

# Walking the Tightrope

---

Variable Income and Limited Liquidity  
Among the US Middle Class

---

by **DAN SILVERMAN**

Date: **DECEMBER 10, 2020**



# Walking the Tightrope: Variable Income and Limited Liquidity Among the US Middle Class

## AUTHOR

---

Dan Silverman\*

## ABSTRACT

---

The Great Recession, and now the economic upheaval surrounding COVID-19, have intensified focus on the financial tightrope that many American families walk. Even outside of a crisis, large fractions of U.S. households face substantial variation in their incomes, with only small buffers of liquid savings. This chapter describes a body of evidence, drawn mostly from administrative data that have recently become available for academic research, showing that large fluctuations in household income are commonplace both across and within years. Even while employed, many U.S. households do not receive very steady streams of income. At the same time, these households maintain low levels of liquid savings. On the day before their paycheck arrives, fewer than 30 percent of households with at least one member working for salary have enough in their checking or savings accounts to cover 10 days of typical spending. In this way, it appears that millions of U.S. households are badly insecure. They have few resources to weather even moderately sized shocks to income or spending. Their reactions to shocks reveal, however, substantial resilience. Household

\*Arizona State University and National Bureau of Economic Research

*Acknowledgements:* Much of the evidence presented here draws on a long-running research collaboration with Michael Gelman, Shachar Kariv, Matthew Shapiro, and Steve Tadelis. I am grateful for the insights they have provided me over more than a decade of work on these subjects. The arguments offered here, while informed by those insights, are my own. Special thanks to Amy Ganz and Melissa Kearney for their many ideas about this chapter and for their challenging and constructive reactions to earlier versions.

spending declines sharply in response to big declines in income, but many find ways to re-arrange obligations and maintain large fractions of consumption. The response of spending to predictable changes in income also indicates that households prefer to rearrange spending when required, rather than reduce consumption to accumulate a buffer of savings. This implies that interventions intended to increase liquid savings buffers will have limited success and household spending will continue to move with the vicissitudes of income. These findings indicate that policy should focus on limiting the uninsured risks families face, rather than try to promote self-insurance through the accumulation of liquid savings.

## Introduction

The Great Recession, and the economic upheaval surrounding COVID-19, have intensified focus on the financial tightrope that many American families seem walk. The consequences of these crises for employment, income, and other indicators of well-being, have stoked concern about the precarious financial position of many working- and middle-class households. Even in normal times, working Americans often face substantial volatility in their incomes and have only very limited liquid savings to buffer it. Especially since the Great Recession, observers have drawn more attention to these patterns and warned that many U.S. households appear to remain badly insecure, perhaps just one missing paycheck or one costly car repair away from major and lasting economic hardship (Lusardi, Schneider, and Tufano 2011; Kristof and WuDunn 2020; Morduch and Schneider 2018; Hasler, Lusardi, and Oggero 2018).

Anonymized administrative data on both income and liquidity, recently made available for academic research, have enriched and clarified the picture on financial insecurity in the United States and elsewhere. In many ways, these new data reinforce the findings of prior surveys, focus groups, and interviews. Long, individual- and household-level histories of tax records, and similar panels of individual-level bank records, show that income varies in important ways, both across and within years. These records indicate that income volatility is not, however, a new phenomenon driven by recent changes in the structure of labor markets and the nature of work. The administrative files show little or no increase in the volatility of income in the past 25 years or so.

Despite the variability of income, or perhaps because of it, recent studies of detailed bank records show that large fractions of U.S. households maintain very limited balances in their liquid checking and savings accounts. By one recent measure, on the day before their paycheck arrives, fewer than 30 percent of families working for salary have enough in their checking or savings accounts to cover 10 days of typical spending.

These same kinds of data reveal, however, another set of facts about income, assets, and spending, that complicate the view of the financial tightrope. In particular, recent examinations of the various ways in which households respond to unexpected changes in income show that many households, despite little liquidity, exhibit substantial resilience. Detailed studies of the responses of federal workers to the big drop in pay from a government shutdown show that households use many different, and sometimes overlooked channels to cope with these often very large income shocks. Households hit by big income shocks do reduce their spending substantially. But they also defer paying bills, often at relatively modest, long-term cost or increases in debt when, as is typical, some forbearance is offered. When the negative income shocks are more persistent, these households often turn to government and social safety nets (see, Looney, Larrimore, and Splinter 2020, in this volume), but also to greater labor supply. When the shocks to income are positive, there too we see adaptation of spending and relatively little evidence of additional saving.

The same financial records that show both nimble responses to seemingly large and unexpected changes in income, and low levels of liquid assets, also reveal substantial sensitivity of spending to the arrival of even highly predictable income. Even for individuals who receive a regular paycheck, spending rises sharply on payday and the four or five days that follow. This is not just a phenomenon of the lowest-income households. Many working-class, middle-class, and high-income households maintain relatively little cash in their savings and checking accounts and spend much more in the few days after payday than they do in the few days before.

This evidence on the nimbleness of households in the face of income shocks, and on the low liquid assets and sensitivity of spending to predictable income across the income spectrum, gives perspective on the financial tightrope so many seem to walk. It suggests that the emphasis of economic theory on precautionary saving and maintaining steady levels of consumption, or “consumption smoothing,” may be misplaced. While the vast majority of working families would undoubtedly prefer greater certainty in their financial lives, they are perhaps understandably unwilling to give up a great deal of consumption in order to obtain it. They are not, in other words, willing to sacrifice a lot of important spending now—a home near good schools, decent clothes for work, replacements for bald tires—in order to secure a large liquid buffer to rely on when income is low. This appears to be due in part because many households do not seem to value consumption smoothing that much, but also because they can often use other mechanisms to help them get by.

From this perspective, well-intended efforts to improve the financial knowledge of working families, or to incentivize them to build liquid savings buffers, or to reduce imperfections in the markets in which they borrow and save, are unlikely to succeed

at getting many off the tightrope. In-depth interviews and first-hand accounts make clear that balancing in such a precarious financial position produces major strain, diminishes many aspects of well-being, and can result in lasting financial damage. At the same time, many working families seem to be remarkably adept at it, are able to lean far without falling off entirely, and do not accumulate large savings buffers even as their incomes rise. It thus seems that, when given the choice between a steady, but substantially lower, level of consumption throughout the year, or a higher average level of consumption that involves the potential for significant highs and lows, many people seem to prefer the more precarious route.

This view suggests that the path to greater financial security for middle-class households is not through efforts that encourage self-insurance with liquid savings or their fintech equivalents. Instead, the evidence on income volatility, liquid savings, and the spending response to income changes indicates that efforts to improve the financial security of middle-class households should focus on reducing the uninsured risk that they face. Those efforts might take traditional forms of social insurance, such as public unemployment and disability insurance, or mandatory paid sick-leave policies. Innovations like the emergency rental assistance program proposed in Ellen, O'Regan, and Ganz (2020) would also fall under this heading. Alternatively, policy could require or encourage employers to bear more of the income risk that their employees now face. Minimum wage policies are one form of reducing such risk at the low end of the earnings distribution. Others include predictable scheduling requirements.

As always, policymakers should consider the social cost of policies and interventions designed to reduce the income risk that workers face. That cost may be borne by taxpayers who fund social insurance programs, or by the workers themselves who may need to accept lower wages, fewer hours, or longer periods of unemployment in exchange for more stable incomes. The costs of self-insurance in terms of near-term consumption forgone are, it seems, too high for many households to accept.

## **1. Benchmark Theories of Spending, Saving, and Consumption**

Economic analyses of household spending, saving, and consumption have been guided, in large part, by theories of consumption smoothing and precautionary saving. These theories, including the permanent income hypothesis (Friedman 1957; Hall 1978), and the buffer-stock saving model (Carroll 1997), are based on the intuitive notion of a diminishing marginal utility of consumption. When consumption is already high, there is less value from having a bit more. Conversely, when the level is low, the value of additional consumption is high. This natural assumption

often implies a preference for consumption smoothing. Better to have average consumption somewhat lower but coming at a steady rate, than to have it higher on average, but subject to serious highs and lows. When income is volatile, the value of consumption smoothing motivates, in turn, precautionary saving. Households seeking to keep their consumption smooth should forgo some spending in order to build and maintain a buffer of liquid assets to spend when income is low.

As analysts use these benchmark theories to understand data on household finances, three central issues emerge. First, the theories focus on smoothing *consumption*, not spending. Data on consumption are, however, much harder to obtain than data on spending. Indeed, it is sometimes difficult to even conceive of what the right measure of consumption of some goods, like housing, cars, or appliances, should be. Testing the theories with spending data must therefore be done with care.

Second, the benchmark prediction of consumption smoothing relies on households not being “liquidity constrained.” Consumption is predicted to remain steady only if households either have enough of a savings buffer built up, or if they have access to credit markets to make up the difference when both income and savings are too low. In reality, many households may be liquidity constrained because credit markets are incomplete and, as a result, the cost of a loan may be too high relative to the value of keeping consumption smooth.

Third, the benchmark theories imply that the spending and consumption responses to income changes should depend on the extent to which those changes are predictable and persistent. In the absence of liquidity constraints, the theories predict that spending should not much respond to long-anticipated changes in income. Spending should not jump, for example, with the first paycheck after a raise associated with years of service on the job. The theory says that increase should have been largely anticipated and consumption smoothed accordingly. Similarly, a one-time and entirely surprising “transitory” change in income should also result in only a modest change in spending. As long as they weren’t liquidity constrained, the household should smooth out most of that shock to income over time. A primary exception to this consumption smoothing rule is when the household thinks of the shock to income as highly persistent or “permanent.” If something changes and from now on the household anticipates its income will be persistently higher (lower), the theory predicts spending to jump (fall) with the news.

## **2. Data Sources on Income and Assets**

It has long been challenging for researchers to measure the variability of individual or household incomes, and the extent to which that variability is unpredictable. It has been similarly difficult to measure the liquid assets that households keep on

hand to cover their expenses when income is low. Evidence about these aspects of household finances have improved, however, in recent years as administrative records have been made available to augment the self-reports of survey respondents. As a result, researchers now have a more complete picture of how variable and unpredictable household income is, and how well-buffered households are against downturns in income. In the appendix to this chapter, I offer some background on how measurement of these elements of household finance have changed. The main text focuses on some key facts that have emerged from the new sources of data.

### 3. Facts about Income Variability from Administrative Data

Recent studies based on administrative data cast doubt on both the idea that being middle class means receiving steady earnings year-to-year and the idea that earnings have become increasingly variable since the 1980s. In one especially influential study using U.S. Social Security Administration (SSA) records, Guvenen, Ozkan, and Song (2014), study the volatility of individual male earnings alone in order to isolate the effects of the macro economy on income volatility, as opposed to the influence of labor supply decisions related to childbirth and child rearing, which are concentrated among women. Figure 1(a), copied from Guvenen, Ozkan, and Song (2014), plots in blue show the standard deviation of “transitory” income shocks—that is, year-to-year changes in income from 1979 to 2011.<sup>1</sup> The plot in red similarly shows the standard deviation of “permanent” income shocks over the same period, measured as five-year changes in income. The average of the first difference tends to track the change in average earnings quite well, so it is usually close to zero. Thus, this figure gives a good sense of the distribution of the percentage changes in annual income year to year.

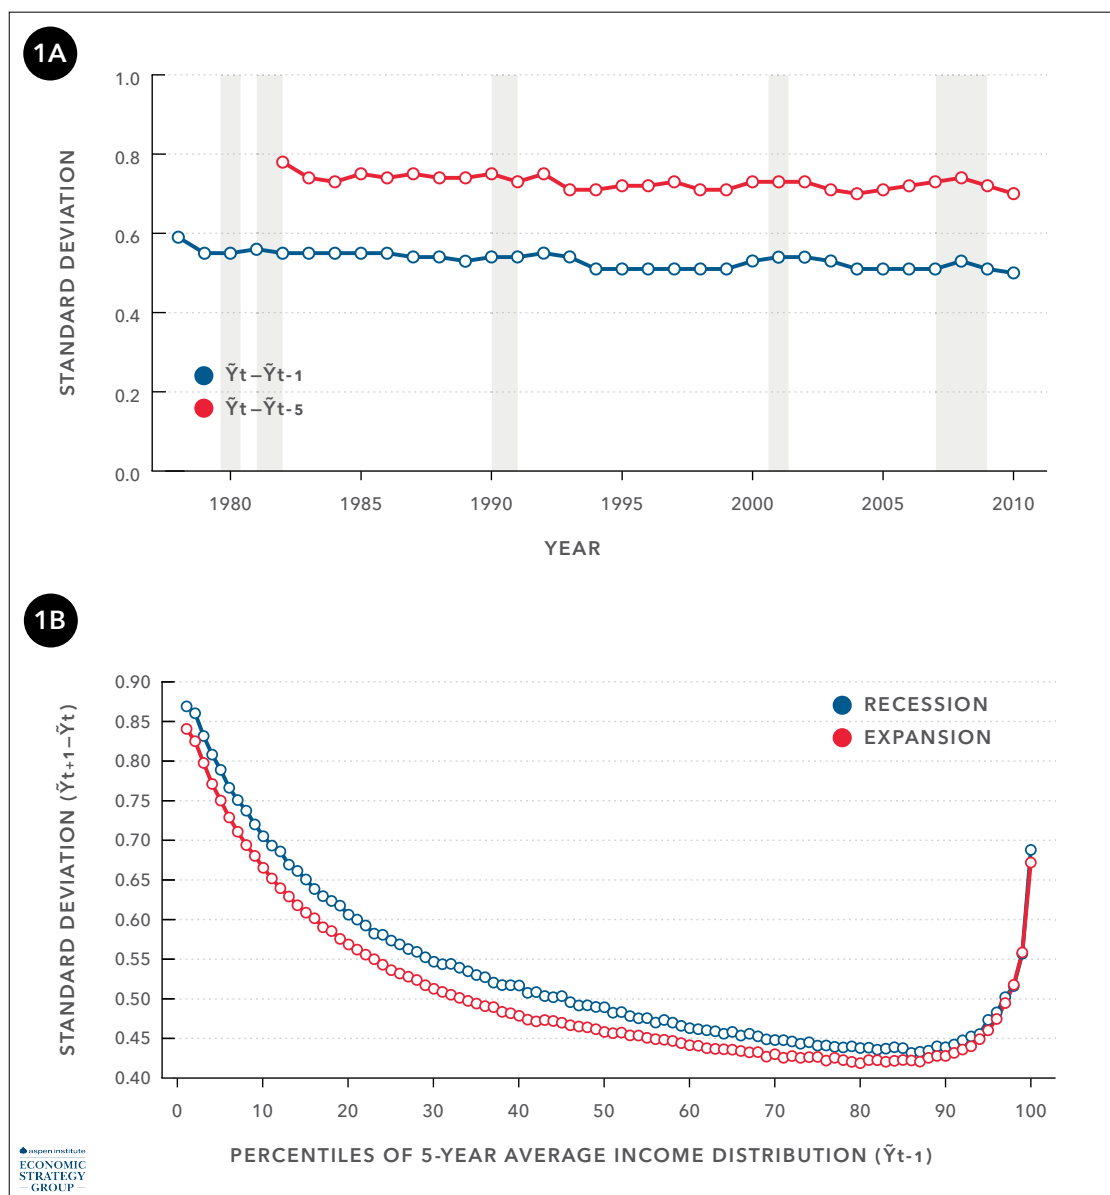
Notably, the standard deviation ranges from approximately 0.50 to 0.60 over the 33 years covered in the study and appears to decline somewhat over time. A standard deviation of log earnings of 0.5 implies about a third of the population experiences an increase in annual income greater than 50 percent or a decline of more than 40 percent.<sup>2</sup> Taking a longer horizon, the plot of the standard deviation of five-year differences in the log of annual earnings tells a similar story.

---

1 Because Guvenen, Ozkan, and Song (2014) are focused on how earnings risk depends on the macro economy, they delineate (in grey) recessionary periods. Other income, such as self-employment, business income, or asset earnings, is not be included. Recall that calculating household earnings from just the SSA files is difficult.

2 To translate the standard deviations into fractions of the population, it helps to know that the distribution of log earnings is approximately normal and, thus, so is the distribution of annual differences in earnings.

**Figure 1: The Standard Deviation of Percentage Changes in Annual Income, Over Time and by Percentile of the Average Income Distribution**



**Source:** Guvenen, Ozkan, and Song (2014)

Figure 1b, also copied from Guvenen, Ozkan, and Song (2014), shows that the variation in annual income is greatest at the very high and low ends of the income distribution. At the 10th percentile of the age-adjusted income distribution, for example, the standard deviation of the log of annual income is approximately 0.7. At the median of the income distribution this number is 0.45, and at the 99th percentile it rises again to nearly 0.7.



These results cast doubt on the idea that being middle class means receiving steady earnings year to year. The tax records show instead that, for many years now, large changes in annual earnings are not rare, including for those around the middle of the earnings distribution.


The focus on men's earnings alone could, however, be misleading about the volatility of household income. Large declines in an individual's annual earnings may, for example, be due to spells of unemployment. According to the Current Population Survey, about 70 percent of unemployed people live with someone who is currently employed, and 61 percent with someone working full time. As a result, the percentage change in the annual income of the whole household might be more modest. Alternatively, if the variation in earnings derives largely from factors besides unemployment, and if the income of household members is highly correlated, then we would expect household income to be about as variable as individual earnings.

DeBacker et al. (2013) assess this question using panel income data from tax returns over the years 1997–2009. First, isolating men's earnings from these tax returns, that study finds a similar, if somewhat smaller, degree of volatility in that source of income than what Guvenen et al. (2014) show.<sup>3</sup> In the tax return data, the standard deviation of percent changes in men's annual earnings rises from around 0.40 to around 0.43 over this period—lower than what the SSA records in Guvenen et al. (2014) reveal, but not dramatically so. Adding up over all sources of income in the household does not change the story. Figure 2, reproduced from DeBacker et al. (2013), shows that the standard deviation in total household income is similarly volatile, ranging from about 0.39 to 0.43 over the period. In this way, the presence of multiple employers and multiple workers does not appear to reduce importantly the volatility of a household's income.<sup>4</sup>

---

3 Splinter (2019) argues that the DeBacker et al. (2013) analysis understates income volatility because it excludes individuals with very low earnings in any of the years under study.

4 Pruitt and Turner (2020) also study income tax returns and find more evidence that spousal earnings function to provide a buffer against the volatility of men's earnings. Especially during recessions, and especially at the extremes of the average household income distribution, they find that the dispersion of the distribution of longer (four-year) changes in household earnings is substantially lower than that for men's earnings alone.

**Figure 2. Standard Deviation of the Percentage Change in Income**

**Source:** DeBacker et al. (2013)

### 3.a. *Uncertain or Simply Variable?*

A key issue surrounding measures of individual and household income volatility is the extent to which these the changes in income are predictable. Some of the increases and decreases in income must be foreseeable by households—but how many, and by when could they have known these changes would occur? Understanding the extent to which movements in income come as a surprise is important because anticipated income increases or decreases have different implications for household finances and well-being than similarly large but surprising changes in income.

Relevant examples of predictable income movement include changes in household earnings on account of seasonal work—say, among construction workers or workers in temporary tourist jobs. These kinds of changes are different from those associated with, for example, an unanticipated decline in hours or sales, or even a layoff, and we would expect households to react to such changes differently. Workers can also often anticipate, at least with some advance warning, a boost to earnings from a job promotion, or an increase in pay associated with years on the job. Reductions in income associated with a move to part-time work can sometimes be anticipated as well. The same is true for movements in individual or household earnings associated with changes in family structure, including marital status or number of children living at home. These often large changes in income may have very different consequences for a household's balance sheet than similar-sized changes attributable to changes in hours or wages at work.

It is challenging to determine from even the best available data the extent to which households anticipate the changes in income that they encounter. Statistical models, augmented by economic theory, can identify changes in income that households ought to have been able to predict, or that they act as if they did predict (e.g., Blundell and Preston 1998; Moffitt and Gottschalk 2002; Primiceri and van Rens 2009; Gelman et al. 2020.) The estimates of the predictable percentage of income volatility that emerge from these approaches are sensitive to the assumptions made and methods used, and they cover much of the range from zero to 100 percent. This makes sense in part because the expectations or behavior of households might reasonably differ from those derived from a particular, even very flexible econometric or theoretical model. Decompositions of income volatility, as in Larrimore, Mortenson, and Splinter (2016), which quantify the extent to which household income movements are associated with easily anticipated events like marriage or the birth of a child, suggest that substantial fractions of annual income volatility are predictable, but a comprehensive and robust quantification remains elusive.

#### **4. Recent Facts about Liquid Asset Holdings**


Even if substantial fractions of the changes in income were predictable, we would still expect households to maintain substantial liquid assets to buffer against unpredictable changes or simply in anticipation of predictable declines in income. For a large fraction of households, especially middle-income ones, this is not the case. Many live, more or less, from paycheck-to-paycheck.

A recent study using de-identified data from the users of a financial aggregation app, examines the levels and high-frequency variation in liquid asset holdings. That paper examines how liquid asset holdings change over the days between paychecks (the pay cycle). Figure 3, taken from that paper, shows median liquid assets over a two-week pay cycle, by terciles of the distribution of liquid assets.<sup>5</sup> To make the measurement of liquid asset levels more comparable across income groups, the measure is expressed as a ratio of checking and savings account balances to average daily total spending.

---

5 In these data, liquid asset balances peak two days after a payday. The balance data are based on funds available, so liquid assets should lag behind payday according to the banks' funds-availability policy. There is at least a one-day lag built into the data because the aggregator application collects balance information during the day, and will reflect a paycheck posted only on the previous day. Note too that liquid asset balances reflect the net inflows and outflows. Spending done just after the receipt of paycheck will therefore lead daily balances to understate gross asset balances right after the receipt of the paycheck.

**Figure 3: Median Liquid Assets over the Paycheck Cycle for the First, Second, and Third Terciles of the Liquid Assets Distribution**



**Source:** Gelman et al. (2018)

This figure shows that households in the top third of the liquid asset distribution are in a different situation from the rest. Over all days, this group maintains a median of 32.1 days of average expenditure in liquid assets. Without tapping into credit, or illiquid forms of savings, or any social or government support, this group could go more than a month without income and still maintain their usual levels of spending. Even on the day before their paycheck arrives, this part of the population has a substantial buffer in their liquid checking and savings accounts.

For the bottom two-thirds of the liquid assets distribution, however, the financial situation looks much less secure. The money in their checking and savings accounts are not adequate for maintaining their typical levels of spending, even for a single pay period. The median number of days of average expenditure maintained in liquid assets is nearly eight for the middle tercile and three for the lowest tercile. Liquid assets are especially low the day just before a paycheck arrives, when the bottom third has a median balance of zero in their checking and savings accounts, and middle third has only four days. Despite, or perhaps because of, substantial income variation, this large part of the population lives, in effect, from paycheck to paycheck.

## 5. How Do Households Respond to Large, Unexpected Income Shocks?

The findings about income volatility and liquidity described above are not easily reconciled with standard frameworks of household finance. If income is subject to large movements, why wouldn't most households maintain a larger buffer and make consumption less subject to changes in income?

Insight can be gathered from studying how household finances respond to large and unambiguously unexpected shocks. One study examined how federal employees responded to the U.S. government shutdown of 2013. At the time, the federal government employed about 2.1 million affected workers, each of whom saw one paycheck in October of 2013 reduced by about 40 percent, or roughly four working days out of the 10 in a typical pay cycle. Only about 800,000 of these workers were furloughed. The rest had to work for at least part of the shutdown, and so we would expect their work-related expenses to be little changed. While they did not know how long the shutdown would last, federal employees should have expected their missing earnings to be recouped as soon as the shutdown ended. In the past, Congress had always done that, even for furloughed employees.

As a group, the federal workforce might reasonably be described as (upper) middle class. At the time, the average federal worker earned about \$82,000 a year. Like the households described above, they also maintained quite limited liquid assets. At least two-thirds didn't have enough in their checking and savings accounts to cover what was missing from this paycheck. So, what did they do?

They adjusted. They spent less, briefly delayed making payments on mortgages, rent, and credit card balances, and most emerged without lasting damage to their finances. As a result, while spending decreased, consumption likely moved much less as households managed their bills in order to meet most of their usual needs.

Figure 4 shows the details. In each panel, the dashed blue line marks the start of the shutdown, the first red line indicates the week when the smaller-than-usual paycheck arrived, and the last red line marks the end of the shutdown and the recouping of that missing pay. In panel (A), we see the response of "non-recurring" expenditures, a large category of spending that is arguably more discretionary in that it excludes things like rent, mortgages, and other bills that are paid in very regular amounts and with very regular frequency. Non-recurring expenditures dips by about a day and a half of average spending in this category and bounces right back up when the missing pay is received. The recurring expenditures category,

which includes rent, mortgages, etc., is different. In panel (B) we see it drops more sharply, by more than two days of average spending, and recovers only incompletely when the missing pay arrives. Panel (C) isolates a particularly important part of recurring spending: mortgage payments. The panel shows too that these payments are delayed substantially by the missing pay, and take some time to recover once the missing pay is recouped.

The rest of the evidence paints the picture of what federal workers did to weather this shock. They didn't bolster spending by making transfers from other accounts to their checking accounts (panel D), and they didn't use their credit cards much more to cover the gap (panel F). Instead, they deferred the payments they would have otherwise made on the credit cards and floated a couple more weeks until they got paid. As with the delayed mortgage payments, this was an important way that households found to meet their consumption needs despite the fact that they had few liquid assets to rely on.

Delaying payments on mortgages, rent, or credit cards can be costly, including late fees and interest on the credit being extended and the potential for damage to credit scores. This study showed, however, little lasting damage from delayed payments, in part because most mortgages and rent typically allow a grace period during which late fees are not charged. For many households, these costs were also avoided because they tend to pay their mortgages and credit card bills whenever they are liquid, not when the bills are due. This sometimes involves making multiple payments per month. As a result, many households affected by the shutdown, who normally would have made a full payment on their bills as soon as they were paid, still had time to wait before those bills came due.

A recent study of a very different episode, this time involving an unexpected rise in disposable income, shows similar adaptation and a sensitivity of spending to income. The study uses administrative data from a financial aggregator to examine how household spending responded to a large decrease in gasoline prices in 2014. Gasoline is an important part of the budget for many, especially middle-class, households. Depending on the survey method, the median household spent \$2,000-\$3,000 per year on gasoline in the months leading up to the price decline. During the October 2014 episode, the price rapidly dropped by almost half.

**Figure 4. Estimated Response of Spending Categories to Government Shutdown**


**Source:** Gelman et al. (2018)

**Notes:** The spending, payment, or transfer category in each panel is normalized by the individual-level daily average for that category.  $N = 3,804$  and  $N = 94,680$  for the treatment and control groups, respectively.

The short- and medium-run elasticity of gasoline spending to changes in price is quite low. Over periods of less than a year, people do not drive much more or less, or switch modes of transportation in response to gasoline prices. As a result, the 2014 gas price drop was like a substantial increase in the after-tax income of households. How did they react?

Figure 5, reproduced from that study, shows that within three to five months, households adapted to this new-found cash made available from gas savings, and spent it on other things. The figure shows, in particular, the estimated marginal propensity to spend (often called the marginal propensity to consume or MPC) of this new-found discretionary income. When the MPC equals zero, that means every penny of savings from a less expensive gasoline is saved; the data do not show it being spent on other things. When the MPC equals 1, that means every penny saved on gas is spent on something else. Figure 5 shows that three months after the price decline, households are already spending the vast majority of those savings on other things. By five months, the point estimate of the MPC is effectively 1.

**Figure 5: The Propensity to Spend an Increase in Discretionary Income Deriving from a Large and Sustained Reduction in the Price of Gasoline**



**Source:** Gelman et al. (2019)

**Notes:** The figure reports estimates of the marginal propensity to consume (MPC) out of a large and sustained reduction in gasoline prices in 2014 as a function of the time, in weeks, since the reduction. Dashed lines show 95 percent confidence interval. See text for further details.



The spending response to this large decline in gas prices thus indicates, again, a tendency for household consumption to track income quite closely. When extra discretionary income arrived in the form of lower gas prices, households tended to spend it rather than save a portion of it in case of a change for the worse.

## **6. How Do Households Respond to Predictable Changes in Income?**


The preceding evidence on the responses of households to unexpected changes in (discretionary) income shows that the spending of many households follows their income more closely than benchmark theories of household finance would predict. Given the evidence on liquid asset balances, this is less of a surprise. If households maintain such a small buffer of cash in their checking and savings accounts, it makes sense that shocks to income will produce a sizeable spending response. The seeming ability of households to maintain much of their usual consumption (if not spending) despite a large, if brief, decline in income helps explain why.

Another reason why spending follows income more closely than benchmark theories might predict is suggested by the response of expenditures to even very predictable changes in income. By way of reminder, the standard model used in neoclassical economics predicts that if changes in income are expected, households will choose to smooth consumption across periods of varying income. But high frequency data on household consumption show that is *not* what households choose to do. These results indicate that for many households consumption smoothing may not be as valuable as those standard theories assume.

Figure 6 is taken from another study of financial aggregator data which examined how spending by different households responded to the arrival of paycheck income. Here, paycheck income is defined so as to be especially predictable—regular in its timing (every two weeks) and amount. As the first panel shows, total spending responds sharply—up nearly 75 percent above average daily expenditure—to the arrival of a paycheck and then settles down for the second week until the next paycheck arrives.


Some of this response of total spending can be explained by the coordination of bills like rent, mortgages, or utilities, with the arrival of paycheck. This kind of coordination is not a refutation of the value of consumption smoothing, just a sensible method of avoiding problems of liquidity. But the second panel shows that non-recurring spending also responds, though less sharply, to the arrival of a paycheck. It too rises by 40 to 50 percent of average daily spending in this category when the check arrives, before returning to more normal levels later in the week. Only spending on items like fast food, in the third panel, seems to follow the standard prediction that consumption should be smooth over the pay cycle.

**Figure 6. Estimated Response of Spending Categories to the Arrival of Predictable Paycheck Income**



**Source:** Gelman et al. (2014)

**Figure 7. Estimated Response of Non-recurring Spending to the Arrival of Predictable Paycheck Income, by Liquid Asset Tercile**



**Source:** Gelman et al. (2014)

Figure 7, taken from that same study, shows an interesting link between smoothing and maintaining a liquid asset buffer. The non-recurring spending of households in the bottom third of the liquid assets distribution is much more responsive to the arrival of a predictable paycheck than is the spending of those with larger buffers. Notably, however, even those in the highest liquid assets tercile, who typically maintain more than a month of average spending in their checking and savings accounts, still spend in response to the arrival of their paychecks. Their non-recurring spending is about 20 percent higher than usual on the day the check arrives and remains higher for about a week before settling down for the week before the next paycheck arrives. Even this group, it seems, does not seek perfectly smooth spending at this frequency.

## **7. Discussion and Implications for Policy**

This chapter described new sources of administrative data on income, spending, liquid assets and debt that provide a novel perspective on the finances of millions of households. Those data show that individual and even household earnings vary in important ways from year-to-year. Despite, or perhaps because of, those large movements in annual income, echoed in higher frequency variation, the same kinds of data show that many middle-class families maintain very little by way of liquid assets. If, for example, they were to miss just one paycheck by surprise, large majorities could not use their checking or savings balances to cover their usual levels of spending until the next paycheck arrives. This is not just a phenomenon of very low-income households. The tendency to maintain relatively few liquid assets is pervasive well into the middle of the income or spending distribution.

While they face large movements in income with little liquid assets, households also display significant resilience to income shocks. Integrated financial records show that middle-class households use often overlooked methods of lowering expenditures without dramatically reducing consumption, and find ways to get through at least short-term but large declines in income. Even in “normal” times, however, spending often reacts sharply to the entirely predictable arrival of income, like a regular paycheck.

Taken together, these findings indicate that standard economic analyses may have over-emphasized the value to many households of maintaining precautionary saving to keep spending and therefore consumption smooth. Most working families would, undoubtedly, be glad to trade the ups and downs of living paycheck-to-paycheck for a steady income and a nice rainy-day fund. But, in reality, there are likely to be substantial costs of making that trade. A steadier income might require accepting a

lower income on average (a lower wage or fewer hours). Saving a substantial buffer often means very careful budgeting with unpleasant and certain sacrifices now in anticipation of the possibility that things may be even worse later. For many, it seems, those costs to total income or nearer-term spending are not worth bearing in part, perhaps, because they can often use other mechanisms to help them get by when income is low.

In this view, policy aimed at promoting greater financial security for the middle class would do better to reduce its emphasis on efforts to encourage self-insurance through accumulating precautionary saving buffers or their fintech equivalents. Financial education and realistic subsidies for liquid savings would seem to have only limited potential to move many families away from having their consumption move closely in time with their income.

Instead, efforts to improve the financial security of middle-class households would seem to do better by focusing on the uninsured risk these households face. Traditional forms of social insurance have this feature. Public unemployment, disability, and health insurance, or mandatory paid sick leave policies help reduce large movements in (effective) income due to uncertain employment or health. Alternatively, public or private policy could lead employers to bear more of these risks that their employees now face. Predictable scheduling requirements or work sharing policies have this feature.

These conclusions derive from analysis of vast, impersonal, administrative datasets reflecting the behaviors and outcomes of millions of Americans. Notably, the detailed, personal, more ethnographic analyses like those in Morduch and Schneider (2018) come to similar policy conclusions, though for different reasons. Those much more intimate and contextualized analyses of middle-class finances also conclude that efforts at financial literacy or facilitating the management of liquid savings are better replaced by policies to shift risk from worker to firm or, perhaps by payment systems that facilitate smoothing even when paychecks fluctuate. This is in part because they think firms have recently shifted more of the financial risk onto their employees, and because managing such high levels of income volatility is too challenging, even for those who are very sophisticated about financial matters (Ogden and Morduch 2017).

The evidence summarized in this chapters suggests that income volatility is not such a new phenomenon and that recent moves toward more unpredictable hours are unlikely to be driving it. Similarly, the findings described here suggest that the lack liquid savings has less to do with the challenges of determining how to accumulate a buffer than with the profound costs doing so, and the alternatives available for weathering downturns in income.

If, indeed, moves to reduce the uninsured risks that middle-class families face are the better path, policy makers must then confront the costs of doing so. Those costs may be to taxpayers who fund social insurance programs, or to the workers themselves who may need to accept lower wages, fewer hours, or longer periods of unemployment in exchange for more stable incomes. The hope is that by providing such insurance at scale the costs will be less than those of self-insurance which, it seems, are too high for many households to accept.

## References

- Blundell, Richard, and Ian Preston. 1998. "Consumption Inequality and Income Uncertainty," *The Quarterly Journal of Economics* 113(2): 603–640.
- Bound, John, Charles Brown, Greg J. Duncan, and Willard L. Rodgers. 1990. "Measurement Error In Cross-Sectional and Longitudinal Labor Market Surveys: Results From Two Validation Studies." In J. Hartog, G. Ridder, and J. Teevnes, eds., *Panel Data and Labor Market Studies*, Elsevier, 1-19.
- Christopher D. Carroll. 1997. Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis, *The Quarterly Journal of Economics* 112(1): 1–55.
- DeBacker, Jason, Bradley Heim, Vasia Panousi, Shanthi Ramnath, and Ivan Vidangos. 2013. "Rising Inequality: Transitory or Persistent? New Evidence from a Panel of U.S. Tax Returns," *Brookings Papers on Economic Activity*, Spring: 67–142.
- Ellen, Ingrid Gould, Katherine O'Regan, and Amy Ganz. 2020. "A Renter Safety Net: A Call for Federal Emergency Rental Assistance." In *Securing Our Economic Future*, edited by Amy Ganz and Melissa Kearney. Aspen Institute Press.
- Friedman, Milton. 1957. *A Theory of the Consumption Function*. Princeton University Press: Princeton, N.J.
- Gelman, Michael, Shachar Kariv, Matthew D. Shapiro, and Dan Silverman. 2020. "Rational Illiquidity and Consumption: Theory and Evidence from Income Tax Withholding and Refunds." Working Paper No. 25757, National Bureau of Economic Research.
- 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.
- . 2018. "How Individuals Respond to a Liquidity Shock: Evidence From the 2013 Government Shutdown," Forthcoming, *Journal of Public Economics*.
- Gelman, Michael, Yuriy Gorodnichenko, Shachar Kariv, Dmitri Koustas, Matthew D. Shapiro, Dan Silverman, and Steven Tadelis. 2019. "The Response of Consumer Spending to Changes in Gasoline Prices." Working Paper No. 22969, National Bureau of Economic Research.
- Güvenen, Fatih, Serdar Ozkan, and Jae Song. 2014. "The Nature of Countercyclical Income Risk." *Journal of Political Economy* 122(3): 621–60.

- Hall, Robert E. 1978. "Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence." *Journal of Political Economy* 86(6): 971–987.
- Hasler, Andrea, Annamaria Lusardi, and Noemi Oggero. 2018. "Financial Fragility in the US: Evidence and Implications." GFLEC Working Paper No. 2018-1.
- J.P. Morgan Chase Institute. 2015. *How Falling Gas Prices Fuel the Consumer: Evidence from 25 Million People*. Available at <https://www.jpmorganchase.com/corporate/institute/document/jpmc-institute-gas-report.pdf>.
- Kaplan, Greg, Giovanni L. Violante, and Justin Weidner. 2014. "The Wealthy Hand-to-Mouth," *Brookings Papers on Economic Activity* 45 (Spring): 77–153.
- Kristof, Nicholas, and Sheryl WuDunn. 2020. *Tightrope: Americans Reaching for Hope*. Knopf Doubleday Publishing Group. New York City, NY.
- Larrimore, Jeff, Jacob Mortenson, and David Splinter. 2016. "Income and Earnings Mobility in U.S. Tax Data." In *Economic Mobility: Research & Ideas on Strengthening Families, Communities & the Economy*, edited by Federal Reserve Bank of St. Louis and the Board of Governors of the Federal Reserve System, 481–516.
- Looney, Adam, Jeff Larrimore, and David Splinter. 2020. "Middle-Class Redistribution: Tax and Transfer Policy for Most Americans." In *Securing Our Economic Future*, edited by Amy Ganz and Melissa Kearney. Aspen Institute Press.
- Lusardi, Annamaria, Daniel Schneider, and Peter Tufano, 2011. "Financially Fragile Households: Evidence and Implications," *Brookings Papers on Economic Activity*, Economic Studies Program, The Brookings Institution 42(1): 83-150.
- Meyer, Bruce D., and Nikolas Mittag. 2019. "Using Linked Survey and Administrative Data to Better Measure Income: Implications for Poverty, Program Effectiveness, and Holes in the Safety Net." *American Economic Journal: Applied Economics* 11(2): 176–204.
- Moffitt, Robert and Peter Gottschalk. 2002. "Trends in the Transitory Variance of Earnings in the United States." *Economic Journal* 112: 68–73
- Morduch, Jonathan, and Rachel Schneider. 2018. *The Financial Diaries: How American Families Cope in a World of Uncertainty*. Princeton University Press: Princeton, N.J.
- Ogden, Timothy, and Jonathan Morduch. 2017. "Too Many Americans Suffer from Financial Instability. Their Employers Can Help Fix It." *Harvard Business Review*, December 14.
- Parker, Jonathan A. 1999. "The Reaction of Household Consumption to Predictable Changes in Social Security Taxes," *American Economic Review* 89(4): 959–973.
- Primiceri, Giorgio and Thijs van Rens. 2009. "Heterogeneous Life-Cycle Profiles, Income Risk and Consumption Inequality." *Journal of Monetary Economics* 56(1): 20–39.
- Pruitt, Seth, and Nicholas Turner. 2020. "Earnings Risk in the Household: Evidence from Millions of US Tax Returns." *American Economic Review: Insights* 2(2): 237–254.
- Splinter, David. 2019. "Income Mobility and Inequality in the United States: Evidence from Tax Data since 1979." Working paper.

## **Appendix:**

### **Background on Sources of Data for Income and Assets**

#### ***A.1. Self-Reports on Surveys Have Long Provided the Best Measures***

In the past, most efforts to measure the variability of different sources of income, or the levels of household liquidity, have relied on self-reports in survey responses. The organizations that collect these surveys expend substantial effort and resources to obtain representative samples and accurate measures, but they face important obstacles on both fronts.

For purposes of measuring income and its changes, the leading survey instruments include the Panel Study of Income Dynamics, conducted by the University of Michigan, the Survey of Income and Program Participation, conducted by the U.S. Census Bureau, and the Current Population Survey also conducted by the Census Bureau on behalf of the Bureau of Labor Statistics. These surveys have long collected household and individual income information via self-reports and mostly by the phone. The cost of implementing high-quality surveys make it difficult to collect income information at high frequency. In addition, and despite both sincere efforts to report accurately and encouragement to use paystubs and tax records, survey participants naturally struggle to recall their incomes with precision. Self-reports are prone to heaping on round values and to relatively stickiness, year-to-year. Validation studies of survey data like these, such as Bound et al. (1990), or Meyer and Mittag (2019) indicate substantial, if mostly non-systematic, measurement error. In Bound et al. (1990), for example, when survey responses were compared to administrative income records, the error to variance ratio was as high as 0.3 and higher-income earners tended to underreport their income and lower-income households tend to overreport.

These errors and biases can be especially important when trying to measure the variability of income. On the one hand, some of what appears like fluctuations could actually represent the challenges of accurate recall—forgetting about some income in one time period, overstating it in another. On the other hand, the tendency to approximate current income by past income might dampen true fluctuations.

Using surveys to collect accurate data on liquid assets is, perhaps, even more challenging than collecting income information. Leading studies of wealth, debt, and liquidity, such as the U.S. Federal Reserve Board's Survey of Consumer Finances (SCF), are conducted mostly in-person in order to obtain a more accurate picture of a household's financial situation. Taking an inventory of a household's accounts and accurately measuring their balances, in person, takes time. The SCF typically requires



about 80 minutes to complete. The expense of getting this kind of information means such studies must be limited in their frequency and sample size. The SCF, for example, is conducted only every three years with a sample of approximately 6,500.

### ***A.2. Measures Drawn from Administrative Records Are Now Common***

The many challenges of obtaining accurate information on income levels and volatility, and on and liquid assets from surveys have, especially over the past 10 years, led researchers to develop administrative datasets for these and other purposes. These new datasets avoid many of the difficulties that surround self-reports and are often enormous in size. In some cases, these data are collected not just from a representative sample of the United States, but include nearly the entire population of U.S. households. In other cases, the datasets are orders of magnitude larger than even the largest survey samples, but include only those individuals or households who hold accounts with certain (large) financial institutions or who choose to integrate their many financial accounts with one financial aggregating platform. The accuracy and size of these administrative data sources are usually limited. Compared with comprehensive surveys, these administrative data sets typically contain relatively little information about the demographic, social, and other characteristics of the individuals and households from which they draw.

The initiatives of several financial services firms have been especially important in improving the measurement of income and liquid asset holdings. Firms like JP Morgan Chase, through its JP Morgan Chase Institute, and the Vanguard Group through its Vanguard Research Initiative, have collaborated with academics and begun using de-identified and aggregated data from their millions of account holders to provide publicly available research on the savings buffers that individuals and households maintain. Financial aggregator firms, such as Mint, Mint Bills, Meniga, and Yodlee have also collaborated with academics to make similar research possible.

There are many distinctions of these data from financial services firms, even when compared with other sources of administrative records such as tax records. Because they often integrate records from an individual's or household's many different accounts, including checking, savings, and credit, these files can give a remarkably complete view of a household's financial situation. They allow, in particular, the simultaneous measurement of income, liquid asset holdings, liquidity more generally, spending, debt, and even credit scores. These data are also often available at very high (even daily) frequency and in real time. They can thus be used to bring a laser-sharp focus to the financial consequences for particular groups of people or specific events and times.

Administrative data from financial services firms have limitations as well. Non-random selection into the provider can be important, especially for aggregators which require motivated opting in. The administrative data from financial services firms also make it difficult to determine the unit of observation. If the data come from just one firm, joint accounts can be identified but the existence of accounts outside the firm usually cannot be ruled out. Aggregators can provide more comprehensive data, but they rarely provide information about whether these accounts are jointly held and, unless aggregator users do it themselves, it is usually impossible to determine which accounts ought to be gathered in one household. For nearly all such data, income measures must be after income tax withholding and any pre-tax purchases or saving. These features make it more difficult to pin down the level of pre-tax income with which many observers are concerned. These relative limitations lead many researchers to specialize somewhat: relying on tax records for measuring income and financial services records for measuring liquidity and spending.

Before de-identified financial services records became available for research, the leading source of administrative data on income was government tax records. For example, the U.S. Social Security Administration's (SSA's) earnings data, used to calculate payroll taxes and benefits, have been made available to qualified researchers and used for many years to study the lifetime dynamics of income. More recently, those records have also been used to study the volatility of income.

Because they are derived from W-2 forms, however, the SSA's earnings file will not include many potentially important sources of income volatility. These sources include self-employment income, business income, asset income, and government benefits. In addition, because the SSA collects these records largely for purposes of calculating payroll tax receipts and eventual Social Security benefits, the records are organized at the individual rather than the household level. These data are therefore best suited to studying fluctuations in the earnings, but not the total income, of individuals and not of households.

Data from income tax returns and associated reporting to the U.S. Internal Revenue Service (IRS) have most of the same advantages of the SSA data and more. The IRS tax return data that have been made available for academic study (under strict confidentiality agreements and data safety protocols) tend to include shorter time periods. The IRS data include both large majorities of U.S. workers plus anyone who files a tax return. Thus, the IRS tax return data allow researchers to study total income (not just earnings) fluctuations. As important, the IRS data are organized for purposes of calculating and collecting income taxes that often depend on family structure. This organization of the data makes it much easier to study income fluctuations at both the individual and the household level.