ten months of income and consumption history, is a bit over one-third of the size of the entire sample described in Section 2. Buffer ratio is the ratio of checking assets to nondurable consumption.

Table 2: 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	0.120	0.091	0.231	0.200	0.233	0.217	0.234	0.237
	(0.003)	(0.002)	(0.009)	(0.009)	(0.009)	(0.008)	(0.009)	(0.008)
(Δ Log Income) x Black		0.080		0.092		0.050		0.010
		(0.004)		(0.006)		(0.005)		(0.005)
(Δ Log Income) x Hispanic		0.040		0.043		0.022		-0.024
		(0.003)		(0.008)		(0.007)		(0.007)
(Δ Log Income) x Checking					-0.252	-0.233		
					(0.009)	(0.008)		
(Δ Log Income) x Liquid(Imputed)							-0.255	-0.258
							(0.010)	(0.009)
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	19,278,989	19,278,989	19,200,198	19,200,198	19,200,198	19,200,198	19,200,198	19,200,198
Adjusted R^2	0.009	0.009	0.001	0.002	0.005	0.005	0.005	0.005

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 (6) and (5). 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 3.6 for details.

Table 3: 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 (8) and (9). 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. See Section 3.7 for details.

Table 4: Welfare Gain of Eliminating Transitory Shocks

	Consumption elasticity of transitory shocks: β	Monthly transitory variance: σ_θ^2	Welfare gain: $\lambda \cong \left[\beta + \beta^2(\gamma-1)\right]\frac{\sigma_\theta^2}{2}$			
			Coefficient of relative risk aversion			
			$\gamma = 1$	$\gamma = 2$	$\gamma = 3$	$\gamma = 4$
Black	0.292	0.16	2.27%	2.93%	3.59%	4.25%
Hispanic	0.243	0.16	1.89%	2.34%	2.80%	3.26%
White	0.200	0.16	1.55%	1.86%	2.17%	2.48%

Note: This table conducts an exercise in the spirit of Lucas (1987) by calculating the welfare gain from eliminating transitory income shocks. The gain is expressed as a fraction of lifetime consumption. The transitory elasticity β is from column (4) of Table 2. The variance of monthly shocks is from Farrell, Greig and Yu (2019). See Section 4 for details.

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

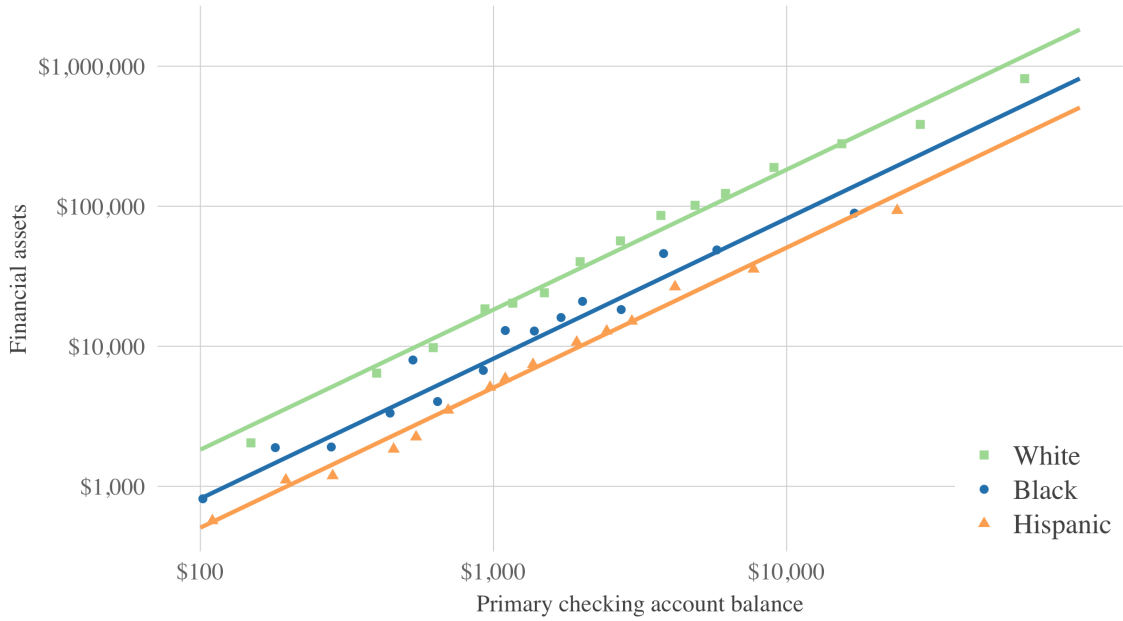Peter Ganong, Damon Jones, Pascal Noel, Diana Farrell, Fiona Greig, and Chris Wheat

Contents

A   Appendix Figures and Tables
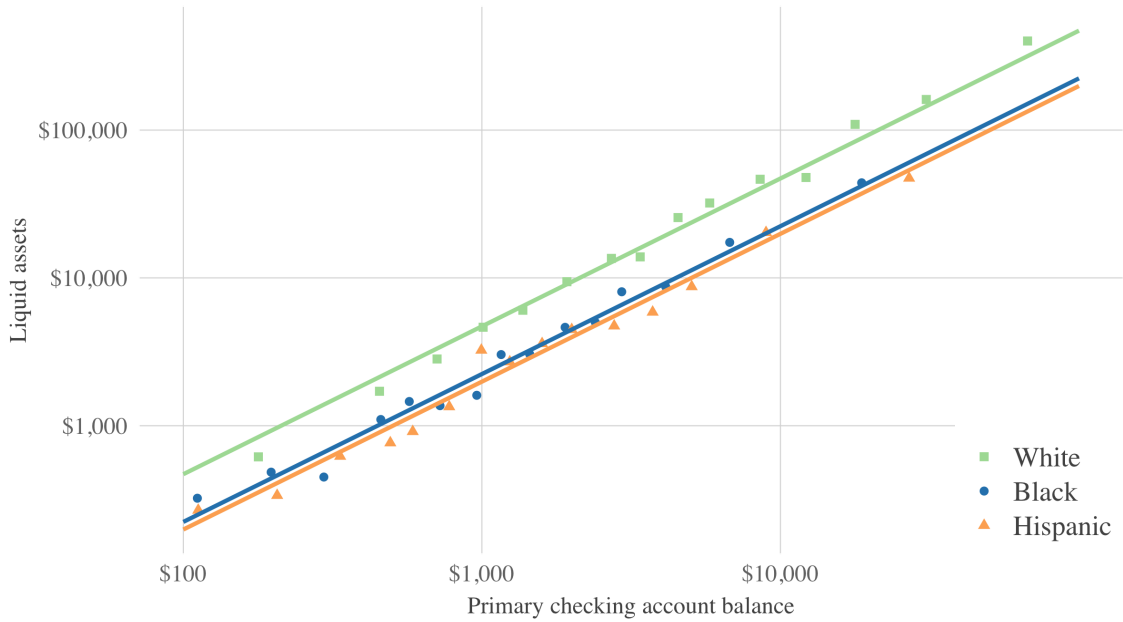
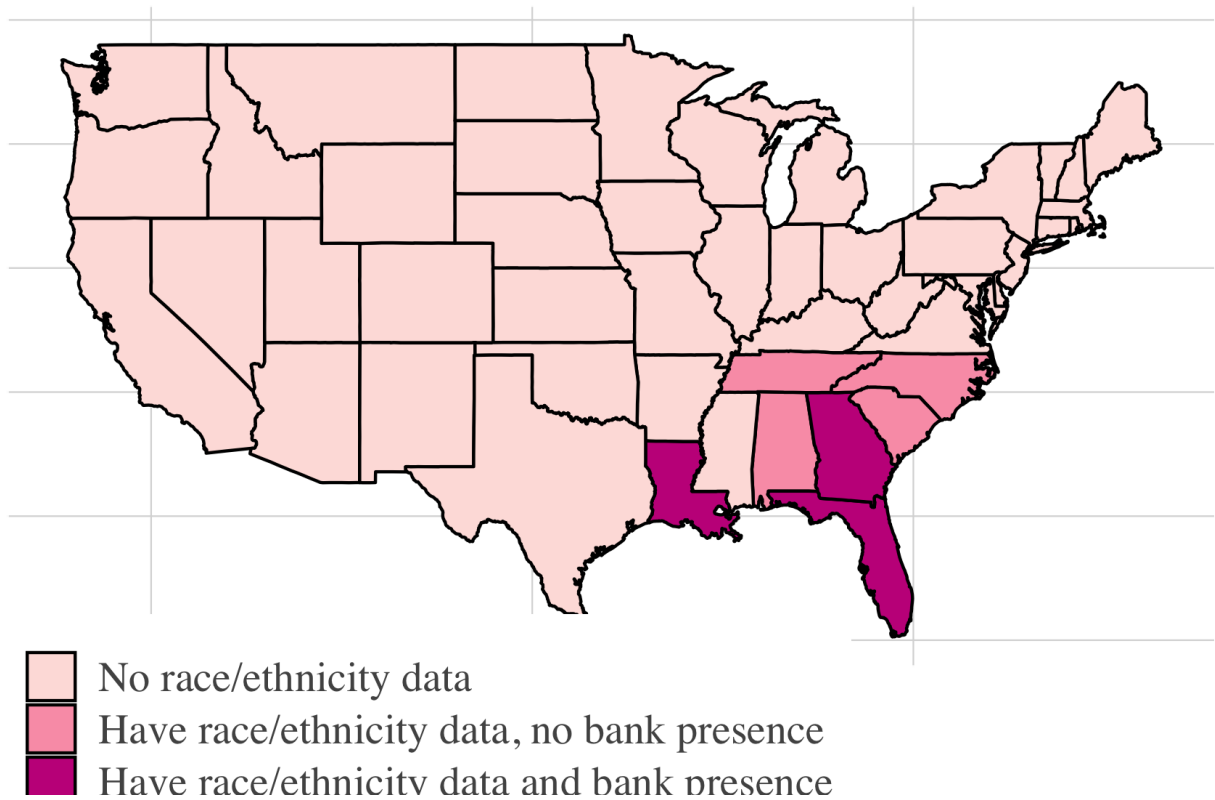Figure 1: Assets and Checking Account Balances

(a) Financial Assets

(b) Liquid Assets

Notes: 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

1

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

☐ No race/ethnicity data
☐ Have race/ethnicity data, no bank presence
☐ Have race/ethnicity data and 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 3: Share of Voters Matched
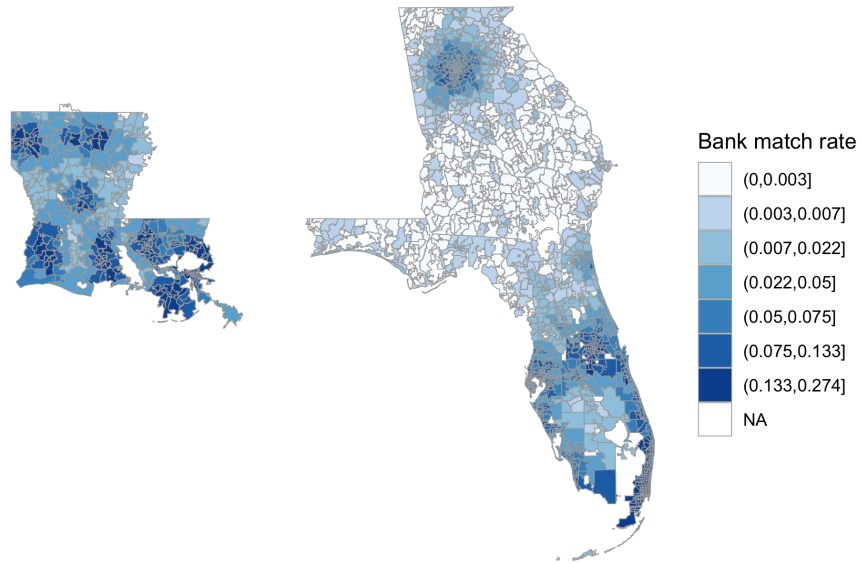
Notes: This figure shows the share of registered voters that have been matched to Chase bank accounts.

Figure 4: Distribution by Race – National versus Sample Frame versus Bank Sample

(a) Income Distribution

National income quintile, by race

National (public use)   3 state banked (public use)   Bank sample

(b) Age Distribution

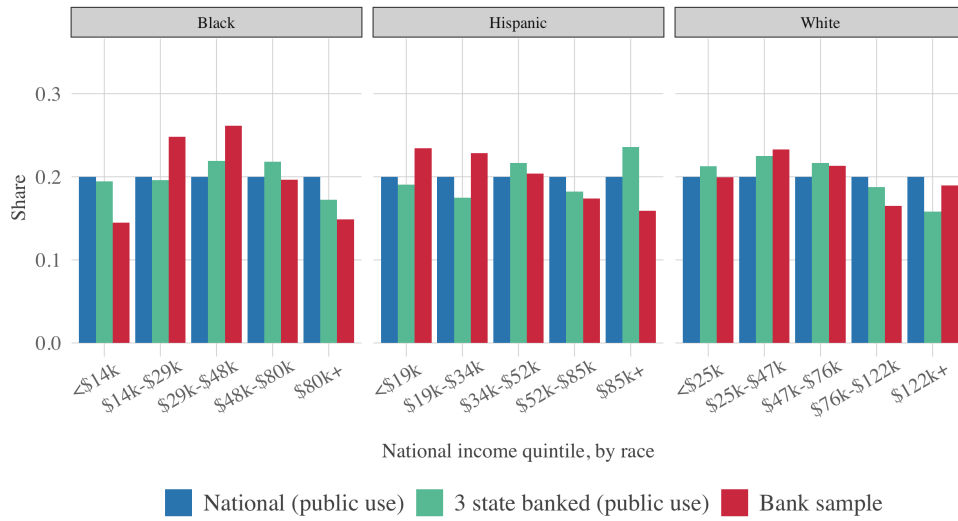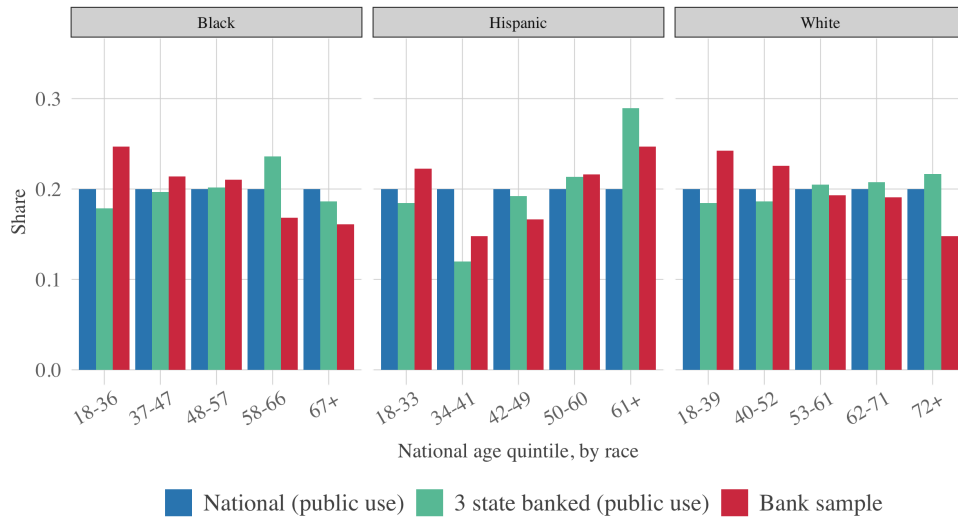National age quintile, by race

National (public use)   3 state banked (public use)   Bank sample

Notes: This figure shows the distribution of income and age in the US 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. Public Use 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.

4

Figure 5: Local Race and Hispanicity Distributions

(a) Share Black
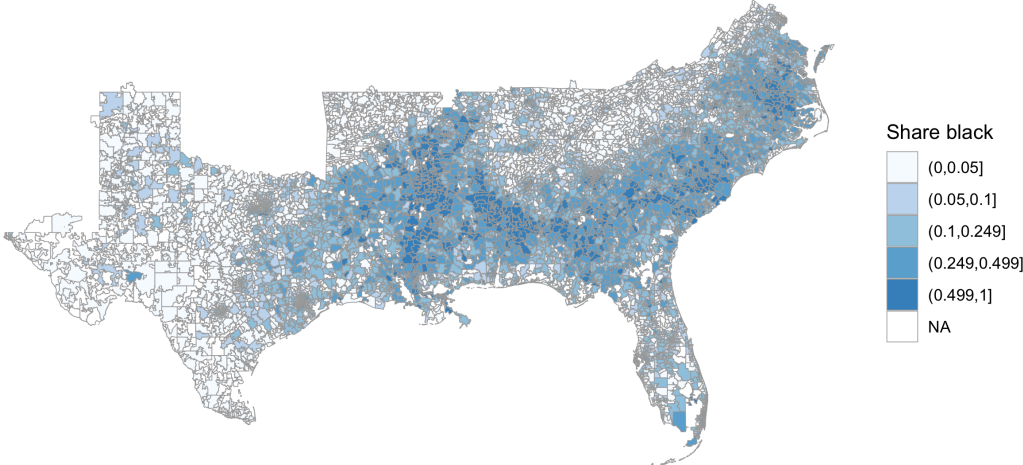
(b) Share Hispanic
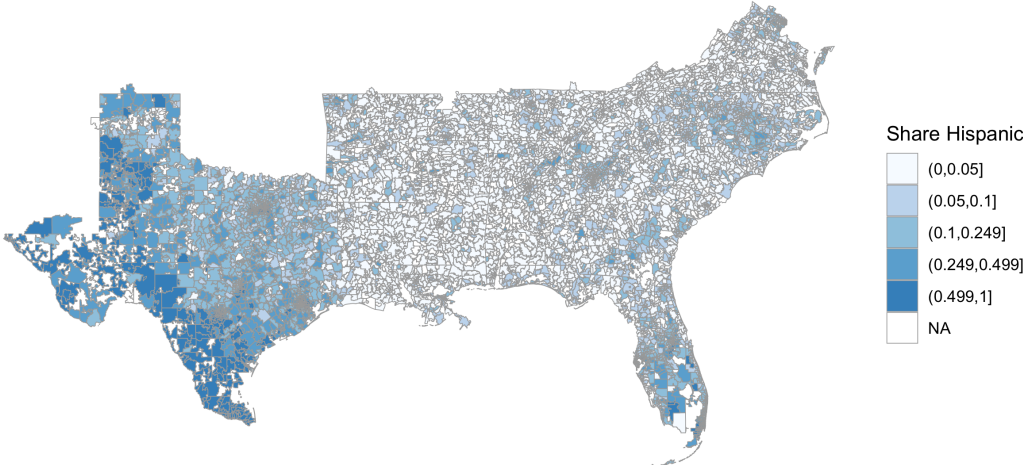
Source: 2010 US Census

Figure 6: Hispanic origin in 3 state sample frame compared to US
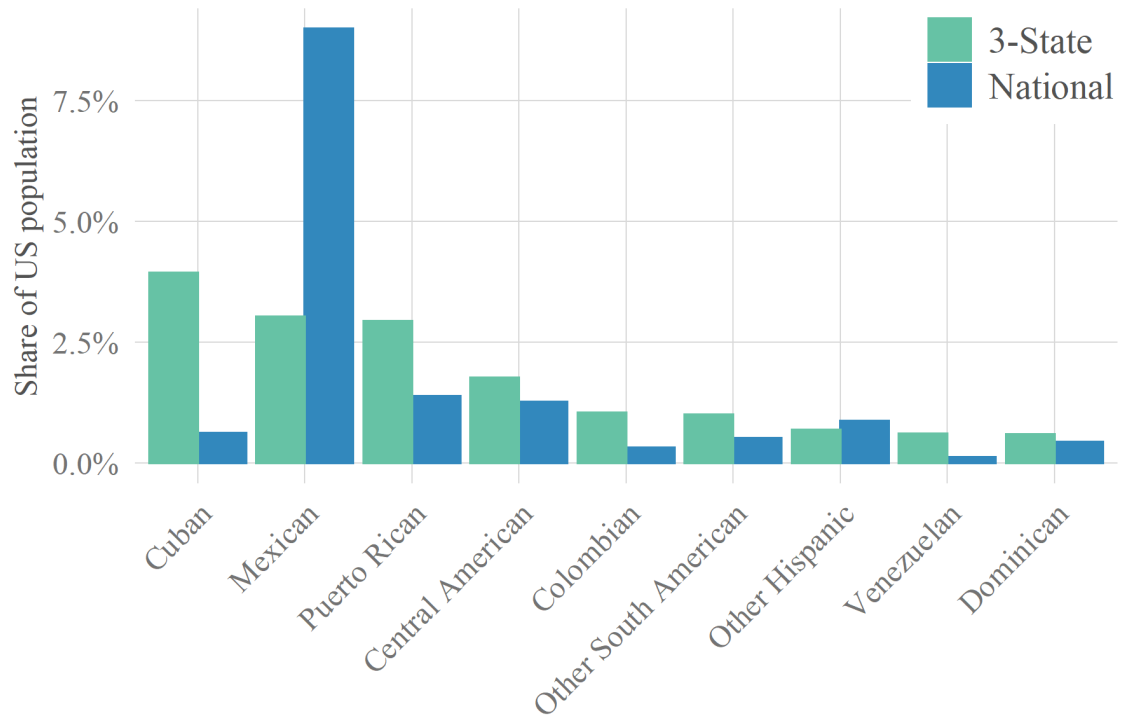
Source: 2018 American Community Survey 1-year estimates, Table B03001.

Figure 7: Financial Vulnerability – National and in Sample Frame

(a) Share Ever Behind on Bills in Last Year

■ National (public use)　■ 3 state banked (public use)
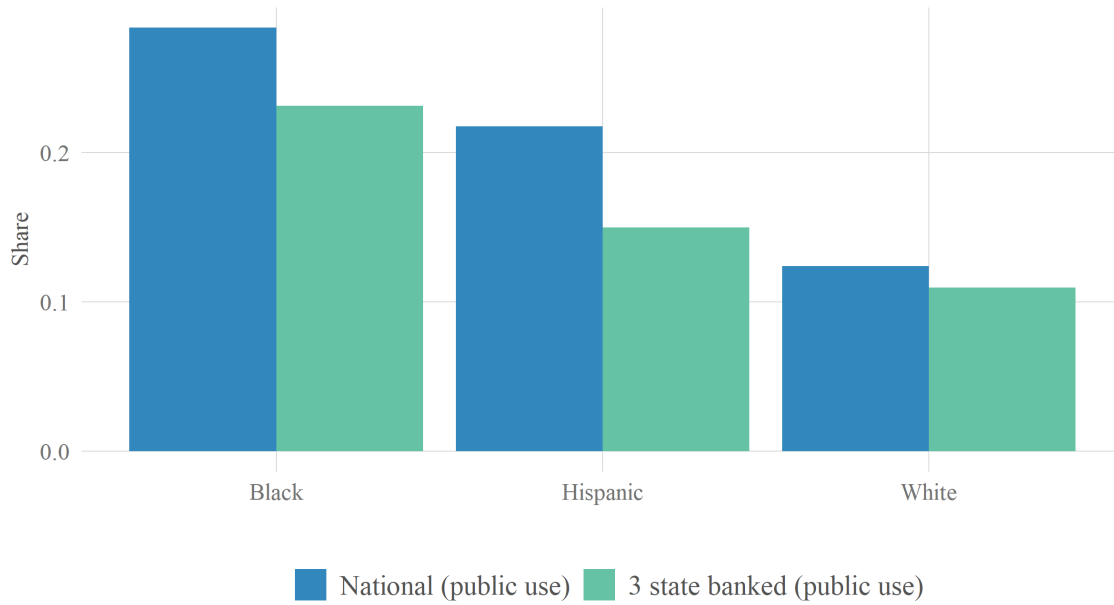
(b) Share with no Emergency Savings

■ National (public use)　■ 3 state banked (public use)
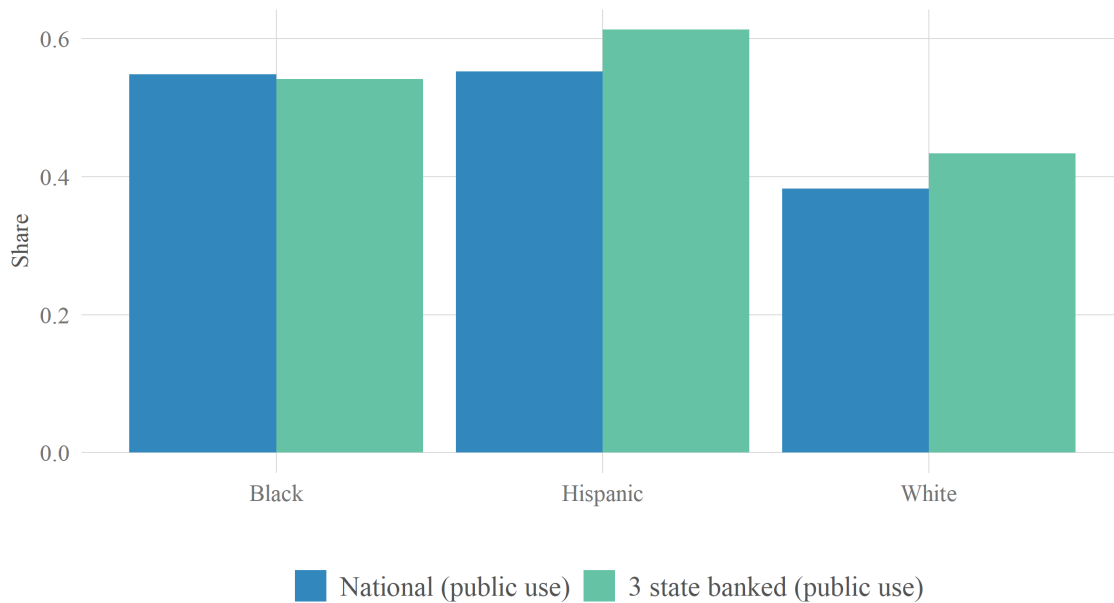
Notes: The questions are "Often times, households find that they are not able to keep up with their bills. Over the last 12 months, was there a time when (you/ or someone else in your household) fell behind on bill payments?" and "Did (you/ or anyone in your household) set aside any money in the past 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 8: Consumption Sensitivity by Race (OLS)

(a) All

Slope:0.12

(b) White

Slope:0.09

(c) Black

Slope:0.17

(d) Hispanic

Slope:0.13
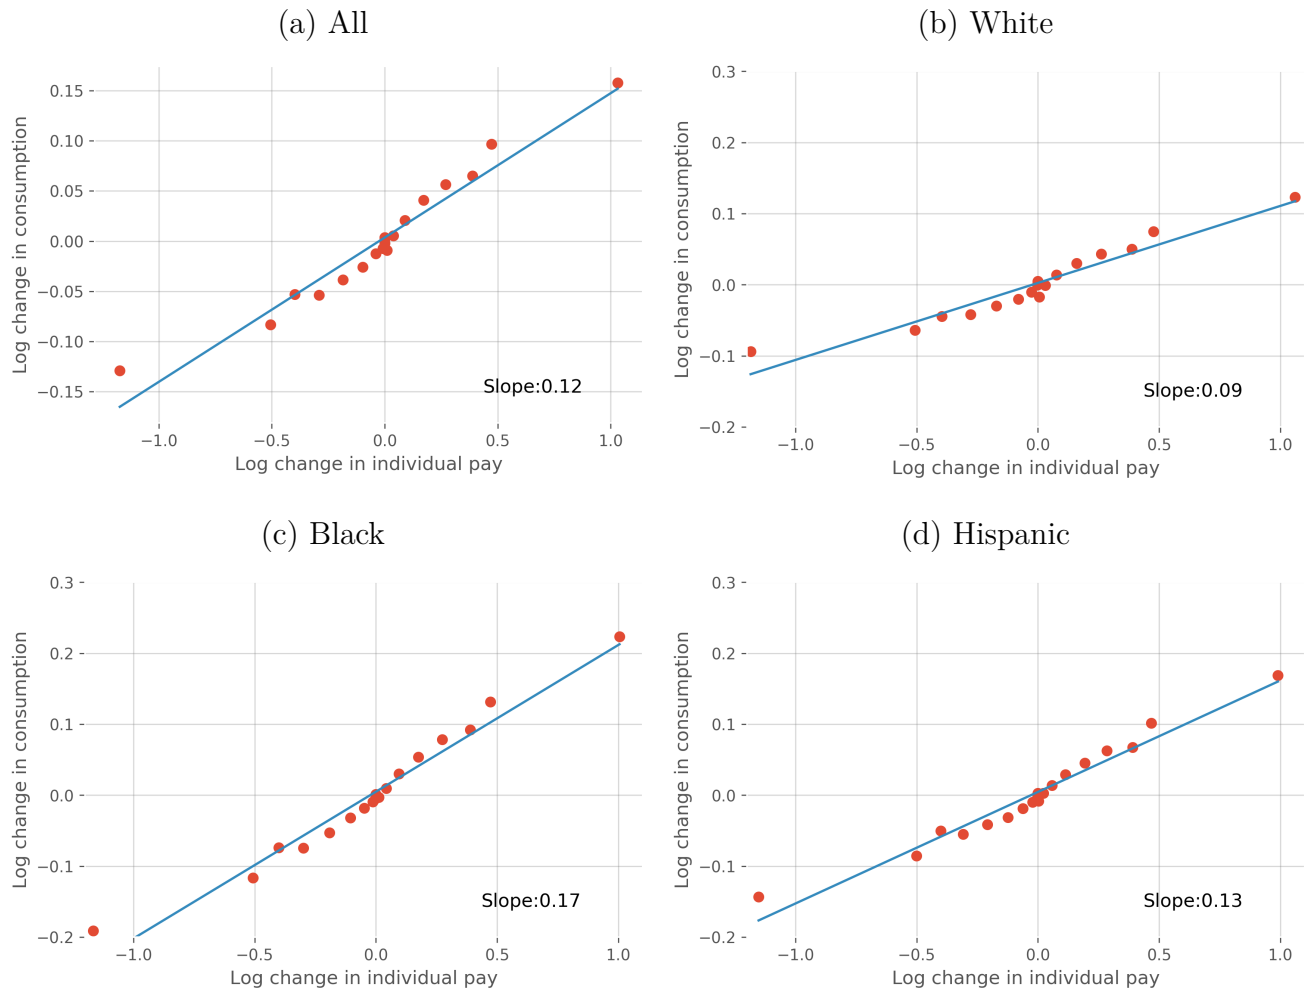
Note: This figure shows Ordinary Least Squares (OLS) estimates of the effect of changes in income on changes in consumption (equations (1) and (2)). The x-axis shows vingtiles of the change in individual pay. The y-axis shows conditional means of the change in consumption. The slope of the line is the elasticity of consumption with respect to income. Panel (a) shows the pooled sample, panel (b) shows white households, panel (c) shows black households, and panel (d) shows Hispanic households. See Section 3.1 for details.

Figure 9: Distribution of Firm Pay Shocks
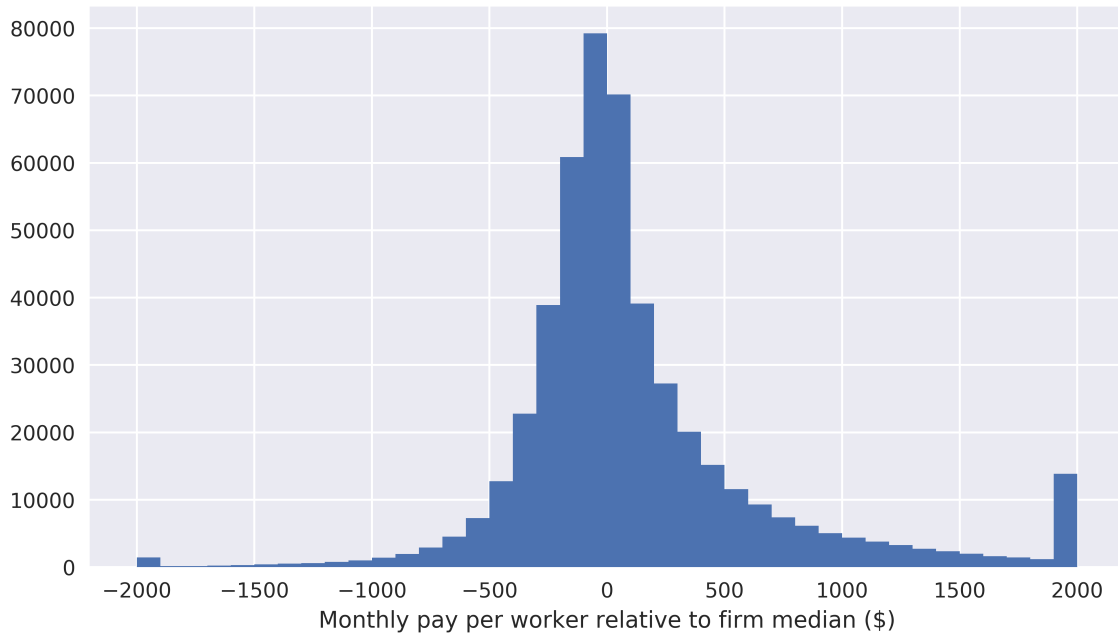
Note: This figure shows a histogram of deviations of firm-wide pay per worker in a given month relative to the firm's median pay per worker across all months. To increase precision, the sample only includes firms with at least 100 employees in every month. Differences are winsorized at +/- $2,000. The average positive deviation is $396 and the average negative deviation is -$227.

Figure 10: Pre-trends in Firm Pay Shock
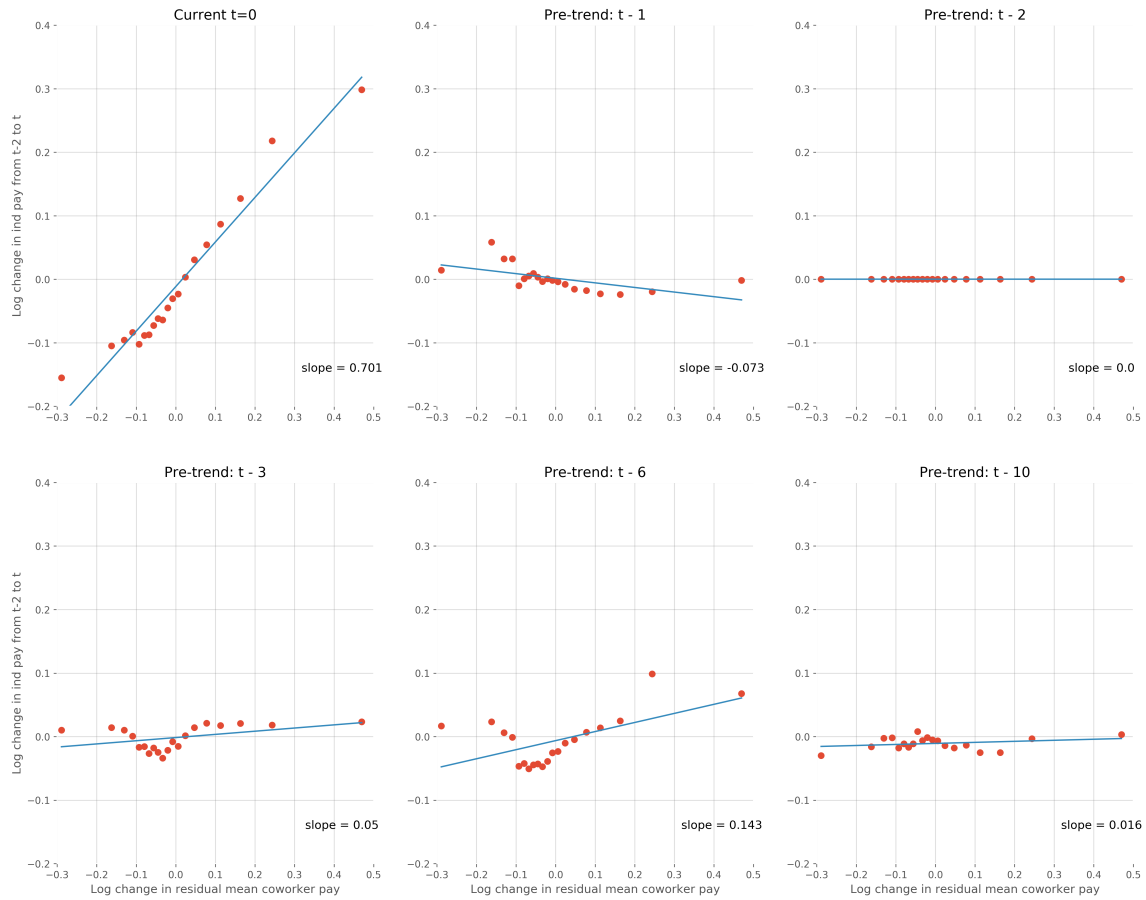
Note: This figure shows the evolution of labor income prior to a firm pay shock. Across all panels, the x-axis plots vingtiles of the change in residualized coworker pay from $t-1$ to t. The y-axis plots conditional means of the change in pay from $t-2$ to t in the top-left panel, $t-2$ to $t-1$ in the top-middle panel, $t-2$ to $t-2$ (which is mechanically equal to zero) in the top-right panel, $t-2$ to $t-3$ in the bottom-left panel, $t-2$ to $t-6$ in the bottom-middle panel, and $t-2$ to $t-6$ in the bottom-right panel.

Figure 11: Persistence of Firm Pay Shock
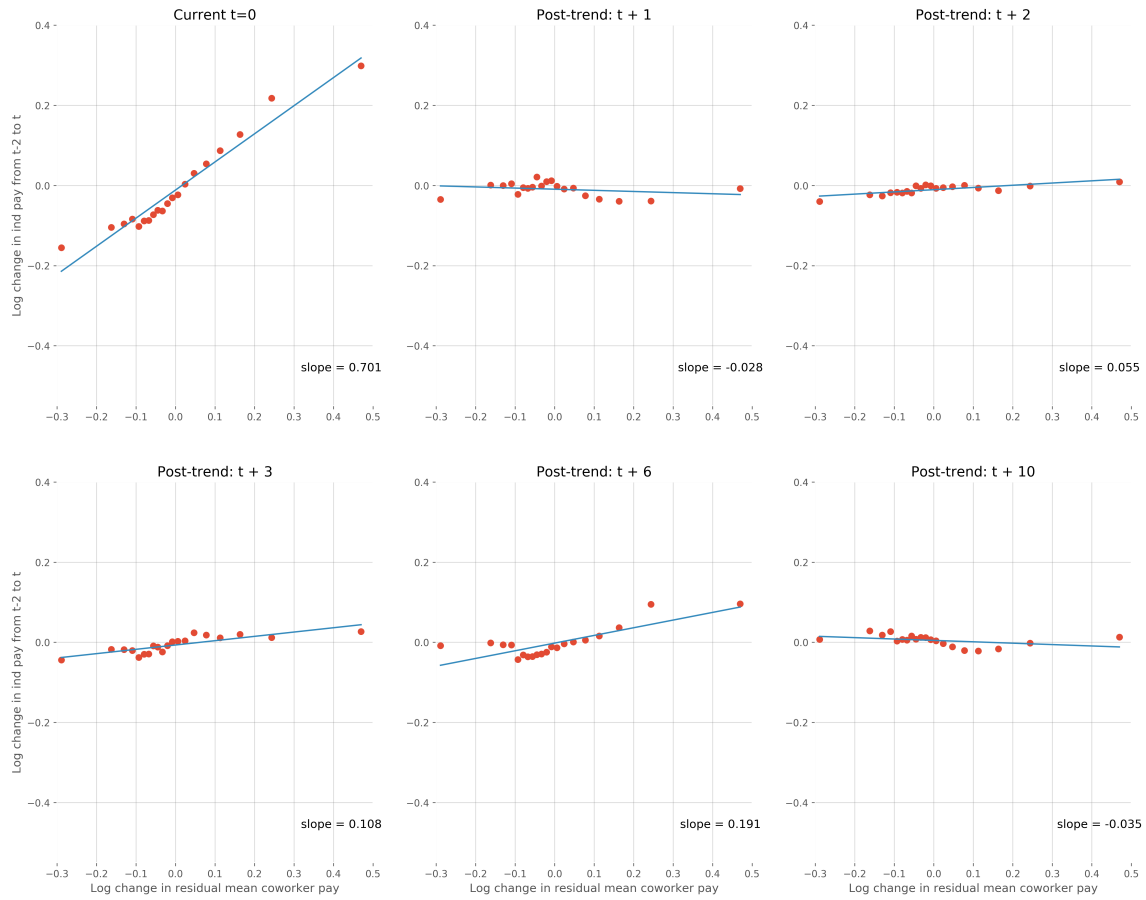
Note: This figure shows the evolution of labor income after a firm pay shock. Across all panels, the x-axis plots vingtiles of the change in residualized coworker pay from $t-1$ to t. The y-axis plots conditional means of the change in pay from $t-2$ to t in the top-left panel, $t-2$ to $t+1$ in the top-middle panel, $t-2$ to $t+2$ in the top-right panel, $t-2$ to $t+3$ in the bottom-left panel, $t-2$ to $t+6$ in the bottom-middle panel, and $t-2$ to $t+10$ in the bottom-right panel.

Figure 12: Standard Error of MPC Estimates in Prior Literature

Notes: Havranek and Sokolova (2020) review 246 estimates of the MPC for nondurables. This figure plots the distribution of standard errors for MPC estimates in their literature review. The vertical bar at one cent reflects the standard error of our preferred MPC for nondurables from Section 3.3.

Figure 13: Pre-trends in Consumption Prior to Firm Pay Shock
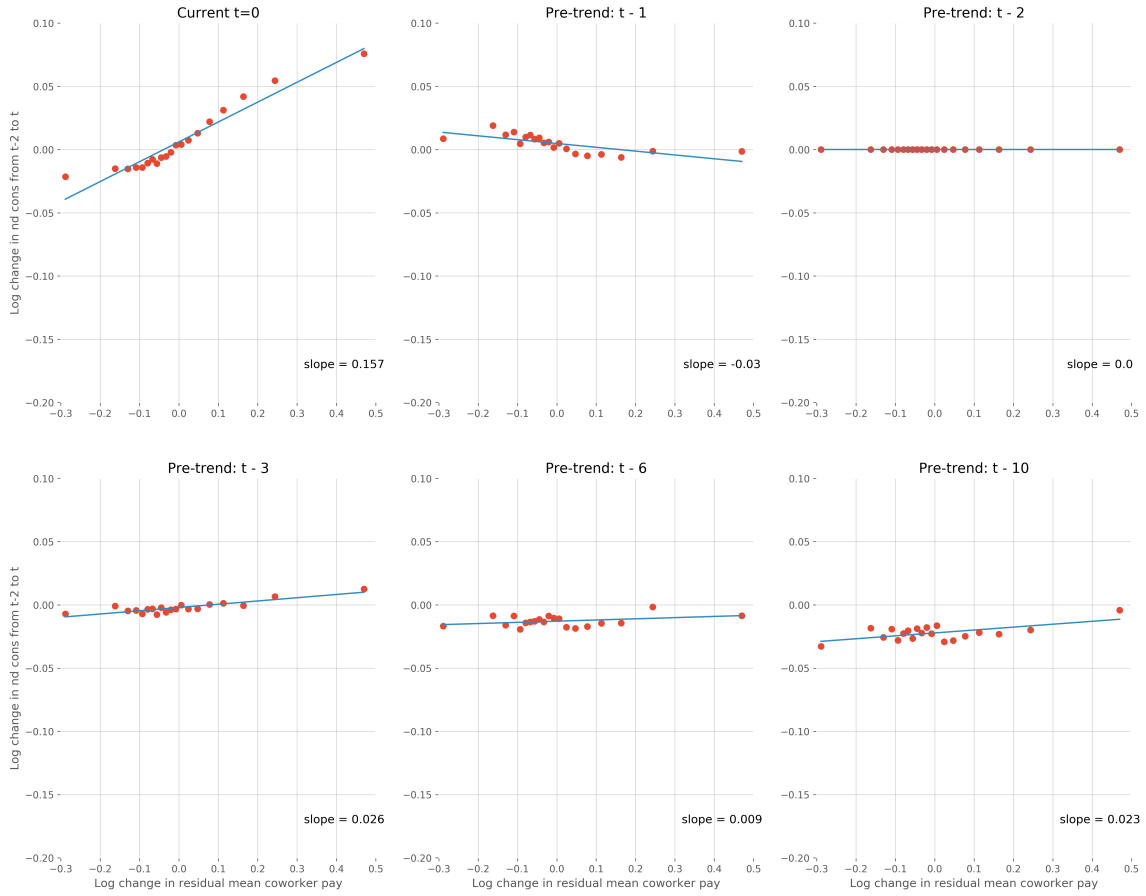
Note: This figure shows the evolution of consumption prior to a firm pay shock. Across all panels, the x-axis plots vingtiles of the change in residualized coworker pay from $t-1$ to t. The y-axis plots conditional means of the change in consumption from $t-2$ to t in the top-left panel, $t-2$ to $t-1$ in the top-middle panel, $t-2$ to $t-2$ (which is mechanically equal to zero) in the top-right panel, $t-2$ to $t-3$ in the bottom-left panel, $t-2$ to $t-6$ in the bottom-middle panel, and $t-2$ to $t-6$ in the bottom-right panel.

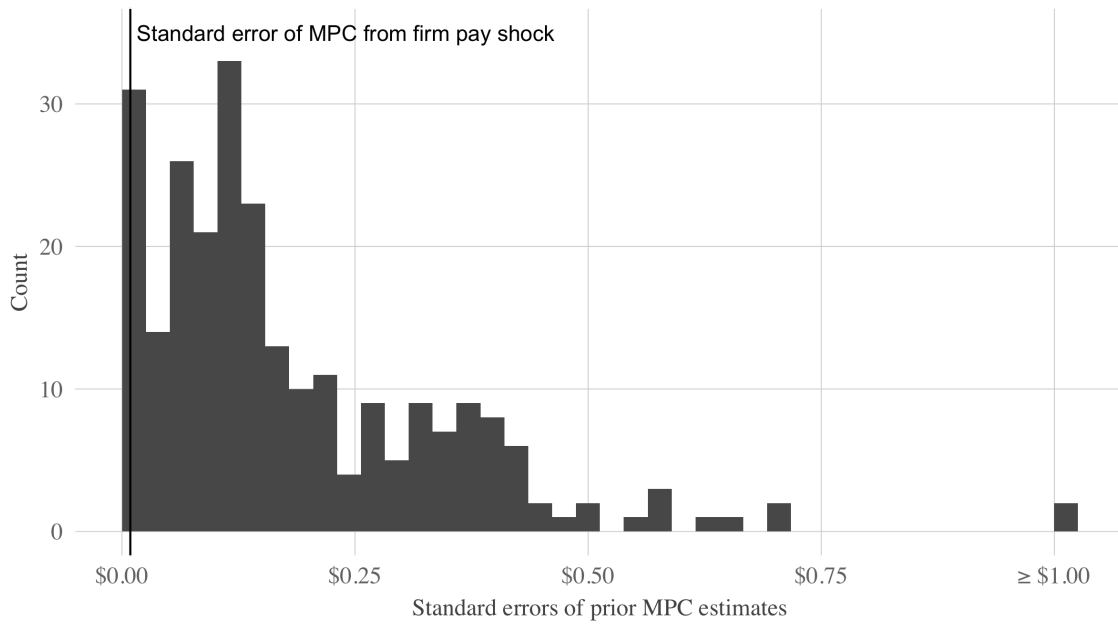Figure 14: Persistence of Consumption After Firm Pay Shocks
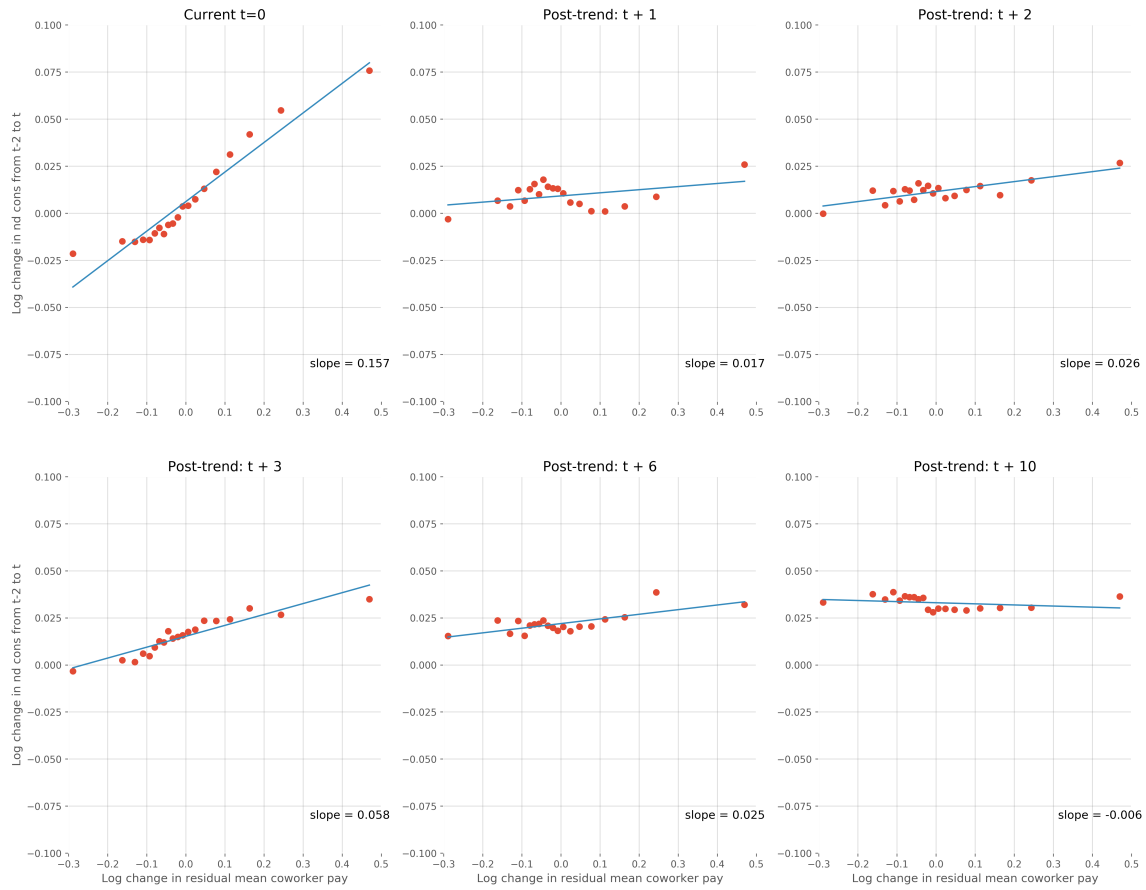
Note: This figure shows the evolution of consumption after a firm pay shock. Across all panels, the x-axis plots vingtiles of the change in residualized coworker pay from $t-1$ to t. The y-axis plots conditional means of the change in pay from $t-2$ to t in the top-left panel, $t-2$ to $t+1$ in the top-middle panel, $t-2$ to $t+2$ in the top-right panel, $t-2$ to $t+3$ in the bottom-left panel, $t-2$ to $t+6$ in the bottom-middle panel, and $t-2$ to $t+10$ in the bottom-right panel.

14

Table 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 checking account balance and total liquid and financial assets by race. See Section 3.6 for details. Sample is 2010, 2013, and 2016 Survey of Consumer Finances.

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

	Florida	Georgia	Louisiana	Sample
White	0.6354	0.5371	0.5979	0.6026
Black	0.1344	0.2964	0.2754	0.1981
Hispanic	0.1628	0.0273	0.012	0.1055
Asian	0.019	0.0204	0.0078	0.018
American Indian	0.0032	0.0011	0.0041	0.0027
Multi-racial	0.0066	NA	NA	0.0039
Other	0.016	0.0123	0.0217	0.0157
Missing	0.0226	0.1054	0.0811	0.0535

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 online Appendix Figure 2.

Table 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 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 6: 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 7: Impact of Income on Consumption by Assets and Race – Additional Specifications

	\multicolumn{8}{c}{Dependent Variable: Δ Log Non Durable Consumption}							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δ Log Income	0.169	0.185	0.205	0.117	0.129	0.150	0.227	0.242
	(0.008)	(0.007)	(0.007)	(0.007)	(0.006)	(0.006)	(0.008)	(0.008)
(Δ Log Income) x Black	0.093	0.050	0.011	0.076	0.036	−0.005	0.031	0.007
	(0.006)	(0.005)	(0.005)	(0.005)	(0.005)	(0.006)	(0.005)	(0.005)
(Δ Log Income) x Hispanic	0.044	0.022	−0.023	0.021	0.001	−0.047	−0.003	−0.045
	(0.007)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.007)	(0.007)
(Δ Log Income) x Checking		−0.234			−0.246			
		(0.008)			(0.010)			
(Δ Log Income) x Liquid			−0.259			−0.270		
			(0.009)			(0.011)		
(Δ Log Income) x Liquid(Broad)							−0.243	
							(0.009)	
(Δ Log Income) x Financial								−0.269
								(0.010)
OLS/IV	IV	IV	IV	IV	IV	IV	IV	IV
Black and Hispanic Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Asset Rank Control		Yes	Yes		Yes	Yes	Yes	Yes
Controls for lagged pay changes				Yes	Yes	Yes	Yes	Yes
Firm-by-calendar-month FE				Yes	Yes	Yes		
Observations	19,200,198	19,200,198	19,200,198	13,105,096	13,105,096	13,105,096	19,200,198	19,200,198
Adjusted R^2	0.005	0.008	0.008	0.009	0.013	0.013	0.005	0.005

Note: This table shows robustness analysis on estimates of the elasticity of consumption with respect to income ($\hat{\beta}$). Columns (1), (2), and (3) replicate columns (4), (6), and (8) from Table 2 in a specification without controlling for any lags. Columns (4) through (8) include controls for 5 lags of changes in coworker pay. Columns (4), (5), and (6) replicate columns (4), (6), and (8) from Table 2 in a specification that adds firm-by-calendar-month fixed effects. Columns (7) and (8) replicate column (8) from Table 2, controlling instead for imputations of assets using the narrow SCF definition of liquid assets and the broader definition of financial assets, respectively. Standard errors are clustered at the firm level. See Section 3.6 for details.

Table 8: Impact of Income on Consumption by Assets and Race – Specification in Levels

	(1)	(2)	(3)	(4)	(5)	(6)
		Dependent Variable: Δ Non Durable Consumption				
Δ Income	0.087	0.067	0.166	0.156	0.176	0.184
	(0.010)	(0.010)	(0.008)	(0.008)	(0.008)	(0.010)
(Δ Income) x Black		0.147		0.051		−0.008
		(0.015)		(0.011)		(0.012)
(Δ Income) x Hispanic		0.080		0.029		−0.034
		(0.013)		(0.010)		(0.011)
(Δ Income) x Checking			−0.375	−0.357		
			(0.020)	(0.020)		
(Δ Income) x Liquid(Imputed)					−0.386	−0.403
					(0.019)	(0.022)
Spec	IV	IV	IV	IV	IV	IV
Black and Hispanic Dummies		Yes		Yes		Yes
Asset Rank Control			Yes	Yes	Yes	Yes
Observations	19,416,653	19,416,653	19,416,653	19,416,653	19,416,653	19,416,653
Adjusted R^2	0.041	0.034	0.082	0.078	0.080	0.082

Note: This table replicates the IV estimates from Table 2 using income and consumption measured in levels rather than logs. See footnote 21 for additional discussion.

Table 9: 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

Public use sources: UI recipiency rate is from Department of Labor, unemployment statistics are for banked households in 2017 Current Population Survey

Table 10: Structural Model Results

	(1) Black Households		(2) Hispanic Households		(3) White Households	
A. Target moments						
	Data	Model	Data	Model	Data	Model
Average elasticity: β	0.29	0.27	0.24	0.23	0.200	0.18
Liquid assets ($)						
30th percentile	85	91	1.5	1.7	920	920
50th percentile	1,000	810	950	1,100	7,200	6,900
70th percentile	3,700	3,400	4,500	3,200	32,000	32,000
B. Estimated parameters						
Borrowing constraint (relative to income): $\underline{a}$	-0.16		-0.30		-0.22	
Low interest rate: r_1	-7.3%		-7.7%		-8.4%	
Medium interest rate: r_2	-2.4%		-0.53%		0.19%	
High interest rate: r_3	0.27%		0.23%		0.35%	
C. Welfare gain from eliminating temporary income volatility						
Average gain (consumption equivalent): $\bar{\lambda}_g$	5.7%		5.0%		3.8%	

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; this explains why the 30th percentile of the liquid asset distribution for Hispanic households is approximately zero. Panel B shows the estimated model parameters. Panel C shows the results from from eliminating temporary income volatility and calculating the welfare gain in lifetime consumption equivalents. All numbers are rounded to two significant digits.

Table 11: Structural Model Parameters

Parameter				Source
A. Common Parameters		All Households		
Coefficient of relative risk aversion: γ		1.0		N/A
Monthly discount factor: δ		0.997		Kaplan and Violante (2010)
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)
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	Black Households	Hispanic Households	White Households	
Unemployment probability: π_U	0.0693	0.0364	0.0344	CPS (2015-2018)
Monthly variance of permanent shocks (when working): σ^2_θ	0.00044	0.00047	0.00047	Calibrated to Carroll and Samwick (1997) & Farrell, Greig and Yu (2019)
Monthly variance of transitory shocks (when working): σ^2_ζ	0.142	0.148	0.148	
Lifecycle income growth: κ_t	-	-	-	CPS (2014-2016)
Probablity of death: $\aleph_t$	-	-	-	CDC (1999-2016)
Average monthly income at time of elasticity estimate	$2,008	$2,378	$3,211	Administrative Bank Data

B Matching bank and voter records

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 "best" voter record. How? Choose record with highest address score. If address scores are a tie, then choose record with highest name.

   (b) Accept 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 voters that identify as black, 99.5 percent for voters that identify as white, and 90.4 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 Hispanic. The 98.9 percent overall agreement rate is better than the agreement rate we would achieve by assigning all customers with race white (89.5 percent), Hispanic (10.7 percent), or black (10.0 percent).

C   Notes on welfare calculations

C.1   Derivation of formula for statistical welfare gain λ

We beginning with equation (12) from Section 4:

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

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:

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

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

which leaves us with:

$$1-(1+\lambda)^{1-\gamma}\,\mathbb{E}\left[e^{\beta\theta_t(1-\gamma)}\right] = 0$$

$$\Rightarrow (1+\lambda)^{\gamma-1} = \mathbb{E}\left[e^{\beta\theta_t(1-\gamma)}\right] \tag{14}$$

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 $\mathbb{E}\left[(1-\gamma)\,\beta\theta\right] = -(1-\gamma)\,\beta\sigma_\theta^2/2$ and $var\,\left[(1-\gamma)\,\beta\theta\right] = (1-\gamma)^2\,\beta^2\sigma_\theta^2$. Thus, we have:

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

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

Taking logs of both sides of (14), we have

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

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

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

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

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

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

As we can see from (15) λ is increasing in σ_θ^2. We can also rewrite equation (15) as follows:

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

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

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

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

C.2   Structural Model

The advantage of our statistical model in Section 4 is that it only requires estimates of differences in consumption elasticities 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 the racial wealth gap. To be useful in our setting, the model needs to match two sets of moments in the data: (1) differences in the distribution of liquid assets across racial groups and (2) variation in consumption elasticities across these groups. Once we are able to fit the data, we can shut off transitory income volatility and calculate the expected lifetime welfare improvement for each agent. Once again, we compare these welfare gains across racial groups.

Our baseline model is a standard buffer-stock model of consumption and saving (Carroll 1997). In each period, a household chooses how much to consume, c_t, in order to maximize expected lifetime utility. The household may save assets in each period, a_t, in a liquid savings account, which earns a gross return R at the beginning of the next period. Households may borrow at the same rate, up to a borrowing limit $\underline{a}$. Labor income, Y_t, follows the same process as in (10) 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 dying at the beginning of each period, $\aleph_{it}$. By the end of the last period, T, household assets must be

non-negative. The optimization problem can be summarized as follows:

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

$$s.t. \quad c_t + a_{t+1} = Ra_t + Y_t,$$

$$c_t \geq 0,$$

$$a_{t+1} \geq \underline{a},$$

$$Ra_T + Y_T - c_T \geq 0,$$

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. The expectation is taken over the permanent and transitory income shocks, and the unemployment and death probabilities. We solve the problem using dynamic programming, building off of the Heterogenous Agents Resources & toolKit (HARK) package (Carroll et al. 2018). 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 the level of mean monthly income by race is set to match that in the JPMCI data. 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 0.99 (Kaplan and Violante 2010). 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 online Appendix Table 11 for a complete list of calibrated parameters and their sources.

---

[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.

In order to capture the high overall consumption sensitivity and the strong asset gradient that we document empirically, the model differs from prior work analyzing buffer stock models in two ways. First, we target the distribution of *liquid* assets, while the standard approach targets total wealth. Second, to match these liquid asset targets, we *estimate* interest rates, while the standard approach calibrates them. This achieves the same end as Laibson, Repetto and Tobacman (2007) and Kaplan and Violante (2014), albeit with less structure. Those two-asset models seek to explain the *causes* of both low liquid and high illiquid asset holdings. In contrast, because we are focused on exploring the *consequences* of low liquid asset holding, a one-asset model suffices.

We estimate the model three times, once for each racial group in our sample, in order to match differences in liquid wealth accumulation over time and consumption elasticities. To match liquid asset levels, we allow for a different distribution of interest rates within each racial group. In particular, households in each group draw from one of three interest rates, with equal probability. The variation in interest rates within- and across groups is meant to capture differences in returns to assets, access to credit, investment opportunities, and other financial frictions that may affect the accumulation of liquid wealth. To further model these differences, we also allow the borrowing constraint to vary by race. An alternative approach would be to allow for heterogeneity across racial groups in preference parameters, such as discount factors or risk preferences. Because interest rates and discount factors effectively perform the same role in the economic model, it is not surprising that we can achieve similar results with this alternative approach. Without any strong reason to choose a model of discount factor heterogeneity over one with interest rate heterogeneity, we proceed with the latter. Moreover, survey evidence from Thompson and Suarez (2015) suggests that if anything, black and Hispanic households are less risk tolerant than their white counterparts, while our model would estimate the opposite in order to fit the data using preference heterogeneity.

Using the simulated method of moments, we estimate the borrowing constraints and interest rate distributions to match the actual 30th, 50th, and 70th percentile of liquid assets and the average consumption elasticities for each racial group. In particular, our simulated asset distributions and elasticities are calculated during age 40 in the model. We choose this age because it is sufficiently old that in one-asset lifecycle models there is little sensitivity of consumption to temporary income shocks. Our target asset distributions are obtained from the Kaplan and Violante (2014) definition of liquid assets in the SCF, while the consumption elasticities come from the estimates in Table 2, column (2).

A summary of our model fit and parameter estimates is presented in Table 10, Panels A and B. We find that a three-point distribution of interest rates is able to capture the distribution of liquid assets fairly well, and equally so across racial groups, with our poorest fit in the case of 70th percentile ofHispanic liquid assets. We are also able to capture the qualitative patterns in consumption elasticities across the racial groups. We estimate borrowing

constraints among Hispanic and white households of 30 percent and 22 percent of current income, respectively. Our model produces a smaller borrowing limit for black households, 16 percent of income, in order to help fit the relatively high consumption elasticity for this group.

The interest rates we infer from the data are non-standard, in that they are *negative* in many cases. The level of liquid assets, relative to income, tend to be lower than a standard buffer-stock model would predict given our other calibrated parameters. Our model therefore uses negative interest rates to fit the bottom and medians of the asset distributions. We interpret these estimates as capturing general differences in frictions in financial markets. More concretely, negative interest 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 lowest interest rates are particularly negative, which is needed to match the lowest level of assets relative to income in each group.

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 comparing the same agent in two different settings, indexed by $j \in \{1, 2\}$. The first case ($j = 1$) is our baseline structural model. In the second case ($j = 2$), we set the standard deviation of monthly transitory shocks and the unemployment probability to zero. We then re-solve the model with an agent that is 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. In each case, we can calculate lifetime expected utility, $V_i^j(c_i)$, and $c_i^j \equiv (c_{i1}, \dots c_{iT})$, the stream of consumption over time. We then calculate the proportional increase in consumption in each period that would leave the agent indifferent between a setting with and without uncertainty:

$$V_i^1 \left((1 + \lambda_i) c_i^1 \right) = V_i^2 \left(c_i^2 \right)$$

Finally, we calculate the average welfare gain, $\bar{\lambda}_g \equiv \sum_{i \in g} \lambda_i$, among households within each racial group $g \in \{black, Hispanic, white\}$.

In Table 10, Panel C, we report the average value of consumption smoothing, $\bar{\lambda}_g$, for each racial group. We find that eliminating transitory income volatility for black households leads to a 5.7 percent gain in welfare. The welfare gain for Hispanic households is a bit lower, at 5.0 percent, and white households enjoy the smallest gain, of only 3.8 percent. Overall, we find that welfare impact of differences in liquid assets and, by extension, the ability to consumption smooth largely mirror the differences in elasticities that we estimate. For example, black households have consumption elasticities that are 50 percent larger that those of white households, and the implied welfare costs of transitory shocks are likewise

50 percent larger for black households. Similarly, we estimate consumption elasticities for Hispanic households that are 20 percent larger than those of white households, and the associated welfare cost of transitory shocks is 32 percent larger for Hispanic households.