Liquid Wealth and Consumption Smoothing of Typical Labor Income Shocks*

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

August 21, 2025

Abstract

We identify exogenous, transitory, and unpredictable shocks to labor income by combining elements from both the quasi-experimental and structural approaches in the consumption smoothing literature. We find that household consumption is highly sensitive to monthly labor income shocks. This suggests that temporary income volatility has a large welfare cost. Furthermore, consumption is most sensitive for households with low liquidity and almost unchanged for households with high liquidity. Our findings that consumption is sensitive in a welfare-relevant setting and that there is a steep, precisely-estimated liquidity gradient help to distinguish between competing classes of consumption models.

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

1 Introduction

Many households face substantial fluctuations in labor income from month to month. For example, the standard deviation of monthly labor income changes is around 30 percent in the United States.[1] At the same time, households report holding few liquid assets that could be used to buffer such volatility. According to the Survey of Consumer Finances, 40 percent of Americans have less than two weeks' worth of income in liquid assets. This combination of high monthly income volatility and low liquid asset holdings raises concerns that it may be difficult for households to smooth consumption from month to month. Yet, perhaps surprisingly, there is little empirical evidence on how typical month-to-month fluctuations in labor income affect consumption and even less evidence about how this consumption sensitivity varies with liquid wealth.

The lack of clear evidence on the consumption sensitivity to monthly income fluctuations has contributed to theoretical uncertainty in explaining why consumption might be sensitive to transitory income changes. The most compelling evidence of sensitivity comes from the response to unusual windfalls of non-labor income.[2] But these windfalls are rare, and data to evaluate their consumption impact is even rarer. It is therefore difficult to distinguish between competing explanations for sensitivity. For example, is sensitivity driven by low liquidity, as in benchmark models, or is it elevated even for high-liquidity households, as in some recent models? Unfortunately, the windfall studies are often underpowered to detect how this sensitivity varies with liquid assets. Furthermore, since windfalls are infrequent, failing to smooth them has small welfare costs. This has generated yet a third potential explanation for sensitivity: maybe sensitivity to windfalls is simply a "near-rational" response to the low stakes in these settings, whereas households would actually smooth welfare-relevant shocks (like monthly income volatility), in line with the Permanent Income Hypothesis.

The goal of this paper is to construct precise estimates of the consumption response to typical labor income shocks and investigate how these vary by liquid wealth. By doing so, we hope to learn about the consequences of monthly income volatility and to test between competing theories of consumption sensitivity.

To precisely measure the consumption response to changes in labor income, we develop an empirical strategy that builds on the strengths of the two main traditions in the consumption smoothing literature. One tradition uses quasi-experimental variation in income. These research designs cleanly identify exogenous shocks, but they rely on unusual windfalls. A second tradition identifies income shocks using covariance restrictions within a structural model of the earnings process.[3] This approach successfully captures variation from more typical income sources and can isolate the unpredictable transitory shocks that are often the target for consumption models. However, this methodology may suffer from concerns about the endogeneity of labor supply, which can vary with consumption needs. Our approach seeks to combine the strengths of both traditions, identifying exogenous, yet typical, transitory, and unpredictable labor income shocks.

First, to isolate *exogenous* changes in income, we develop an instrument in the tradition of the

[1] See Ganong et al. (2025) for U.S. evidence, Druedahl, Graber, and Jørgensen (2023) for Denmark, and Brewer, Cominetti, and Jenkins (2025) for the U.K.

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

[3] See Arellano et al. (2023), Attanasio and Pavoni (2011), Blundell, Pistaferri, and Preston (2008), Blundell, Low, and Preston (2013), Choi, McGarry, and Schoeni (2016), Etheridge (2015), Guvenen et al. (2014), and Hall and Mishkin (1982).

quasi-experimental literature. Simply regressing consumption changes on all income changes could lead to bias in the presence of unobserved shocks that affect income and also directly affect consumption (such as sickness or vacations). We therefore instrument for a worker's change in income with their coworkers' change in income. We make the identifying assumption that employer-wide changes in pay capture shocks that are unrelated to an individual worker's endogenous labor supply decision (Koustas 2018). Although this instrument plausibly isolates exogenous shocks, one challenge in interpreting the consumption response is that changes in coworker pay may be predictable and persistent.

Second, to isolate *unpredictable, transitory* changes in income, we implement techniques from the covariance restrictions tradition in the consumption-smoothing literature. We construct a version of the instrument that is purged of predictable variation arising from seasonality and the number of paychecks per month. We further narrow the instrument to capture just the effect of transitory shocks. Specifically, we instrument for this period's income change using the change in income from one period ahead (Hall and Mishkin 1982; Blundell, Pistaferri, and Preston 2008). This instrument identifies transitory shocks when the firm component of the income process is the combination of unit root permanent shocks plus fully transitory shocks.[4] The intuition is that a transitory shock this period will revert next period, whereas a permanent shock will not. Using the period-ahead income change therefore isolates only the effect of this period's transitory shock.

Although such a procedure for isolating unpredictable transitory shocks has been widely used in the strand of the consumption-smoothing literature that relies on covariance restrictions, its identifying assumptions have seldom been validated using the standard methods used in the analysis of natural experiments. To bridge this gap we compare the pre-trends in income and consumption for workers with high values of the instrument ("treated") and low values of the instrument ("control"). We show that testing for differential pre-trends between treatment and control—the canonical test used by the modern event-study natural experiment literature to validate a parallel trends assumption—also provides a way to validate the key identifying assumption in the covariance restrictions literature. We find that the pre-trends for the two groups are similar, consistent with the view that the instrument captures income changes which are unpredictable to the worker. We further examine the post-trends in income for the treated and control groups and show that they follow the patterns one would expect under the assumed income process.

Our empirical strategy therefore connects the natural experiment approach to measuring consumption sensitivity with the covariance restriction approach. The combined approach isolates shocks that are exogenous (as in the natural experiment literature) as well as typical, unpredictable, and transitory (as in the covariance restriction literature). Furthermore, our pre- and post-trend tests show that it is possible to translate the identifying assumptions from the covariance restrictions literature into an event study framework that can be evaluated using standard techniques from the natural experiment literature.

To implement this empirical strategy we build a dataset with information on labor income, consumption, and liquid wealth using de-identified administrative bank account records from Chase. Relative to existing datasets with these variables, the three key strengths of this dataset are a large sample size of 1.3 million U.S. households, reliable measures of liquid wealth, and employer identifiers

[4] Although this simple income process is consistent with our data, we show that our estimates are also robust to assuming more flexible income processes.

which enable our main research design.

We have three reduced-form empirical findings. First, we find that month-to-month labor income fluctuations meaningfully affect consumption for the average household. In our main specification, we find an elasticity of 0.22, meaning that an unpredictable, transitory 10% increase in income in one month increases nondurable consumption in the same month by 2.2%. Translating this elasticity into a marginal propensity to consume (MPC), we find a monthly nondurables MPC of 0.10 and a quarterly nondurables MPC of 0.20.

Second, we find that households' consumption smoothing varies sharply with their liquid assets. Households with few liquid assets are an order of magnitude more sensitive than households with ample liquid assets. Moreover, this downward-sloping gradient is precisely estimated.

Third, comparing relatively predictable to relatively unpredictable income shocks for the same households, we find that responses to predictable changes in income are more muted. Although our headline estimate is an elasticity of 0.22, when we include predictable variation in pay from seasonality and from changes in the number of paychecks received per month, the elasticity falls to 0.11.

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

The model implies that temporary income volatility has a substantial welfare cost. The magnitude of the cost varies between 0.6% and 1.6% of lifetime consumption, depending on the assumed level of risk aversion and whether we incorporate the effect of income changes on consumption beyond the first month. Relative to the benchmark view in Lucas (1987) that a cost of half a percent of lifetime consumption is "large," these estimates suggest there could be a substantial welfare cost.

In addition to our empirical contribution documenting the sensitivity of consumption to typical labor income shocks, this paper makes four contributions to the theoretical literature on consumption smoothing. All four contributions are made possible by our focus on typical labor income fluctuations. This focus enables us to precisely estimate the MPC and its gradient with respect to liquidity, to evaluate the consumption effects of a frequently occurring and therefore welfare-relevant type of income variation, and to compare the effects of predictable and unpredictable income shocks for the same households.

First, our finding that MPCs fall sharply as liquidity increases confirms a central prediction of benchmark models that has been previously difficult to test. In many benchmark models of consumption behavior, the key state variable—and sometimes the only state variable—is the household's liquid assets (Aiyagari 1994; Carroll 1997; Kaplan and Violante 2014; Laibson et al. 2024). In these models, households prefer to have smooth consumption and consumption sensitivity arises only in the absence of sufficient assets. A central implication of these 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).

However, it has been difficult to test this central prediction of benchmark models using the windfall-based identification strategy in prior work.[5] The possibility of high MPCs for high-wealth

[5]See Boehm, Fize, and Jaravel (2025), Kueng (2018), and Parker et al. (2013) for examples of windfall papers that

households—informed in part by statistical imprecision regarding the sign of the gradient with liquidity—has spurred the development of several recent theoretical models which seek to explain why even households with substantial liquid wealth do not smooth their consumption (Bianchi, Ilut, and Saijo 2023; Campbell and Hercowitz 2019; Ilut and Valchev 2023; Lian 2023; Massenot and Foucault 2022; Mijakovic 2024). In this paper we find that the MPC out of typical income fluctuations is ten times larger for low-asset households than for high-asset households. Moreover, because we are able to construct our instrument for every worker in every month, this gradient is precisely estimated. Our estimates therefore provide renewed empirical support for this central prediction of benchmark models.

Second, our use of typical labor income fluctuations allows us to test and reject a prediction of consumption models based on a version of "nearly-rational" decision rules. This strand of alternatives to benchmark models treats consumption sensitivity not as a result of low-liquidity, but as a natural response to the low stakes of windfalls. This argument is based on the observation first developed in Cochrane (1989) that there is close to zero welfare cost of failing to smooth consumption in the face of small or infrequent income changes. This interpretation implies that households *would* smooth frequently-occurring income shocks, in line with Friedman's Permanent Income Hypothesis, because the cost of failing to smooth such shocks is large (Fuchs-Schündeln and Hassan 2016).

The design in this paper addresses the Cochrane (1989) critique by studying a source of frequently-occurring income fluctuations that, if not smoothed, have a substantial welfare cost: typical labor income shocks. Most working-age households get the vast majority of their income from labor, and as a consequence most income volatility comes from labor income instead of non-labor income. Yet we find that household spending is sensitive even to typical labor income volatility. The results in this paper thus favor the low-liquidity interpretation of consumption sensitivity over some versions of the near-rationality interpretation.[6]

Third, our research design allows us to resolve ambiguity in the consumption smoothing literature around the average level of the MPC. As noted in Kaplan and Violante (2010), reliable measures of the average consumption response to income shocks are central for macroeconomic models. Although we have mentioned that there is a high degree of uncertainty about the *gradient* of the MPC with respect to liquidity due to statistical imprecision, there is also uncertainty about the average *level* of the MPC arising from a more subtle source: publication bias. A recent meta-analysis by Havranek and Sokolova (2020) finds that published MPC estimates may dramatically overstate the true MPC after correcting for publication bias. Such concerns may be muted for our design because the standard error is less than one cent, and therefore even economically insignificant consumption responses will be statistically significant at conventional levels.

The paper's fourth contribution to the theoretical consumption-smoothing literature is to directly

provide seminal evidence of high average sensitivity but have difficulty distinguishing between a positive or a negative gradient with respect to liquidity. One windfall paper with evidence of a gradient in at least part of the liquid asset distribution is Fagereng, Holm, and Natvik (2021). That paper estimates MPCs using lotteries, ranks households by quartiles of liquid assets, and finds that the bottom quartile has a higher MPC than quartiles two, three, and four, though it is unable to distinguish between a positive or negative gradient among the top three quartiles. The paper interprets the high average MPC to lotteries through the lens of a luxury bequest motive. In work subsequent to ours, Baker et al. (2023) documents an MPC gradient out of pandemic stimulus checks. For evidence on the liquidity gradient outside of natural experiments or outside of transitory shocks, see Baker (2018), Blundell, Pistaferri, and Preston (2008), Crawley and Kuchler (2023), Gelman (2021), Graham and McDowall (2025), and Souleles (1999).

[6] Our results are consistent with models such as Kaplan and Violante (2014) and Laibson et al. (2024) in which low liquidity arising in part from negative income shocks *causes* high sensitivity. They are also consistent with models where an alternative channel that is correlated with wealth (e.g., inter-household transfers in Chiteji and Hamilton 2002, or permanent heterogeneity as in Aguiar, Bils, and Boar 2021, Epper et al. 2020, or Ganong et al. 2024) contributes to both low liquidity and high MPCs.

test and confirm the hypothesis that households use advance information about income changes to smooth their consumption. The classic excess sensitivity test based on the Euler equation dating back to Hall (1978) is that households' consumption should not respond to predictable income shocks (whereas households should respond to unpredictable income shocks). Indeed, the Jappelli and Pistaferri (2010) review article separates predictable shocks from unpredictable shocks as *the key dividing line* in the empirical consumption literature. However, it has been difficult to directly assess the hypothesis that the consumption response is larger when a shock is less predictable because it requires comparisons across distinct natural experiments that differ both in terms of predictability and on many other dimensions as well. Because our methodology captures both predictable and unpredictable variation in labor income for the same households, it offers more direct evidence on the role of advance information. We find larger responses to less predictable shocks than to more predictable ones, consistent with this classic conjecture of consumption models.[7]

The paper proceeds as follows. Section 2 describes our research design and identification strategy. Section 3 describes the data. Section 4 describes our empirical results. Section 5 quantifies the welfare cost of income volatility using a simple model. Section 6 concludes.

2 Research Design

Our aim is to estimate the sensitivity of consumption to transitory, unpredictable, and exogenous shocks to typical labor income. Our empirical strategy can be summarized as using two-stage least squares:

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

$$\Delta y_{it} = \phi - \rho \Delta y_{j(-i,t),t+1} + \nu_{it}. \tag{2}$$

where i indexes households, t indexes time (months), c is the log of monthly nondurable consumption, y is the log of monthly labor income, and j indexes firms.

The crucial step in our empirical strategy is the first stage in equation (2). In this step, we instrument for a household's income change this period with $-\Delta y_{j(-i,t),t+1}$: the negative *period-ahead* income change of their *coworkers*.[8] This approach builds on two distinct traditions in the consumption-smoothing literature. The first tradition uses covariance restrictions to isolate transitory, unpredictable shocks. This is reflected in our use of the period-ahead income change. The second tradition uses quasi-experiments to isolate exogenous shocks. This is reflected in our use of a coworker instrument.

The rest of this section details this empirical approach. We begin by defining the underlying object of interest, and address some of the idiosyncrasies of applying our method in a monthly setting. We then discuss the ideal experiment together with alternative empirical approaches that try to capture

[7] We do not literally find zero response to predictable income shocks. The starkest form of this prediction is that no variables which are known at $t-1$ can predict the change in consumption from $t-1$ to t (Hall 1978). However, Jappelli and Pistaferri (2010) notes some ways that the starkest form of the prediction might fail without providing a compelling rejection of the Euler equation. These include liquidity constraints and non-separable preferences over time use and consumption (Attanasio and Weber 1995). When looking at the consumption effects of more predictable income shocks, we find a strong gradient with respect to liquidity, suggesting that liquidity constraints are one relevant channel in our setting.

[8] As we show below, the change is multiplied by negative one because the income growth in period $t+1$ is inversely related to transitory shocks in period t.

the ideal experiment. Next, we present diagnostic tests to assess the identifying assumptions of our preferred approach. Finally, we address the endogeneity of labor income by using pay changes for coworkers.

2.1 Environment and Object of Interest

A traditional way to think about the consumption response to income shocks is through the lens of the classic permanent-transitory income process

$$\begin{aligned} y_{it} &= z_{it} + \varepsilon_{it} \\ z_{it} &= z_{i,t-1} + \eta_{i,t} \\ \Rightarrow \Delta y_{it} \equiv y_{i,t} - y_{i,t-1} &= \eta_{i,t} + \Delta\varepsilon_{it} = \eta_{i,t} + \varepsilon_{it} - \varepsilon_{i,t-1}. \end{aligned} \tag{3}$$

Under this income process, log income in each period consists of two independently and identically distributed (iid) mean-zero shocks ε_{it} and η_{it} with nonzero variances σ_ε^2 and σ_η^2. ε_{it} is transitory while η_{it} is permanent. These assumptions have several testable predictions, which we discuss and then test below. We also consider alternative income processes with a more general auto-regressive component in Appendix B.3, and show that our results are robust. Moreover, given the concerns raised in Commault (2022), we confirm in Appendix B.4 that our data are consistent with an assumption of no serial correlation in the transitory component of income. It is important to note that in our preferred specification, described in Section 2.5, we instrument for income using variation in coworker pay, which isolates firm-wide shocks to income. In our final specification we therefore only require that the structural assumptions above apply to the firm component of the income process. The worker-specific income process can in principle be more general.

Next, we specify a consumption function:

$$\Delta c_{it} = \beta_\eta \eta_{it} + \beta_\varepsilon \varepsilon_{it} + \zeta_{it}, \tag{4}$$

allowing for different responses to permanent and transitory shocks to income. This consumption function can be derived from a life-cycle model of consumption and savings with constant relative risk aversion preferences (Blundell, Low, and Preston 2013). In this paper, the key object of interest is β_ε, the consumption response to a transitory income shock.

Income growth and consumption growth in equations (3) and (4) are best understood as changes *net of predictable movements*, thus explaining the absence of any intercepts. In a classic implementation of the covariance restrictions approach, Blundell, Pistaferri, and Preston (2008) study annual earnings, where predictable movements are best explained by time trends and age profiles. They therefore residualize their variables from these trends.

We implement a similar approach to purge predictable variation in our setting with monthly earnings. At the monthly frequency there are two important sources of predictable movements in income. First, there is firm-specific, recurring annual seasonality. This could be driven, for example, by annual bonuses or lulls in operations that occur at the same time every year in a firm and that, importantly, differ in timing across different firms. We account for this seasonality by creating a linear predictor of monthly income growth based on income growth from 12 months prior. We calibrate this linear predictor using a simple optimization routine. The remaining residual no longer features a correlation between income growth today and that of 12 months ago. The second important source of predictable variation is driven by the number of paychecks within a calendar month. Workers who are paid fortnightly, for example, receive two paychecks in most months but receive an extra third paycheck in some months, a phenomenon that is knowable in advance. To address this second predictable component, we specify our instrument for income in terms of pay per paycheck, which abstracts from the number of paychecks. Additional detail on these adjustment procedures is described in Appendix B.

In the next two sections, we assume that equations (3) and (4) hold and introduce additional assumptions that are sufficient to identify the consumption response to a transitory shock β_ε. We also discuss testable predictions that enable us to assess if the equations are an accurate summary of the underlying income process.

2.2 Estimators: Ideal Experiment and Prior Literature

The ideal experiment for identifying the consumption response to a transitory shock is a randomized controlled trial (RCT), which we refer to as the "β_ε RCT" because it identifies the β_ε parameter in equation (4). A sample of households are randomly assigned to a treatment group ($D_i^\varepsilon = 1$) or control group ($D_i^\varepsilon = 0$). Members of the treatment group receive a one-time boost to log income of μ_ε in period t^*. That is, experimental income, $y_{it}^{RCT}(D_i^\varepsilon)$ is

$$y_{it}^{RCT}(D_i^\varepsilon) = \begin{cases} y_{it} + D_i^\varepsilon \mu_\varepsilon & \text{if } t = t^* \\ y_{it} & \text{otherwise} \end{cases} \tag{5}$$

where y_{it} is log income for each worker, as above. By construction, the income shock in the experiment is orthogonal to the rest of the income history. Let consumption still be characterized by equation (4). Importantly, the treatment transfer is *unannounced* prior to period t^*. It follows immediately that using the treatment indicator as an instrument for income in period t^* identifies the parameter β_ε.

Result RCT (β_ε RCT)

Under random assignment of a shock to log income in period t^*, we are able to identify β_ε using the following ratio:

$$\hat{\beta}_{\varepsilon,RCT} \equiv \frac{\mathbb{E}\left(c_{i,t^*} \,|D_i^\varepsilon = 1\right) - \mathbb{E}\left(c_{i,t^*} \,|D_i^\varepsilon = 0\right)}{\mathbb{E}\left(y_{i,t^*}^{RCT} \,|D_i^\varepsilon = 1\right) - \mathbb{E}\left(y_{i,t^*}^{RCT} \,|D_i^\varepsilon = 0\right)} = \beta_\varepsilon$$

where y_{it}^{RCT} is log income, inclusive of the treatment payment.

This proof and all others below are provided in Appendix C.

In contrast to the ideal experiment, an ordinary least squares regression, $\Delta c_{it} = \alpha + \beta_{OLS} \Delta y_{it} + \zeta_{it}$, will not identify any of the structural parameters of interest because of three distinct problems. First, Δy_{it} mixes permanent and transitory shocks to income. Benchmark theories of consumption predict very different responses to permanent and transitory shocks, and so an estimate that mixes the effect of both types of shocks will be hard to interpret (Jappelli and Pistaferri 2010). Second, Δy_{it} includes both new transitory shocks (ε_{it}) as well as predictable reversions of past transitory shocks ($\varepsilon_{i,t-1}$);

according to economic theory, the former should affect consumption but the latter should not.[9] Third, Δy_{it} is a combination of exogenous shocks, such as a change in shift scheduling or a demand shock to one's firm, and endogenous variation in labor supply, such as a pre-planned vacation. In the latter case, $\hat{\beta}_{OLS}$ will include jointly determined movements in consumption and income by the household or worker. For example, $\hat{\beta}_{OLS}$ will be biased downward if the worker takes time away from paid work during a period of higher-than-normal consumption (such as a vacation or a health emergency) or biased upward if a worker increases their labor supply to finance a specific consumption need. We will address each of these challenges to estimation in turn below.

These problems have motivated two approaches in the prior literature. First, one common way to estimate the consumption response to transitory income is via a natural experiment with windfall income. These are typically analyzed using the same assumptions—that windfalls are orthogonal to the rest of income and unanticipated by recipient—and the same estimator described above for the "β_ε RCT". As discussed in the introduction, analyses of windfalls have been fertile ground for consumption researchers, but the ability to draw conclusions for consumption-savings models has been limited by the fact that such windfalls are inherently infrequent. They are not the type of shocks that households usually experience and statistical precision is often limited by small sample sizes.

The second common approach to measuring consumption smoothing assumes a set of orthogonality conditions that allow for identification of β_ε. The following assumptions are sufficient for such identification:

Assumption IID (IID Income Shocks)

The shocks η_{it} and ε_{it} are independent and identically distributed (iid) at the (i, t) level:

$$(\eta_{it}, \varepsilon_{it}) \sim IID$$

and are further independent of each other

$$\eta_{i,t} \perp \varepsilon_{it} \quad \forall t$$

The key implication of this assumption is that the shocks are serially uncorrelated, and are therefore unpredictable given the values of past shocks. Although our assumptions are formulated as orthogonality conditions, we follow the prior literature (e.g., Jappelli and Pistaferri 2010) in calling this the "covariance restriction" approach. This is because covariance restrictions are sufficient to identify consumption sensitivity to temporary and permanent income shocks.[10] We make the stronger assumptions described here because they yield additional empirical implications that can be inspected to help validate the model.

In addition, the approach requires the following:

Assumption CE (Consumption Exogeneity)

$$\zeta_{i,t} \perp (\eta_{i,t+s}, \varepsilon_{i,t+s}) \quad \forall s$$

This final assumption embeds two important implications. First, it imposes the restriction that households do not anticipate the shocks to income, at least not in any way that affects consumption.

[9] See Appendix D.

[10] Specifically, identification is possible under a weaker set of assumptions than Assumption IID, namely that $\text{cov}(\Delta c_{it}, \eta_{i,t+1}) = \text{cov}(\Delta c_{it}, \varepsilon_{i,t+1}) = 0$. Kaplan and Violante (2010) refers to these as "No Foresight" assumptions.

While we can test whether the income shocks are "predictable" by the econometrician, we can only assume that households do not have private information about future shocks. Second, it implies that unobservable determinants of consumption are not correlated with contemporaneous shocks to income. In other words, it assumes that all observed income changes (net of predictable changes) are exogenous. This assumption would fail if, for example, a worker took a vacation or had a health shock that changed both their labor income and their consumption. In Section 2.5 below, we show how our use of an instrument allows us to achieve identification under a weaker exogeneity assumption. For the rest of Section 2.2, however, we maintain this stronger assumption.

Under these orthogonality conditions, we can show the following identification result, previously derived in Hall and Mishkin (1982) and Blundell, Pistaferri, and Preston (2008), which we refer to henceforth as "HM-BPP":[11]

Result HM-BPP (Identification of β_ε)

If Assumptions IID and CE hold, the ratio of the covariance between consumption growth and the (negative) one-period income growth lead, and the covariance between income growth and the (negative) one-period income growth lead, is equal to β_ε (Hall and Mishkin 1982; Blundell, Pistaferri, and Preston 2008):

$$\hat{\beta}_{\varepsilon,HM-BPP} \equiv \frac{\operatorname{cov}\left(\Delta c_{it}, -\Delta y_{i,t+1}\right)}{\operatorname{cov}\left(\Delta y_{it}, -\Delta y_{i,t+1}\right)} = \frac{\operatorname{cov}\left(\Delta c_{it}, \varepsilon_{it}\right)}{\operatorname{var}\left(\varepsilon_{it}\right)} = \frac{\beta_\varepsilon \sigma_\varepsilon^2}{\sigma_\varepsilon^2} = \beta_\varepsilon$$

The term $\operatorname{cov}\left(\Delta c_{it}, -\Delta y_{i,t+1}\right) / \operatorname{cov}\left(\Delta y_{it}, -\Delta y_{i,t+1}\right)$ in Result HM-BPP can be interpreted as an instrumental variables (IV) regression of Δc_{it} on Δy_{it}, instrumenting for Δy_{it} with the negative of its one-period lead, which in turn is equivalent to the coefficient from an OLS regression of Δc_{it} on ε_{it}. We provide intuition behind the use of $-\Delta y_{i,t+1}$ as the instrument in the next section.

2.3 An Event Study Approach to Assessing the Credibility of Covariance Restrictions

The HM-BPP result requires strong assumptions—namely that income consists of a permanent unit root component and a fully transitory component, that these permanent and transitory shocks are IID and therefore unpredictable, and that these shocks are exogenous with respect to consumption.

We now show that it is possible to assess the validity of these assumptions using tests that are very similar to those used in the quasi-experimental literature to assess the identification assumptions in event study research designs. Intuitively, if the instrument for income is predictable to the econometrician then we should see systematic differential pre-trends for income. Even if it is unpredictable to the econometrician, if it is anticipated by the household then we might see differential systematic pre-trends for consumption. And finally, if the income process really is the combination of fully transitory shocks and unit root permanent shocks then this will have implications for post-trends of income.

As a means of building intuition, we present results for a discretized version of the HM-BPP instrument. We focus on a particular time period, t^*, as our "event time" and assign workers to "treated" and "control" groups based on their change in income between periods t^* and t^*+1:

$$D_{i,t^*} \equiv \mathbf{1}(-\Delta y_{i,t^*+1} > M_{-\Delta y}) \tag{6}$$

[11] In Appendix C, we derive this result in more detail.

where $M_{-\Delta y} \equiv median(-\Delta y_{i,t^*+1})$. We label the group with the steeper income decrease from t^* to t^*+1 as treated because the transitory shock in period t^* is negatively correlated with the change in income from period t^* to t^*+1. Separating workers into treatment and control groups allows us to use difference-in-differences and event study methods as a bridge between covariance restrictions and a quasi-experiment. This exercise is in the spirit of Lamadon, Mogstad, and Setzler (2022), who use a difference-in-difference framework to illuminate their structural identification strategy.[12] Finally, in order to map our predictions to common event study approaches, we normalize log income relative to a pre-event period:

$$\tilde{y}_{i,t} \equiv y_{i,t} - y_{i,t^*+r}, \quad r < 0 \tag{7}$$

Under the assumptions from HM-BPP, five testable predictions follow for the binary instrument. Appendix E reports analogous results for the continuous case.

First, income in the treated group is higher than in the control group at t^*:

Testable Prediction 1 (First Stage: Positive Income Difference at t^*)

If Assumption IID holds, the average difference in (normalized) log income between the treated and control group in period t^* is equal to the average difference in the transitory shock in period t^*, and is positive:

$$\mathbb{E}\left(\tilde{y}_{i,t^*} \middle| D_{i,t^*}=1\right) - \mathbb{E}\left(\tilde{y}_{i,t^*} \middle| D_{i,t^*}=0\right) = \mathbb{E}\left(\varepsilon_{i,t^*} \middle| D_{i,t^*}=1\right) - \mathbb{E}\left(\varepsilon_{i,t^*} \middle| D_{i,t^*}=0\right) > 0$$

This difference, which can be thought of as a first stage regression, shows that the instrument isolates variation in the transitory shock in period t^*.

Second, there are parallel trends in pre-event income between the treated and control groups:

Testable Prediction 2 (Parallel Pre-event Income Trends)

If Assumption IID holds, the average difference in (normalized) log income between the treated and control group during periods before t^* is zero, i.e. there are parallel trends:

$$\mathbb{E}\left(\tilde{y}_{i,t^*+k} \middle| D_{i,t^*}=1\right) - \mathbb{E}\left(\tilde{y}_{i,t^*+k} \middle| D_{i,t^*}=0\right) = 0 \text{ for } k < 0$$

This test is reminiscent of the test in Blundell, Pistaferri, and Preston (2008) footnote 20 for whether the first difference of income covaries with lags of the first difference in income. Our test differs in that we use normalized log income instead of first differences and we make use of a discrete version of the instrument for better comparability with typical event study diagnostics.

Although our focus is on the response of consumption to the transitory shock in period t^*, the model also produces predictions about the path of income beyond the event period:

Testable Prediction 3 (Negative Income Difference at t^*+1)

If Assumption IID holds, the average difference in (normalized) log income between the treated and control group in period t^*+1 is equal to the average difference in the permanent and transitory

[12] They impose covariance restrictions on the process for firm value-added shocks and use an event study framework to show that the shocks to value-added appear to be *permanent*. We use a similar strategy to show that shocks to worker income appear to be *transitory*.

shocks in period t^*+1, and both average differences are negative.

$$\mathbb{E}\left(\tilde{y}_{i,t^*+1}\left|D_{i,t^*}=1\right.\right)-\mathbb{E}\left(\tilde{y}_{i,t^*+1}\left|D_{i,t^*}=0\right.\right)=\underbrace{\mathbb{E}\left(\eta_{i,t^*+1}\left|D_{i,t^*}=1\right.\right)-\mathbb{E}\left(\eta_{i,t^*+1}\left|D_{i,t^*}=0\right.\right)}_{<0}$$
$$+\underbrace{\mathbb{E}\left(\varepsilon_{i,t^*+1}\left|D_{i,t^*}=1\right.\right)-\mathbb{E}\left(\varepsilon_{i,t^*+1}\left|D_{i,t^*}=0\right.\right)}_{<0}$$
$$<0$$

Testable Prediction 4 (Negative Income Difference at t^*+2 Onward)

If Assumption IID holds, the difference in (normalized) log income between the treated and control group in periods $t>t^*+1$ is equal to the average difference in the permanent shock in period t^*+1, and this difference is negative.

$$\mathbb{E}\left(\tilde{y}_{i,t^*+k}\left|D_{i,t^*}=1\right.\right)-\mathbb{E}\left(\tilde{y}_{i,t^*+k}\left|D_{i,t^*}=0\right.\right)=\mathbb{E}\left(\eta_{i,t^*+1}\left|D_{i,t^*}=1\right.\right)-\mathbb{E}\left(\eta_{i,t^*+1}\left|D_{i,t^*}=0\right.\right)$$
$$<0 \quad \text{for } k>1$$

We do not make use of these post-period patterns to identify β_ε. Nonetheless, they produce additional moments that allow us to further assess the plausibility of our structural assumptions and the nature of permanent and transitory shocks in this model. Put another way, these predictions show how we might deconvolute our instrument into three distinct shocks. For example, while there are other moments of the data that can be used to characterize the nature of permanent shocks (see Blundell, Pistaferri, and Preston 2008), Prediction 4 points to clear visual evidence that can be recovered with standard event study techniques. We report the empirical results for each of these tests in Section 4.2.

As in the case of a difference-in-differences estimator or event study, these four testable predictions show that the discretized version of our instrument implies patterns that can be visually inspected using panel data. Figure 1a presents a stylized event study plot of the difference $\mathbb{E}\left(\tilde{y}_{i,t^*+k}\left|D_{i,t^*}=1\right.\right)-\mathbb{E}\left(\tilde{y}_{i,t^*+k}\left|D_{i,t^*}=0\right.\right)$ for different values of k for the income process in equation (3). Figure 1a illustrates a sawtooth pattern for the difference in income between our treatment and control groups. This reflects the patterns described in the testable predictions.

To understand why these sawtooth income dynamics are predicted under the assumed income process, it is useful to revisit the definition of the instrument $-\Delta y_{i,t^*+1}$. Equation (3) shows that this instrument can be decomposed as:

$$-\Delta y_{i,t^*+1}=-\eta_{i,t^*+1}-\varepsilon_{i,t^*+1}+\varepsilon_{it^*}.$$

In this environment, there are three types of shocks that can lead to an above-median value of $-\Delta y_{i,t^*+1}$ (i.e. to a drop in income from t^* to t^*+1) :

1. A positive temporary income shock at t^* (ε_{i,t^*})

2. A negative temporary income shock at t^*+1 (ε_{i,t^*+1})

3. A negative permanent income shock at t^*+1 (η_{i,t^*+1})

A positive temporary income shock at t^* (shock 1) will cause income to rise at t^*, consistent with

Figure 1: Stylized Patterns

(a) Event Study Plot of $\mathbb{E}\left(\tilde{y}_{i,t}\,\middle|D_{i,t^*}=1\right)-\mathbb{E}\left(\tilde{y}_{i,t}\,\middle|D_{i,t^*}=0\right)$

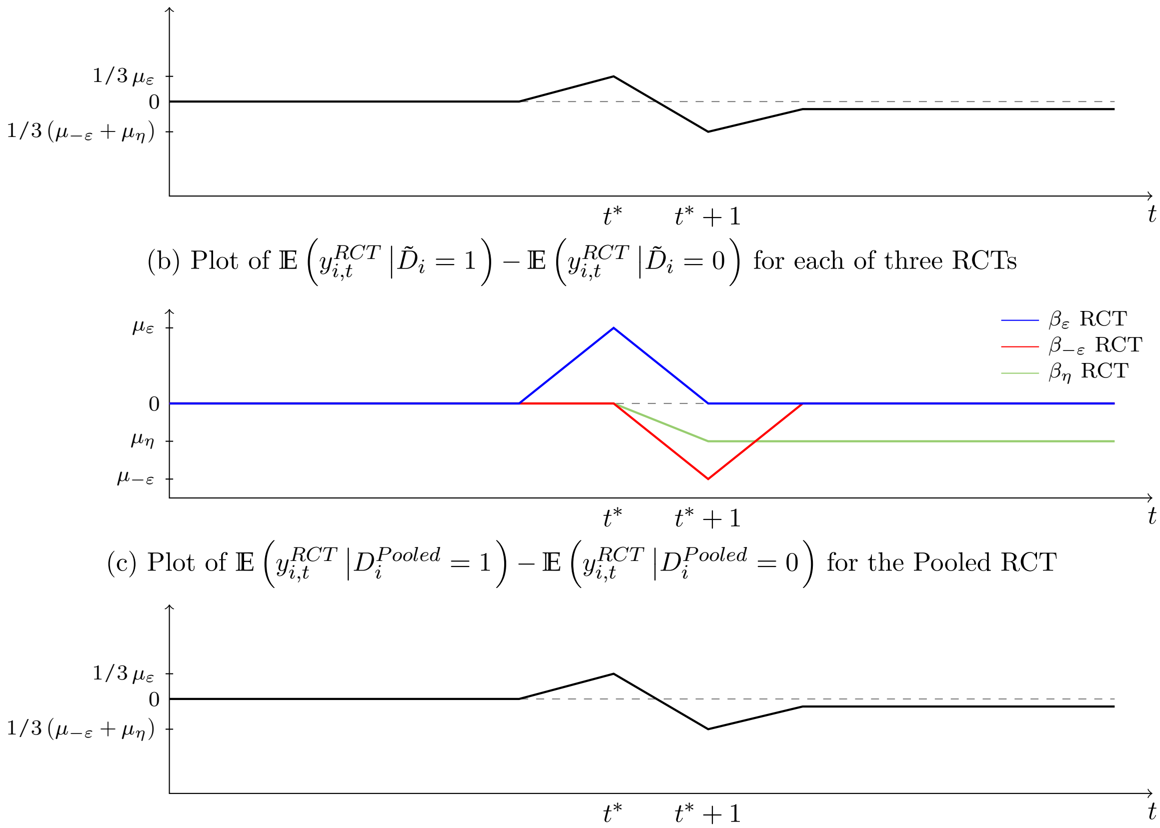

Notes: This figure shows hypothetical event study patterns for the binary version of our instrument for temporary income changes and how they facilitate a deconvolution of three distinct structural shocks to income. Panel (a) shows pre- and post-event patterns for the difference between average, normalized income $\tilde{y}_{i,t}$ in a group where $-\Delta y_{i,t^*+1}$ is above-median and a group where $-\Delta y_{i,t^*+1}$ is below-median. Panel (b) shows the difference in average income for three hypothetical RCTs described in Section 2.4: β_ε increases income at t^* (the experiment of interest) while two nuisance experiments, $\beta_{-\varepsilon}$ and β_η, decrease income at t^*+1 and at every period from t^*+1 forward, respectively. Panel (c) shows that an equal-weighted sum of income in the three RCTs of panel (b), i.e. the Pooled RCT, is equal to the path of income in panel (a), i.e. the path associated with our discrete instrument.

Testable Prediction 1. Negative temporary or permanent income shocks at t^*+1 (shocks 2 and 3) will cause income to fall at t^*+1, consistent with Testable Prediction 3. And a negative permanent income shock at t^*+1 (shock 3) will cause income to remain depressed beyond t^*+1, consistent with Testable Prediction 4. Because each of these three shocks is iid, they are unrelated to income prior to t^*, so there should be no pre-event trends in income, consistent with Testable Prediction 2. Crucially for identification, out of these three shocks, only the first one, ε_{i,t^*}, is realized and apparent as of t^*. The other two shocks are as of yet unanticipated. The causal impact of just this transitory shock can then be isolated by measuring the response of consumption at t^*.

These patterns also help clarify why the period-ahead income change serves as a useful instrument for isolating this period's transitory shocks. Two kinds of shocks can occur in any given period t^*. A transitory shock at t^* will revert at t^*+1. The resulting income change from t^* to t^*+1 is picked up by the instrument (i.e. $-\Delta y_{i,t^*+1}\propto\varepsilon_{i,t^*}$). But a permanent shock at t^* will not revert at t^*+1.

Because a contemporaneous permanent shock does not cause income to change from t^* to t^*+1, this type of shock is effectively excluded from the instrument. Hence the only contemporaneous shock isolated by the instrument is the transitory shock.

Although the first four testable predictions implied by the HM-BPP assumptions concerned the path of income, a fifth and final testable prediction concerns consumption. In particular, the assumptions imply parallel pre-trends for normalized consumption between the treated and control groups[13]:

Testable Prediction 5 (Parallel Pre-event Consumption Trends)

If Assumptions IID and CE hold, the average difference in (normalized) log consumption between the treated and control group is zero during periods $t < t^*$, i.e. there are parallel trends:

$$\mathbb{E}\left(\tilde{c}_{i,t^*+k} \,\middle|\, D_{i,t^*}=1\right) - \mathbb{E}\left(\tilde{c}_{i,t^*+k} \,\middle|\, D_{i,t^*}=0\right) = 0 \textit{ for } k < 0$$

This test is similar to the test in the footer of Blundell, Pistaferri, and Preston (2008) Table 5 for whether the first difference of consumption covaries with the leads of the first difference in income. The intuition for the test is that because the shocks picked up by the instrument are unpredictable prior to t^*, there should not be any consumption pre-trends. Testable Prediction 5 therefore provides a way to assess whether the household acts as if the income shock really is a surprise.

Overall, this analysis demonstrates that event-study graphs of income and consumption can be used to test the assumptions underlying the covariance-restriction methodology and also to visually assess the magnitude of consumption sensitivity to a transitory shock. In the next section, to clarify how this design captures the same parameter as the ideal experiment, we show that it is possible to rewrite HM-BPP as the product of a set of experiments.

2.4 Intuition for Covariance Restrictions via a Mixture of RCTs

The idea that a researcher can capture temporary income shocks at time t using $-\Delta y_{i,t+1}$ as an instrument may be counter-intuitive to some readers. To build intuition for this result, we will show a chain of equivalences between the ideal experiment and HM-BPP. Readers already familiar with the identification approach in HM-BPP may skip this section.

Within the context of the environment described in Section 2.1, we now assume that three separate experiments have occurred. The first experiment is the original treatment and control sample from the β_ε RCT in Section 2.2, with treatment now denoted as $\tilde{D}_i^\varepsilon$. We also assume that two additional experiments have occurred which we will refer to as the "$\beta_{-\varepsilon}$ RCT" and the "β_η RCT." In each experiment, assignment to treatment is not announced ahead of time.

Each RCT captures one of the three structural shocks that collectively constitute $\Delta y_{i,t+1}$. We assume, without loss of generality, that each RCT features the same number of treatment and control group members.[14]

In the second experiment, a different sample is randomly assigned to a treatment or control group, indexed by $\tilde{D}_i^{-\varepsilon} \in \{0,1\}$. Members of the treatment group receive a one-time *negative* transfer in period t^*+1 of $\mu_{-\varepsilon} < 0$:[15]

[13] Normalized consumption is defined in a similar way to income, $\tilde{c}_{i,t} \equiv c_{i,t} - c_{i,t^*+r}, \quad r < 0$.

[14] In general, for a given set of shares in each RCT, there exists appropriately-scaled transfers that preserve our results.

[15] In practice, it is hard to imagine an RCT where money is taken away from the treatment group. The study could alternatively give a cash transfer to both treatment and control groups in every period, and reduce the transfer for the treatment group, only in period t^*+1.

$$y_{it}^{RCT}(\tilde{D}_i^{-\varepsilon}) = \begin{cases} y_{it} + \tilde{D}_i^{-\varepsilon}\mu_{-\varepsilon} & \text{if } t = t^* + 1 \\ y_{it} & \text{otherwise} \end{cases} \tag{8}$$

Additionally, there is a third experiment, the "β_η RCT." In this group, treatment and control are indexed by $\tilde{D}_i^\eta$, and the treatment group receives a *negative* transfer of $\mu_\eta < 0$ in period $t^* + 1$ and every period thereafter:

$$y_{it}^{RCT}(\tilde{D}_i^\eta) = \begin{cases} y_{it} & \text{for } t \leq t^* \\ y_{it} + \tilde{D}_i^\eta \mu_\eta & \text{for } t \geq t^* + 1 \end{cases} \tag{9}$$

Figure 1b plots the difference in average log income between the treatment and control groups for each of the three RCTs.

Finally, for pedagogical purposes, we add an absurd assumption: the $\tilde{D}$ labels for who was in each RCT have been lost! Instead, the econometrician only has a dataset that is pooled across the three experiments with a single treatment indicator:[16]

$$D_i^{\text{Pooled}} = \begin{cases} 1 & \text{if } \tilde{D}_i^\varepsilon = 1, \text{or } \tilde{D}_i^{-\varepsilon} = 1, \text{or } \tilde{D}_i^\eta = 1 \\ 0 & \text{if } \tilde{D}_i^\varepsilon = 0, \text{or } \tilde{D}_i^{-\varepsilon} = 0, \text{or } \tilde{D}_i^\eta = 0 \end{cases} \tag{10}$$

In this (admittedly unrealistic) environment, we describe a chain of equivalences between the ideal experiment of Result RCT and the estimator in Result HM-BPP. The first link in this chain is that the ideal experiment has a more complicated counterpart, namely the pooled experiment, which adds two new experiments, but still identifies the same parameter. The second link in the chain is that this more complicated, pooled experiment generates the same variation in income as the binary version of HM-BPP. Finally, the binary version of HM-BPP is a simplified version of the classic HM-BPP estimator.

The first link in the chain is that with this pooled experiment, one can still recover β_ε:

Result Pooled (Parameter Equivalence of Pooled RCT and Simple RCT)

Using the Pooled RCT, we are able to identify β_ε using the following ratio:

$$\begin{aligned} \hat{\beta}_{\varepsilon,Pooled} &\equiv \frac{\mathbb{E}\left(c_{i,t^*} \,\middle|\, D_i^{\text{Pooled}} = 1\right) - \mathbb{E}\left(c_{i,t^*} \,\middle|\, D_i^{\text{Pooled}} = 0\right)}{\mathbb{E}\left(y_{i,t^*}^{RCT} \,\middle|\, D_i^{\text{Pooled}} = 1\right) - \mathbb{E}\left(y_{i,t^*}^{RCT} \,\middle|\, D_i^{\text{Pooled}} = 0\right)} \\ &= \frac{\frac{1}{3}\left[\mathbb{E}\left(c_{i,t^*} \,\middle|\, \tilde{D}_i^\varepsilon = 1\right) - \mathbb{E}\left(c_{i,t^*} \,\middle|\, \tilde{D}_i^\varepsilon = 0\right)\right]}{\frac{1}{3}\left[\mathbb{E}\left(y_{i,t^*}^{RCT} \,\middle|\, \tilde{D}_i^\varepsilon = 1\right) - \mathbb{E}\left(y_{i,t^*}^{RCT} \,\middle|\, \tilde{D}_i^\varepsilon = 0\right)\right]} \\ &= \beta_\varepsilon \end{aligned}$$

where y_{it}^{RCT} is log income, inclusive of the treatment payment.

The key step in the proof is going from the observed experiment in line one to the unobserved experiment in line two. At time t^*, the only difference in income known to members of *any* of the

[16] To account for the fact that the three RCT samples are mutually exclusive, $\tilde{D}_i^\varepsilon = 1$ will be used below as shorthand for (1) originally being a member of the "β_ε RCT" and (2) being in the treatment group, and $\tilde{D}_i^\varepsilon = 0$ will likewise identify those analagous control group members. We can think of there being a third value, $\tilde{D}_i^\varepsilon = \emptyset$ for those who were not originally members of the "β_ε RCT." We use a similar convention for the other two RCTs. Everyone in the pooled sample will have exactly one non-empty indicator.

treatment or control groups is the transfer to members of the β_ε RCT. Thus, the only difference in consumption patterns in that period is driven by that subgroup. The subsequent differences in income for the other two subgroups are not known or anticipated prior to period t^*+1. The fact that the behavioral differences are driven by a subset in period t^*, a third to be exact, cancels out in the numerator and denominator of the estimator. We therefore recover the estimate as if we had a sample featuring only the ideal RCT group.

The second link in the chain is that this pooled RCT has the same dynamic pattern of income as the binary HM-BPP instrument:

Result Pooled RCT-HMBPP (Income Equivalence of Pooled RCT and HM-BPP)

If Assumption IID holds, the average difference in (normalized) log income based on the binary instrument, $D_{i,t^*} \equiv \mathbf{1}(-\Delta y_{i,t^*+1} > M_{-\Delta y})$, is equivalent to the average difference in log income between the treatment and control groups of the pooled experiment, in all time periods:

$$\mathbb{E}\left(\tilde{y}_{i,t}\left|D_{i,t^*}=1\right.\right)-\mathbb{E}\left(\tilde{y}_{i,t}\left|D_{i,t^*}=0\right.\right)=\mathbb{E}\left(y_{i,t}^{RCT}\left|D_i^{\text{Pooled}}=1\right.\right)-\mathbb{E}\left(y_{i,t}^{RCT}\left|D_i^{\text{Pooled}}=0\right.\right)\forall t$$

where the treatments in each of the three sub-RCTs are defined as:

$$\begin{aligned}
\mu_\varepsilon &\equiv 3\left(\mathbb{E}\left(\varepsilon_{i,t^*}\left|D_{i,t^*}=1\right.\right)-\mathbb{E}\left(\varepsilon_{i,t^*}\left|D_{i,t^*}=0\right.\right)\right)\\
\mu_{-\varepsilon} &\equiv 3\left(\mathbb{E}\left(\varepsilon_{i,t^*+1}\left|D_{i,t^*}=1\right.\right)-\mathbb{E}\left(\varepsilon_{i,t^*+1}\left|D_{i,t^*}=0\right.\right)\right)\\
\mu_\eta &\equiv 3\left(\mathbb{E}\left(\eta_{i,t^*+1}\left|D_{i,t^*}=1\right.\right)-\mathbb{E}\left(\eta_{i,t^*+1}\left|D_{i,t^*}=0\right.\right)\right)
\end{aligned}$$

and y_{it}^{RCT} is log income, inclusive of the treatment payment.

This result can also be seen visually in Figure 1c. The figure shows the difference between income in the treatment and control groups for the pooled RCT. It is the exact same pattern as Figure 1a. This is because the change in income captured by $-\Delta y_{i,t^*+1}$ can be thought of as simply being the weighted sum of the three hypothetical RCTs in Figure 1b.

Finally, the binary instrument $D_{i,t^*} \equiv \mathbf{1}(-\Delta y_{i,t^*+1} > M_{-\Delta y})$ is simply a discretized version of the continuous instrument, $-\Delta y_{i,t^*+1}$. Both the binary and continuous versions of the instrument identify β_ε:

Result Wald (Identification of β_ε with Wald Estimator)

If Assumptions IID and CE hold, the ratio of the difference in (normalized) log consumption between the treated and control group and the difference in (normalized) log income between the treated and control group in period t^* is equal to β_ε:

$$\hat{\beta}_{\varepsilon,\text{WALD}} \equiv \frac{\mathbb{E}\left(\tilde{c}_{i,t}\left|D_{i,t^*}=1\right.\right)-\mathbb{E}\left(\tilde{c}_{i,t}\left|D_{i,t^*}=0\right.\right)}{\mathbb{E}\left(\tilde{y}_{i,t}\left|D_{i,t^*}=1\right.\right)-\mathbb{E}\left(\tilde{y}_{i,t}\left|D_{i,t^*}=0\right.\right)}=\beta_\varepsilon$$

The binary version is useful conceptually because it clarifies the connection to an experiment but comes at the cost of throwing away identifying variation in the data. The path of income in Figure 1a describes equally well the difference between treatment and control for the binary case as it does for covariances with a continuous instrument (see Appendix E for more details).

This chain of equivalences demonstrates that under the structural assumptions of HM-BPP, this methodology recovers the same parameter as the ideal experiment. Moreover, these structural assumptions can be evaluated in an event study difference-in-differences framework similar to the frameworks used in the quasi-experimental literature.

2.5 Isolating Plausibly Exogenous Income Shocks

While the above method helps to isolate shocks that are transitory, we do need to make strong assumptions about the exogeneity of these shocks. This is where methods from the other strand of the consumption smoothing literature, which exploits natural experiments to isolate exogenous shocks, is helpful. To disentangle changes in income from endogenous labor supply decisions, we use variation in pay at the firm-level. This approach is inspired by the labor economics model of firm effects pioneered in Abowd, Kramarz, and Margolis (1999, henceforth AKM). We first extend the income process by decomposing the transitory and permanent shocks into a firm component and a unique worker-specific component:

$$\begin{aligned} y_{i,t} &= z_{i,t} + \varepsilon^w_{i,t} + \varepsilon^f_{j(i,t),t} \\ z_{i,t} &= z_{i,t-1} + \eta^w_{i,t} + \eta^f_{j(i,t),t} \end{aligned} \tag{11}$$

where η^w and ε^w are worker-specific shocks, $j(i,t)$ is a worker i's firm at time t, and $\varepsilon^f_{j(i,t),t}$, and $\eta^f_{j(i,t),t}$ are the firm effects at time t. Equation (11) follows Lachowska et al. (2023) in generalizing AKM to allow for *time-varying* firm effects. This results in the following expression for income growth:

$$\Delta y_{i,t} = \eta^w_{i,t} + \eta^f_{j(i,t),t} + \Delta\varepsilon^w_{i,t} + \Delta\varepsilon^f_{j(i,t),t} \tag{12}$$

We next extend the consumption function to explicitly include firm-specific and worker-specific transitory shocks separately:

$$\Delta c_{i,t} = \beta_\eta \left(\eta^w_{i,t} + \eta^f_{j(i,t),t}\right) + \beta_\varepsilon \left(\varepsilon^w_{i,t} + \varepsilon^f_{j(i,t),t}\right) + \zeta_{it} \tag{13}$$

Equation (13) makes a substantive restriction that the consumption response to firm-wide temporary shocks $\varepsilon^f_{j(i,t),t}$ is the same as the response to individual-specific temporary shocks ε_{it}. This assumption could be violated in the presence of mental accounting or peer effects (De Giorgi, Frederiksen, and Pistaferri 2020).[17]

Building on our previous assumptions, we make the following amended ones to incorporate the firm component of the shock:

Assumption IID2 (IID Income Shocks II)

The worker-specific shocks η^w_{it}, and ε^w_{it} are independent and identically distributed (iid) at the (i,t) level:

$$(\eta^w_{it}, \varepsilon^w_{it}) \sim IID;$$

[17] De Giorgi, Frederiksen, and Pistaferri (2020) documents that idiosyncratic changes in the income of one peer (e.g an income shock to a "coworker of the spouse of my coworker") affect consumption via peer effects. One standard story motivating such a behavioral channel is a "keeping up with the Joneses" logic where my consumption choices are driven in part by a desire to keep up with the consumption of my higher-consuming peers. It is unclear if this story applies to the firm pay shocks we study, where most workers in the firm experienced a similar income shock.

and are independent of each other:

$$\eta^w_{i,t} \perp \varepsilon^w_{i,t} \quad \forall t;$$

and the firm-specific shocks $\eta^f_{j(i,t),t}$ *and* $\varepsilon^f_{j(i,t),t}$ *are iid at the* (j,t) *level:*

$$(\eta^f_{j(i,t),t}, \varepsilon^f_{j(i,t),t}) \sim IID;$$

and are independent of each other

$$\eta^f_{j(i,t),t} \perp \varepsilon^f_{j(i,t),t} \quad \forall t$$

These three assumptions are, again, formal statements of the iid properties of the various shocks. In addition, we have a modified restriction on how income shocks relate to current consumption:

Assumption CE2 (Consumption Exogeneity II)

$$\zeta_{i,t} \perp (\eta^f_{j(i,t),t+s}, \varepsilon^f_{j(i,t),t+s}) \quad \forall s$$

and if $j(i,t) = j(i',t+1)$ *then*

$$\zeta_{i,t} \perp (\eta^w_{i',t+s}, \varepsilon^w_{i',t+s}) \quad \forall s$$

These assumptions require that firm-level and coworkers' individual level income shocks are orthogonal to innovations to worker consumption. Assumption CE2 does not require that *all* changes in firm income are orthogonal to the consumption error. For example, if some workers choose to work in education because they have a preference for leisure (and spending less money) in the summer, this will generate predictable seasonality in earnings. As long as our seasonal adjustment procedure (described in Appendix B.1) is successful, changes in income from say a school's summer vacation will not show up in the instrument. Nevertheless, Assumption CE2 would be violated if workers sort to firms based on private information regarding future non-seasonal firm shocks. An alternative violation—a type of exclusion restriction violation—would occur if it were the case that changes in ε^f affect earnings via hours worked and hours worked directly affect consumption above and beyond the impact on earnings.[18]

Note, however, that Assumption CE2 is weaker than HM-BPP's Assumption CE in one crucial way. It is still satisfied even if a worker's consumption choices are endogenous to their own specific temporary changes in income, i.e. $\text{cov}(\zeta_{it}, \varepsilon^w_{it}) \neq 0$. Unobserved shocks that affect both income and consumption (like a vacation, a health shock, or a desire to increase labor supply to finance a specific purchase) will not be a violation of this assumption. This is the same endogeneity problem that the quasi-experimental literature seeks to address.

To isolate the firm component of labor income changes, we construct the negative of the leave-out-mean instrument for income changes using coworkers at the firm j:

$$-\Delta y_{j(-i,t),t+1} \equiv -\Delta \ln \left(\frac{1}{N_{j(-i,t),t+1}} \sum_{i'|j(i',t+1)=j(i,t),i'\neq i} Y_{i',t+1} \right) \tag{14}$$

where $N_{j(-i,t),t+1}$ is the number of coworkers at firm j in period $t+1$. We leave out worker i from the mean to avoid mechanical correlation between this instrument and monthly income changes for agent

[18] In Section 4.3 we report results from tests for a direct effect of hours worked on consumption.

i. This approach of using coworker's income variation is similar in spirit to the research design used in Koustas (2018).[19] Our innovation is to combine this with the HM-BPP research design, yielding the following identification result with a one-period lead, coworker instrument:

Result ID (Identification of β_ε II)

If Assumptions IID2 and CE2 hold, β_ε is identified by using the (negative) one-period lead, coworker income growth instrument:

$$\hat{\beta}_{\varepsilon,HM-BPP-Co} \equiv \frac{\text{cov}\left(\Delta c_{it}, -\Delta y_{j(-i,t),t+1}\right)}{\text{cov}\left(\Delta y_{it}, -\Delta y_{j(-i,t),t+1}\right)} = \frac{\text{cov}\left(\Delta c_{it}, \varepsilon^f_{j(i,t),t}\right)}{\text{var}\left(\varepsilon^f_{j(i,t),t}\right)} = \frac{\beta_\varepsilon \sigma^2_{f\varepsilon}}{\sigma^2_{f\varepsilon}} = \beta_\varepsilon$$

where $\sigma^2_{f\varepsilon}$ is the variance of the transitory firm shock.

Result ID is the key methodological result of the paper, bridging the covariance restrictions literature and the quasi-experimental literature. It isolates transitory income shocks using the change in income from t to $t+1$ and has a coherent theoretical interpretation, building on HM-BPP. At the same time, it uses changes in firmwide pay, $-\Delta y_{j(-i,t),t+1}$, as an instrument for changes in own pay and uses pre-event tests to probe the validity of identifying assumptions, following the quasi-experimental literature.

3 Data

We study consumption smoothing using bank account data from the JPMorgan Chase Institute (JPMCI). This de-identified dataset has three advantages over existing data sources: it has a large sample size with millions of households, it includes reliable measures of liquid wealth, and it contains employer identifiers, which can be used to construct plausible instruments for income fluctuations.

Our analysis sample is drawn from the 50.1 million households with a Chase checking account in the JPMCI data from August 2018 to December 2022. The unit of observation is household-by-employer-by-month.[20]

Our analysis centers on three variables from this dataset: income, spending, and assets. JPMCI categorizes each checking account inflow by whether it appears to reflect income. Most of our analysis focuses on labor income, for which we observe the counterparty who made the direct deposit. This field enables us to identify other workers with Chase bank accounts who are paid by the same employer. This is useful because we use fluctuations in firm pay to instrument for changes in a household's total labor income. Figure A-1a shows that the distribution of labor income for our analysis sample is similar to the distribution of labor income for typical employed households in the Current Population Survey (CPS), with a slight over-representation in the middle of the income distribution.

[19] It is also complementary to another strand in the consumption-smoothing literature which instruments for earnings changes using changes in observable firm characteristics such as changes in union contracts (Shea 1995), value added (Fagereng, Guiso, and Pistaferri 2017), and stock prices (Baker 2018).

[20] For our purposes, a household is a unit of checking accounts that are administratively linked. This will tend to capture inflows and outflows for single-adult households and multi-adult households where members share linked accounts. In principle, a household with multiple, non-linked accounts at Chase, as in the case of a married couple where each spouse holds independent accounts, will appear as separate observations in the data. However, Ganong and Noel (2019) shows that 79% of households actively link accounts, consistent with survey evidence suggesting that most people link their bank accounts administratively when they get married (TD Bank 2014). We do not capture income sent to non-Chase bank accounts. As we note below, we estimate in the Survey of Consumer Finances (SCF) that two-thirds of households have checking accounts only at one financial institution.

The second key variable is a measure of spending on nondurable goods and services. Spending is measured from debit and credit card transactions, cash withdrawals, and electronic transactions captured through the bank account. Ganong and Noel (2019) constructs a measure of nondurable spending in the JPMCI data, and we adopt that definition in this paper. Examples of nondurable spending include groceries, food away from home, fuel, utilities, clothing, medical co-pays, and payments at drugstores. Altogether, transactions that can be affirmatively categorized as nondurable comprise an average of 26 percent of checking account outflows.

Ganong and Noel (2019) shows that this measure of nondurable spending is more comprehensive than what can be measured in U.S. survey data, but nevertheless has some limitations which are relevant for this paper. First, we only observe data on bank account and credit card spending at a single financial institution. We estimate using the 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 were to have 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. Analysis of the SCF indicates that checking account balance is a useful proxy for total liquid assets. Average log liquid assets increase one-for-one with average log checking account balances, as shown in Figure A-2. Figure A-1b shows that the distribution of checking account balances for our analysis sample is similar to the distribution of checking account balances for employed households in the SCF.

Because of computational constraints, we use a 10% subsample of households. However, we are able to construct a variable for the pay of all coworkers before we draw the 10% subsample so this restriction does not lead to additional noise in our estimates of the change in firm pay. We also impose sample screens to ensure that we focus on employed households whose primary bank accounts are at Chase and for whom we reliably measure labor income (see Appendix F for more details). Our final analysis sample after taking the random subsample and imposing these screens includes 1,327,214 households and 27,881,033 household-employer-months. Table A-1 shows summary statistics for this analysis sample for each of our key variables (labor income, nondurable consumption, and checking balances).

We also use three supplementary datasets to understand the degree to which changes in coworker pay are correlated with own pay and the source of this correlation. The first data set is the Continuous Wage and Benefit History (CWBH) which captures quarterly wage records for unemployment insurance (UI) claimants from seven U.S. states. We are able to construct the change in coworker pay because the dataset includes information on total firm pay and total number of employees at the firm for each UI claimant. The second dataset is time clock data detailing when workers start and stop their shifts.[21] The company's software is primarily used by small service-sector firms such as coffee shops.

The third supplementary dataset is made up of de-identified administrative earnings records from an anonymous U.S. payroll processor (henceforth PayrollCompany). Following Ganong et al. (2025), our analysis utilizes payroll data covering the period from 2010 to 2023, during which time PayrollCompany serviced approximately two million client firms. We observe paycheck information for all

[21] This dataset is from Homebase (joinhomebase.com).

workers employed at firms included in the dataset, between two and four million employees monthly. Although PayrollCompany primarily serves small firms, Ganong et al. (2025) shows that these data are generally representative of the level of earnings volatility, the prices of labor, the quantities of labor, and contract types for U.S. workers overall.

4 Results

This section implements the HM-BPP-Coworker research design to estimate the causal impact of exogenous, transitory, and unpredictable labor income shocks on consumption. We begin by presenting empirical tests of the research design's identification assumptions. We then present the paper's three main empirical results: first, monthly labor income shocks meaningfully affect consumption for the average household; second, consumption sensitivity varies sharply with a household's liquid assets; third, households are less responsive to relatively more predictable sources of income variation.

4.1 First Stage Estimate and Sources of Variation

The HM-BPP-Coworker research design relies on a number of identifying assumptions. We first investigate instrument relevance. We find that there is an economically and statistically significant relationship between own pay and the period-ahead change in coworker pay, which will become the first stage in the IV estimate below. We estimate

$$\Delta y_{it} = \phi - \rho \Delta y_{j(-i,t),t+1} + \nu_{it}. \tag{15}$$

Firm pay shocks estimated using coworkers have strong predictive power for individual pay. The first stage relationship is reported in Table A-2 column (1) and depicted graphically in Figure A-3a.[22] We estimate ρ of 0.153 with a standard error of 0.007, which means that a 10% decrease in period-ahead coworker pay is associated, on average, with a 1.5% increase in own pay this period.[23]

This strong relationship between changes in own pay and changes in firm pay is a robust finding across many different datasets. We re-estimate equation (15) in payroll data, time clock data, and in U.S. tax data, each of which are described in more detail in Section 3. Although the JPMCI estimate of the change in coworker pay relies on the subset of coworkers who also happen to have bank accounts with Chase, all three of these alternative datasets capture the change in pay for the universe of coworkers at a given firm. We report these estimates in Table A-4. In the payroll data and the time clock data we obtain estimates of ρ of 0.17 and 0.14, respectively, very similar to the estimate of 0.15 from the JPMCI data. The tax data are only available at the quarterly level. We estimate ρ

[22] Figure A-3a uses a linear spline to depict the somewhat nonlinear relationship between coworker pay and own pay changes. Since the second stage elasticity estimate of 0.25 using the linear spline in Figure A-3c is very similar to the elasticity estimate of 0.22 shown in Table 1 below in a specification without the spline, we use the less structural approach without a spline for the remainder of the analysis. Table A-3 reports summary statistics for the coworker instrument and the fitted values from the first stage regression.

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

separately for the seven states in the tax data and find estimates of 0.42 to 0.61, as compared to 0.58 at the quarterly level in the JPMCI data. Other researchers have also found a strong first stage when estimating similar regressions: Lachowska et al. (2022) report coefficients between 0.56 and 0.65 for earnings per hour while Koustas (2018) reports an F-statistic of 104 for earnings per fortnightly pay period (no coefficient or standard error reported). These strong relationships suggest that firm pay may be a useful instrument for own pay in any employer-employee matched dataset.

We can leverage one of these alternative datasets to provide more detail on the sources of the pay shocks captured by the first stage in this research design. Although the JPMCI bank account data cannot separate between different components of worker pay, payroll data can. In particular, PayrollCompany data separately records changes in base wages (which is the hourly wage rate for hourly workers and the per-period contracted regular pay for salaried workers), total earnings, and (for hourly workers) total hours. Figure G-1 separates the first stage variation in PayrollCompany data into these various sources. It demonstrates that variation in hours worked drives most of the variation for hourly workers (who make up 60% of U.S. workers), while variation in non-base salary compensation drives most of the variation for salaried workers. In contrast, almost none of the variation picked up by the instrument is attributable to changes in base wages. This is not surprising, since base wages are sticky (Grigsby, Hurst, and Yildirmaz 2021) and the instrument is specifically constructed to isolate transitory rather than permanent changes in pay.

4.2 Diagnostic Tests Using Pre-and Post-Event Trends in Income and Consumption

Next, to assess whether the other identification assumptions are satisfied, we implement the event-study diagnostics for income and consumption proposed in Section 2.3. Define treatment following equation (6) except using coworker income instead of own income as:

$$D_{i,t^*} \equiv \mathbf{1}\{-\Delta y_{j(-i,t^*),t^*+1} > M_{-\Delta y}\} \tag{16}$$

Figures 2a and 2b study income and consumption for workers in the treatment group $D_{i,t^*} = 1$ and in the control group $D_{i,t^*} = 0$ for $s \leq t^*$.[24] Four patterns are readily apparent.

First, the most striking feature of Figure 2a is that income increases for the treatment group and falls for the control group at time 0 (equivalent to time t^* in our formal notation). This clarifies that comparing observations with different values of treatment D_{i,t^*} captures the difference between workers whose income goes up versus workers whose income goes down. Thus, this is akin to a dose-response difference-in-difference design, comparing relatively more treated (positive transitory shock) workers to relatively less treated (negative transitory shock) workers.

Second, Figure 2b shows that consumption diverges sharply between the treatment and control groups in period 0, coincident with the divergence in income between the two groups. This pattern is consistent with a causal impact of income on consumption.

Third, income and consumption are trending similarly for the treatment and control groups in the year before treatment is realized. Thus, it appears that parallel pre-trends is satisfied. Finally, there is a secular trend where income and consumption are trending up over time for both the treatment and control groups.

[24] This figure and all subsequent event studies imposed a balanced panel restriction. We subset to households with positive labor income and positive consumption in all months between $t = -12$ and $t = 0$

Figure 2: Time Paths of Income and Consumption Using Binary Instrument

(a) Income

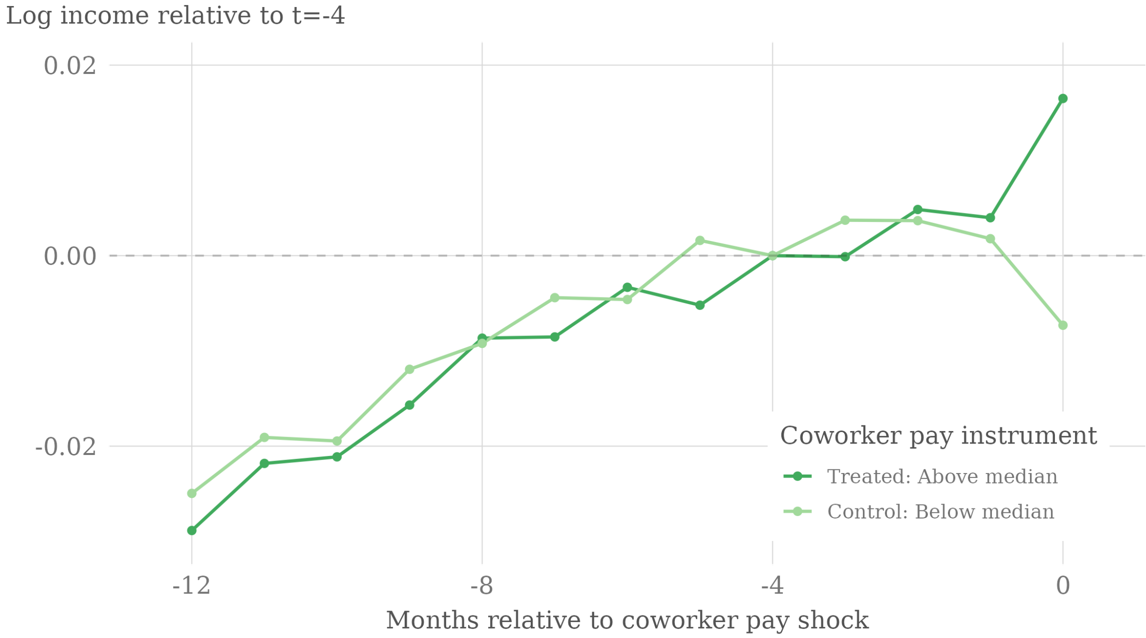

(b) Consumption

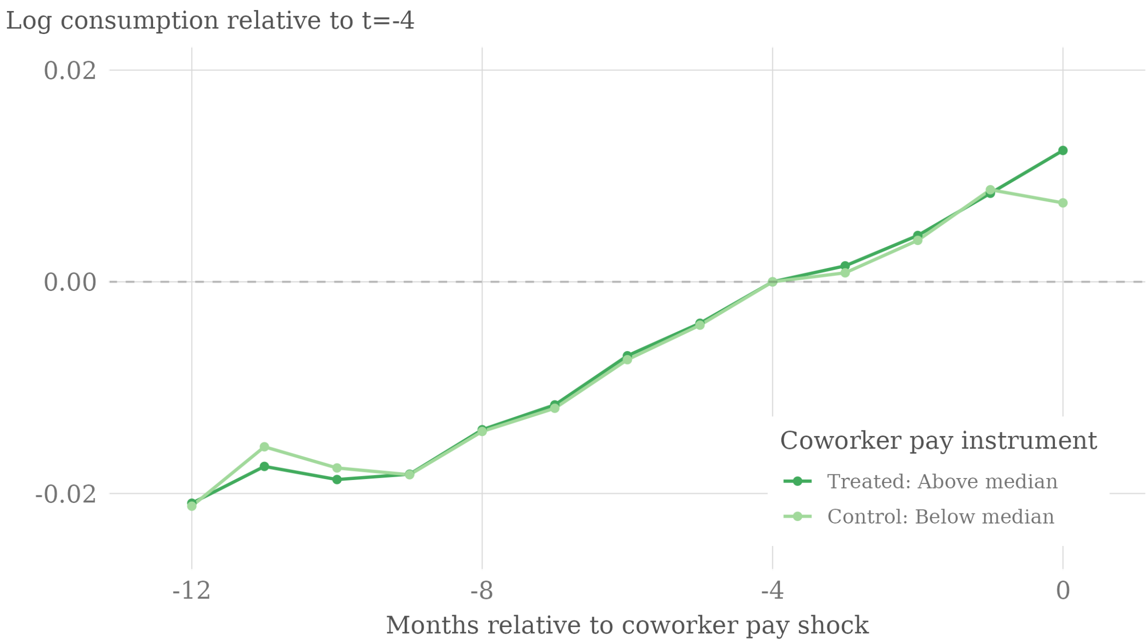

Notes: This figure shows the time paths of income and consumption using a binary instrument as defined in equation (16). The first panel shows the levels of log seasonally adjusted income relative to $t = -4$ for the treated and control groups. In these plots event time 0 corresponds to time t^* (the date of instrument realization) in the formal notation. The second panel shows the levels of log consumption relative to $t = -4$ for the treated and control groups.

These observations motivate us to generalize equation (15) to show the dynamics of income using an event study-like specification:

$$y_{i,t^*+s} = \sum_{k=-12,k\neq -4}^{12} \rho_k \times \mathbf{1}\{k = t^* + s\} \times D_{i,t^*} + \phi_s + \nu_{i,t^*+s} \tag{17}$$

where s captures time relative to treatment t^* and the ϕ_s controls absorb the upward secular trend in income shown in Figure 2a. Although some event study specifications normalize ρ_{-1} to zero, we instead normalize ρ_{-4} because it makes it easy for the reader to evaluate changes in income in the months immediately prior to t^*.[25]

Figure 3 reports estimates and confidence intervals from equation (17). We use two different definitions for D_{i,t^*}. The first is the binary definition from equation (16), which clarifies the connection to an event study with two groups. The second is $D_{i,t^*} = -\Delta y_{j(-i,t),t^*+1}$, which allows for continuous treatment intensity and is our preferred version for estimating consumption sensitivity leveraging all the variation in the data. To ease interpretation, in each panel we report both the event study coefficients that we estimate from the data and event study estimates that are generated when we simulate income using equation (11).[26] This simulated income series provides a benchmark for what the event study estimates would show if the assumptions underlying the empirical approach were satisfied.

Figure 3 empirically tests (and confirms) each of the four theoretical predictions for income detailed in Section 2.3 and summarized in Figure 1, as well as the continuous counterparts of those predictions detailed in Appendix E. First, as in Testable Prediction 1 ("First Stage"), we see a substantially higher income for treatment than control at t^*.[27] This supports the claim that the research design captures a transitory income shock at t^*. Second, as in Testable Prediction 2 ("Parallel Pre-Event Income Trends"), we see no differential pre-trend in income between treatment and control for $s < 0$. This indicates that shocks to coworker pay are unpredictable from the econometrician's perspective.

The third and fourth testable predictions concern post-trends in income. As in Testable Prediction 3 ("Negative Difference at $t^* + 1$"), we see substantially lower income for treatment than control at $t^* + 1$. And as in Testable Prediction 4 ("Negative Difference from $t^* + 2$ Onward"), we see slightly lower income for $s \geq 2$.[28] Both of these post-trends are what we would expect under the assumed permanent-transitory income process; they are the result of transitory and permanent shocks at $t^* + 1$. Although the instrument captures three types of structural shocks, giving rise to this sawtooth pattern in income, the crucial point for identification is that the only one of these three shocks realized as of date 0 is the transitory shock at time t^*. The method therefore calls for examining consumption only through date 0 to assess the causal impact of just this transitory shock.

We assess the fifth and final testable prediction of the identifying assumptions by examining pre-trends in consumption in Figure 4, which we plot alongside the path of income. Concretely, we re-

[25]Normalizing $\rho_{-4} = 0$ is equivalent to normalizing log income by income in period $t^* - 4$, i.e. setting $r = -4$ in the results of Section 2 or Appendix E.

[26]We provide detail on this simulation in Appendix H. The key point is to explore whether an income process with the general transitory and permanent shock structure that we assume generates a sawtooth pattern in income consistent with the pattern we observe in the data.

[27]For the continuous case, $\rho_0 = 0.17$ which is very similar to the estimate of $\rho = 0.15$ above. They are not identical because the former uses a dependent variable of $y_{i,t^*} - y_{i,t^*-4}$ while the latter uses a dependent variable of $y_{i,t^*} - y_{i,t^*-1}$.

[28]The fact that income is only slightly lower is consistent with the interpretation that at the monthly level, temporary firm shocks are much more important than permanent firm shocks.

Figure 3: Income Event Study Around Instrument Realization

(a) Binary Instrument

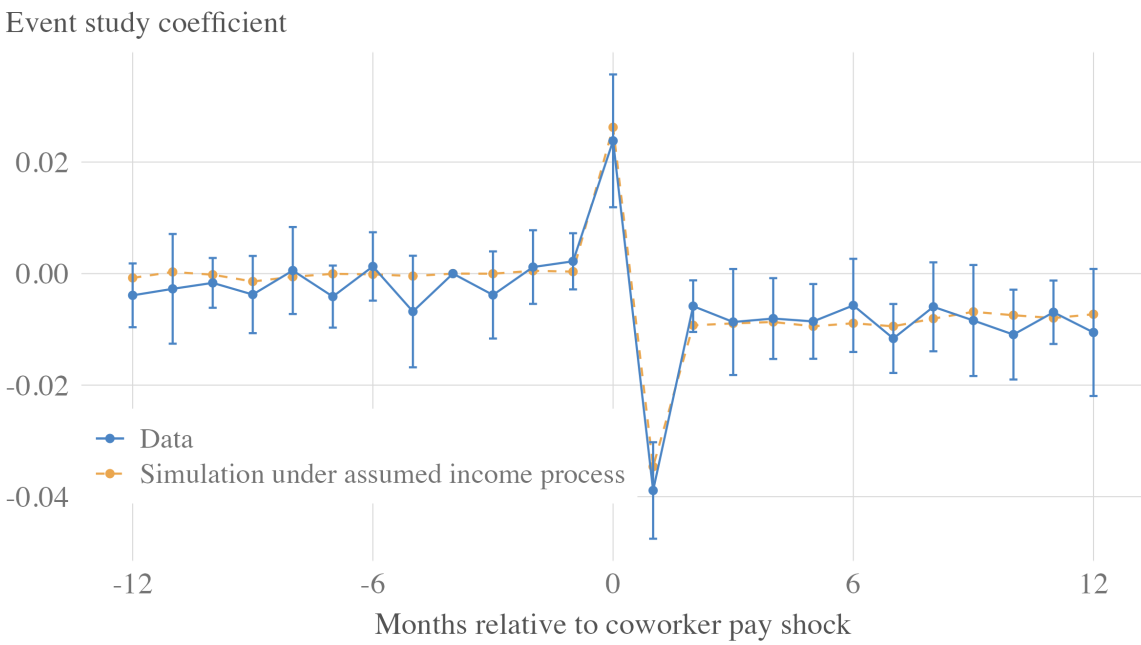

(b) Continuous Instrument

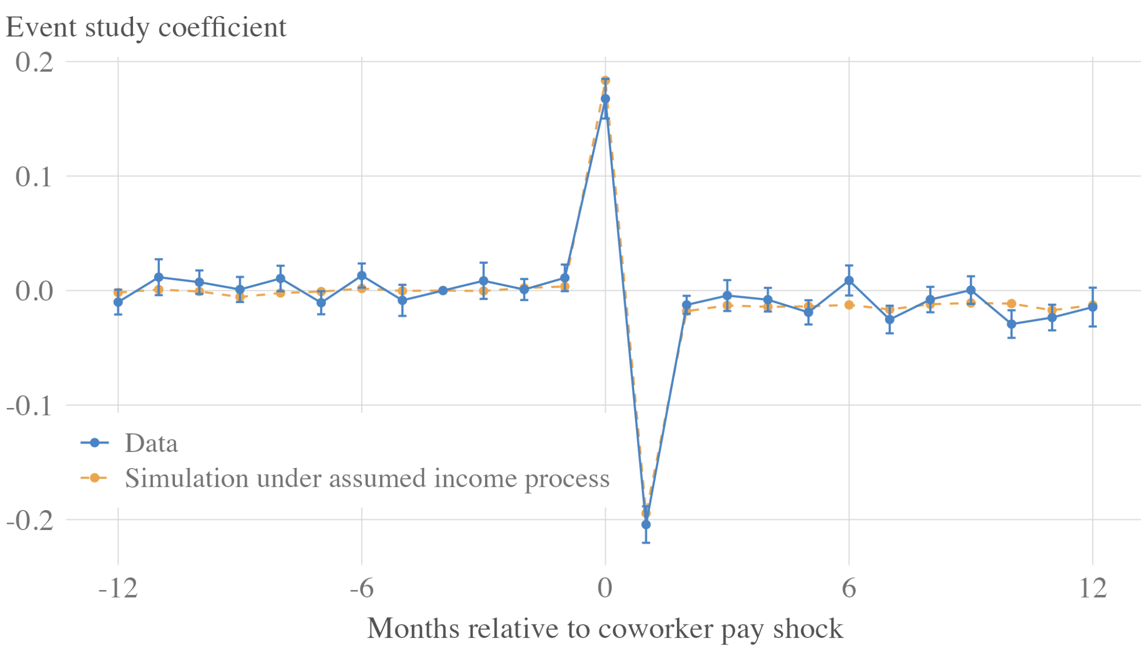

Notes: This figure shows the time path of event study–like coefficients ρ_k from estimating equation (17) using a binary instrument (as defined in equation 16) in the first panel, and a continuous instrument (as defined in equation 14) in the second panel. The regressions are estimated using seasonally adjusted income changes. The two figures compare estimates obtained from the JPMCI data to estimates from simulating income using equation (11). The simulation includes both worker and coworker shocks (see Appendix H for details). Both panels include uniform 95% confidence intervals for the empirical estimates.

Figure 4: Income and Consumption Event Studies Around Instrument Realization

(a) Binary Instrument

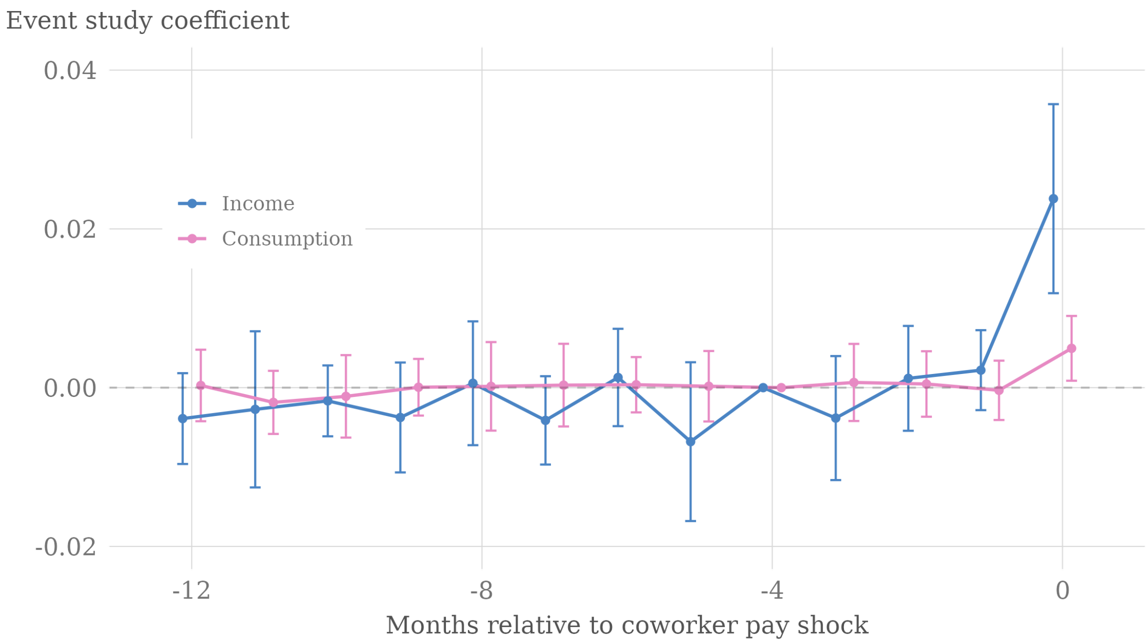

(b) Continuous Instrument

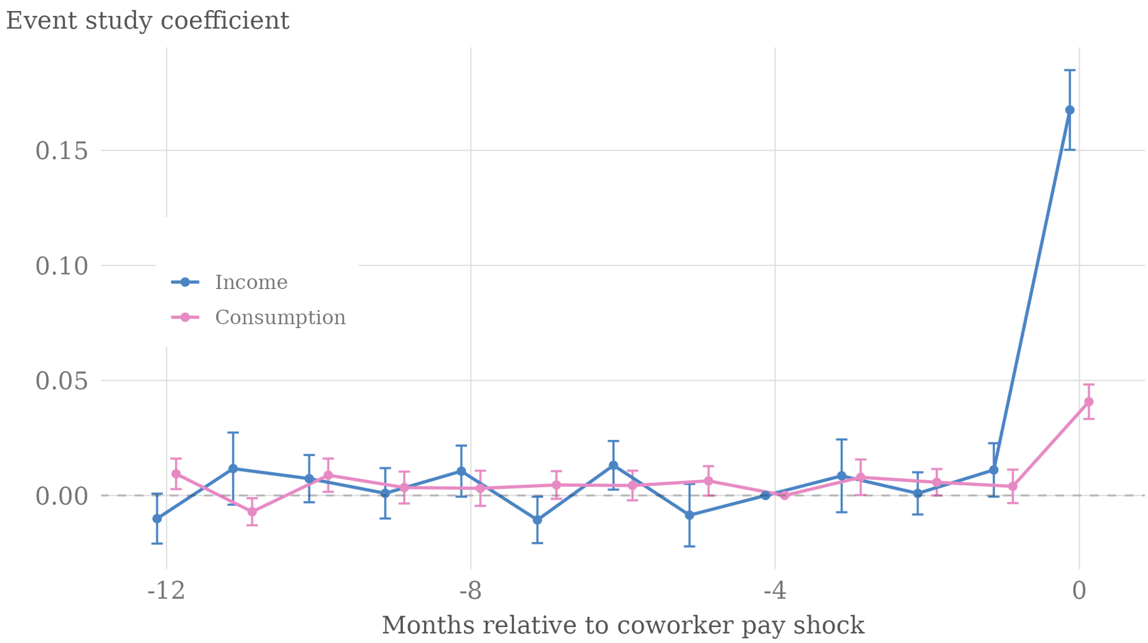

Notes: This figure shows the time path of event study-like coefficients ρ_k for income and consumption from estimating equation (17) using a binary instrument (as defined in equation 16) in Panel (a), and a continuous instrument (as defined in equation 14) in Panel (b). The estimates in Panel (a) recover the difference between the treatment and control groups depicted in Figure 2. Both panels include 95% uniform confidence intervals.

estimate equation (17) substituting c_{i,t^*+s} as the dependent variable. Testable Prediction 5 ("Parallel Pre-Event Consumption Trends") predicts that consumption trends similarly for workers who will see an increase in income at $t = 0$ and those who will see a decrease. The figure supports this prediction. The finding that workers do not appear to adjust consumption before the pay shock is realized is consistent with the assumption that this shock is unanticipated. Of course it could still be that the household has some degree of advance information, but is unwilling or unable to adjust consumption ahead of the income shock. The farthest we can go is to show that the household at least *acts* as if the shock is unanticipated.

The absence of pre-event trends for income and consumption together with the sawtooth post-trends for income are consistent with the identifying assumptions underlying our empirical approach. These findings motivate a focus in the rest of the section on the contemporaneous effects on consumption at time t associated with the change $-\Delta y_{j(-i,t),t+1}$.

4.3 Causal Impact of Income on Consumption

The event study patterns shown in Figures 2 and 4 provide visual evidence that consumption rises in response to a transitory income shock. A simple Wald estimator can be used to quantify the effect of income on consumption from the binary event study in Figure 2. The relative income change between treatment and control is 2.16% and the relative consumption change is 0.52%, implying an elasticity of 0.52 / 2.16 = 0.24.

To obtain quantitative estimates of consumption sensitivity using the *full* support of firm pay shocks, we estimate equations (1) and (2), repeated here for convenience:

$$\Delta c_{it} = \alpha + \beta \Delta y_{it} + \zeta_{it} \tag{18}$$

$$\Delta y_{it} = \phi - \rho \Delta y_{j(-i,t),t+1} + \nu_{it}. \tag{19}$$

As we discuss in Section 2.1, in order to remove predictable changes in pay the coworker instrument $-\Delta y_{j(-i,t),t+1}$ is measured as monthly pay per paycheck and is the residual after removing recurring annual seasonality. Household labor income changes Δy_{it} are also seasonally adjusted. In a number of cases, when a household i has labor income from more than one employer j in a given month—because there are two workers in the household or because one worker has two jobs—then the household-month enters the regression twice. In each case changes in consumption Δc_{it} and labor income Δy_{it} are aggregated at the household level, while the instrument $-\Delta y_{j(-i,t),t+1}$ is employer-specific. Thus we define the unit of observation as the household-employer-month and use two-way clustered standard errors by firm j and household i.

Table 1 column (1) reports our headline estimate of the consumption response to a transitory income shock. We estimate an elasticity of nondurable consumption to income of 0.22 ($\hat{\beta}$=0.22), with a standard error of 0.015.[29,30] This estimate, which leverages the full support of firm pay shocks using the regression approach in equations (18) and (19), delivers the same conclusion as the simple Wald estimate from the binary event in Figure 2. This similarity reinforces the connection between

[29] Figure A-3b shows a binscatter of the reduced-form and Figure A-3c shows the second stage.

[30] The result is robust to allowing for an alternative income process that features a more general autoregressive component. In Appendix B.3 we derive an estimator for the elasticity to transitory shocks under an income process with an autoregressive component with coefficient less than one. When we implement this estimator in our data we recover an elasticity of $\hat{\beta} = 0.20$ with a standard error of 0.013, similar to our headline estimate.

Table 1: Impact of Income on Consumption

	(1)	(2)	(3)	(4)	(5)
Δ Log Income	0.221	0.274	0.183	0.189	0.109
	(0.015)	(0.018)	(0.009)	(0.008)	(0.003)
(Δ Log Income) × Checking		−0.475	−0.334	−0.348	−0.166
		(0.047)	(0.022)	(0.019)	(0.006)
Coworker Instrument	Period-Ahead Pay Per Check	Period-Ahead Pay Per Check	Pay Per Check	Pay Per Check	Total Pay
Seasonal Adjustment	Yes	Yes	Yes	No	No
Type of Income Variation Captured by Instrument					
New Transitory Shock	Y	Y	Y	Y	Y
Predictable Reversion of Transitory Shock			Y	Y	Y
Permanent Shock			Y	Y	Y
Predictable Recurring Annual Changes				Y	Y
Predictable Pay Schedule Variation					Y

Notes: This table reports estimates of the elasticity of consumption with respect to income ($\hat{\beta}$). Columns (1) and (2) show IV estimates of the effect of income on consumption using equations (19) and (18). Column (2) controls for checking account buffer. The checking buffer variable is parameterized as $BufferRank/N$ - 0.5, so that the variable is scaled from −0.5 for the lowest asset buffer household to 0.5 for the highest asset buffer household. Column (3) shows IV estimates using a contemporaneous coworker pay per paycheck instrument, rather than a period-ahead one. Column (4) shows IV estimates using a contemporaneous coworker pay per paycheck instrument without seasonal adjustment. Column (5) shows IV estimates using a contemporaneous coworker total pay instrument without seasonal adjustment. All specifications use data from 27,881,033 household-employer-months. Standard errors, which are clustered two ways at the household-job level, are reported in parentheses.

Figure 5: Dynamic and Cumulative MPCs

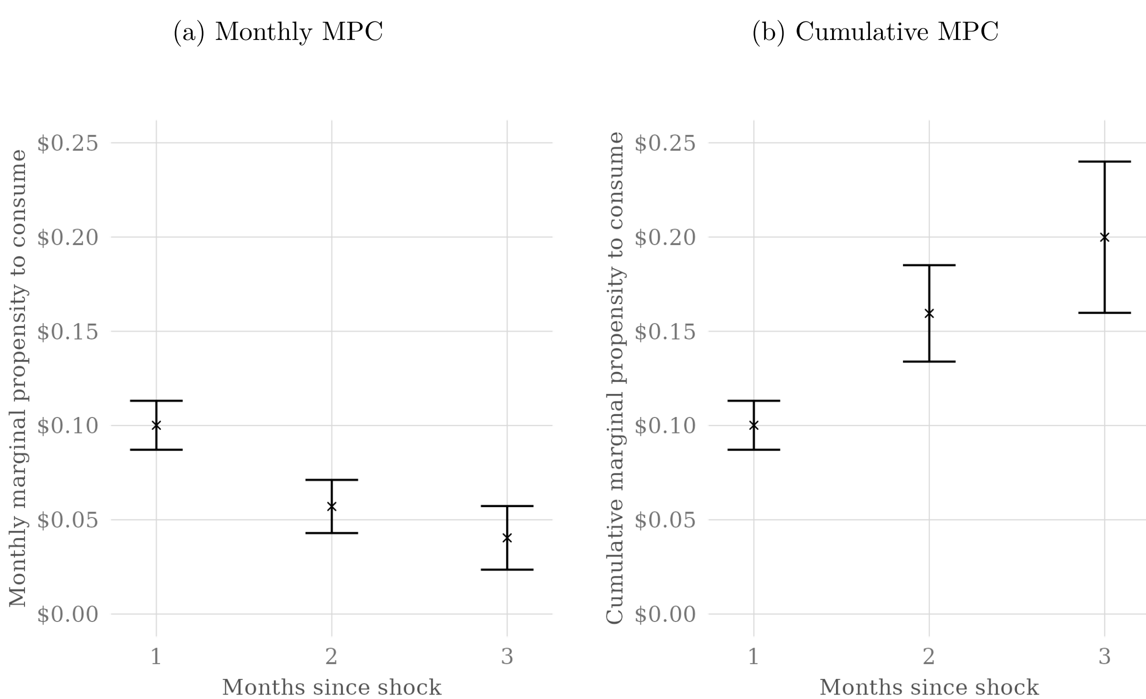

Notes: This figure shows the monthly and cumulative MPC out of a transitory income shock. Panel (a) plots the month-by-month MPCs, while Panel (b) plots the cumulative MPC out of the shock. Ninety-five percent confidence intervals are reported. Standard errors are two-way clustered at the household–job level. Point estimates and standard errors are reported in Table A-5. See Appendix J for further details.

our empirical approach and a standard event study difference-in-differences framework used in the quasi-experimental literature, as foreshadowed in Section 2.4, Result Wald.

We also report our consumption sensitivity estimate in terms of an MPC since some of the prior literature and many structural models focus on MPCs rather than elasticities. We translate our elasticity into an MPC by rescaling our estimate by the ratio of mean nondurable consumption to mean income in our sample.[31,32] We estimate a monthly nondurable consumption MPC of 0.10, with a standard error of less than one cent (see Figure 5 and Table A-5).[33]

We implement four tests which indicate that the consumption sensitivity we document is not driven by high-frequency substitution of spending across months, by changes in time use, by spending categories with weak links between expenditure and actual consumption, or by local macroeconomic trends. First, we investigate the concern that monthly consumption sensitivity estimates may simply be picking up inter-temporal substitution—an increase in current consumption could be coupled with offsetting reductions in the near future. To assess this concern, we follow previous literature in tracing out the dynamic response of consumption over multiple periods, i.e. the intertemporal marginal propensities to consume (iMPCs) (Auclert, Rognlie, and Straub 2024; Boehm, Fize, and Jaravel 2025; Parker et al. 2013). We extend the consumption function in equation (4) to allow consumption to more flexibly respond to a transitory shock in the contemporaneous month and in the two subsequent months and estimate those dynamic responses.[34]

The left panel of Figure 5 plots the MPC in the month of a transitory shock and in the two subsequent months, while the right panel plots the cumulative MPC. Contrary to an offsetting of the initial response, we find that the consumption response continues to build over the subsequent two months, translating to a cumulative quarterly MPC of 0.20.[35] This shows that households are not just responding to firm pay shocks by reallocating spending from one month to the next. Furthermore, it suggests that households are not simply dipping into their stockpiles while keeping true consumption smooth, since Baker, Johnson, and Kueng (2021) shows that such inventory management happens at a time horizon shorter than one quarter.[36]

Second, we find little evidence that spending fluctuations are being driven by work-related expenses. We define a spending category as work related if it exhibits a larger-than-median drop at retirement (Aguiar and Hurst 2013). We find that the spending response for non-work-related expenses is 0.26, similar to the overall spending response (Table A-6). This is consistent with the view that changes in spending are driven by changes in income rather than changes in time use, as required

[31] An alternative strategy would be to estimate equations (18) and (19) in levels instead of logs. However, this is not our preferred strategy because a dollar-based specification upweights the consumption response of high-income, high asset households (Yitzhaki 1996), whose shocks are larger in dollar terms (but not in percentage terms). This concern arises because our research design captures the average response to shocks of varying sizes. In contrast, previous strategies estimating shocks that are similar in dollar terms across groups do not have this bias in estimating MPC measures using a levels specification.

[32] As discussed in Ganong and Noel (2019), average nondurable spending measured in JPMCI data is in between nondurable spending benchmarks from the Consumer Expenditure (CE) Survey and the National Income and Product Accounts (NIPA) measure of Personal Consumption Expenditures (PCE). Readers who seek to scale our MPC estimates to match NIPA benchmarks following the procedure in Broda and Parker (2014) should multiply our MPCs by 1.5. We discuss this scaling in more detail in Appendix I.

[33] When calculating the standard error of the MPC we do not incorporate sampling error of the scaling factor, which would make a negligible difference given the large sample size.

[34] See Appendix J for further details and identification proofs.

[35] We lack the statistical precision to say much about responses farther out than one quarter.

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

by the exclusion restriction in Assumption CE2.[37]

Third, we show that spending is sensitive across a range of alternative measures of consumption. The most restrictive measure of nondurable spending we consider is one where we subset only to grocery spending. This is useful because many groceries are consumed shortly after they are purchased. Table A-6 shows that grocery spending is highly sensitive, with an elasticity of 0.22, the same as the elasticity for all nondurable spending categories. We also consider a measure of "strict nondurables" (categorized following the taxonomy in Lusardi 1996), which is broader than groceries but narrower than total nondurable spending. This in-between measure is also sensitive, with an elasticity of 0.15.

We also follow the method developed by Laibson, Maxted, and Moll (2022) to map our empirical estimate of the marginal propensity to spend on nondurables (the "nondurable MPX" using their terminology) into estimates of the marginal propensity to spend on all goods and services (the "total MPX") and the model-consistent, or "notional," MPC. These estimates are reported in Table A-7. The total spending MPX scales up our nondurable MPX to account for the response of durables. We find a quarterly total spending MPX of 0.72. The notional MPC is the MPC that should be used as a target for standard consumption models, and incorporates the fact that spending on durable goods generates a consumption flow over multiple periods. We find a quarterly notional MPC of 0.23.

Finally, we find no evidence that the consumption impacts we document are driven by unobserved macroeconomic shocks rather than by changes in pay. If changes in pay were correlated with local macroeconomic shocks, and these shocks directly affected consumption, this would be a violation of the exclusion restriction in Assumption CE2. However, Table A-8 shows that our estimates of consumption sensitivity are nearly unchanged when we add time-by-location fixed effects, narrowing in on variation across firms exposed to the same aggregate shocks.

Overall, the finding that consumption is sensitive not just to infrequent windfalls (as has been documented in prior work), but even to frequent monthly labor income shocks has two implications. First, it suggests that monthly labor income volatility may have large welfare costs. Second, this behavior is inconsistent with some "near rational" explanations of consumption sensitivity. We discuss both of these points in more detail in Section 5.3.

4.4 Consumption Smoothing by Liquid Assets

Having documented that consumption is sensitive to temporary income shocks, the next natural question is "why?" We provide new evidence on this question by constructing precise estimates of consumption smoothing across the distribution of liquid assets. This analysis follows in a tradition going back to Zeldes (1989) of measuring heterogeneity in consumption smoothing by asset holdings. This is useful because a key source of theoretical uncertainty is whether consumption sensitivity is concentrated among low-asset households or whether it also extends to high-asset households.

To analyze the role of liquid assets, we construct each household's financial buffer as the ratio of average checking balance to average nondurable consumption in the six months prior to the payroll shock. We normalize assets by nondurable consumption because the key state variable in models of household consumption is the ratio of assets to permanent income (which we proxy for using lagged consumption).[38]

[37]Section 4.4 provides additional evidence consistent with Assumption CE2. It documents that consumption sensitivity varies sharply with liquid assets, even though the effect on time use is presumably similar across the normalized liquid asset distribution.

[38]Since the asset buffer measure is normalized, this does not mechanically imply that low *buffer* households necessarily

We begin with a parsimonious regression model of how assets affect consumption smoothing and then move on to specifications that model the role of assets more flexibly. To integrate data from across the entire distribution of liquid assets in one simple specification, we re-estimate equations (18) and (19), interacting both the instrument and the change in income with liquid assets:

$$\Delta c_{it} = \alpha + \beta \Delta y_{it} + \beta^{asset} \Delta y_{it} a_{it} + \zeta_{it}$$
$$\Delta y_{it} = \phi - \rho \Delta y_{j(-i,t),t+1} - \rho^{asset} \Delta y_{j(-i,t),t+1} a_{it} + \nu_{it}.$$

where a_{it} is each household's normalized asset rank (scaled between 0 and 1 and then centered by subtracting 0.5). This specification means that β captures the causal effect of an income shock at the median of the asset distribution, while $\beta + 0.5\beta^{asset}$ captures the effect at the top of the asset distribution and $\beta - 0.5\beta^{asset}$ captures the effect at the bottom of the asset distribution.

The results are shown in Table 1 column (2). We estimate β^{asset} of -0.48 with a standard error of 0.05. This indicates a sharp and statistically precise negative relationship between consumption sensitivity and a household's liquid asset buffer.[39]

To compare our results to prior estimates, in the remainder of the section we report the marginal propensity to consume (MPC) nondurables by quantiles of assets. We do this by estimating equations (18) and (19) separately for different liquid asset bins, and then scaling our elasticity estimates by the ratio of average nondurable consumption to average income in each bin.[40]

A key source of uncertainty in the prior literature has been whether and to what extent consumption sensitivity to temporary income shocks varies with liquid assets. Although many windfall papers provide convincing evidence of high sensitivity for average households, the extent of heterogeneity in this sensitivity by liquid assets has been difficult to assess. Figure 6a illustrates this statistical uncertainty by reporting liquid asset heterogeneity estimates from three seminal windfall papers (Boehm, Fize, and Jaravel 2025; Kueng 2018; Parker et al. 2013).[41] Estimates from all three of these papers are consistent with an MPC that is falling as liquid assets rise, as in standard models. However, all three papers' estimates are also consistent with an MPC that is constant (e.g., liquid assets are irrelevant for understanding consumption smoothing behavior), and even with an MPC that is rising (e.g., an MPC that is *higher* for high-asset households). Moreover, in addition to the uncertainty in empirical windfall estimates, there is also significant uncertainty about whether there is any liquidity-MPC gradient in the survey literature which asks households to predict their response to hypothetical windfalls (Crossley et al. 2025). It is quite understandable how, given this uncertainty, a range of theoretical models have emerged, some of which feature a large role for liquidity while others feature little role

have low *levels* of checking account balances. Instead, low buffer households are those whose checking account assets are low relative to their typical consumption levels. Nevertheless, there is a strong correlation between asset buffers and asset levels, as can be seen in Table A-10. Furthermore, our key results on the relationship between assets and consumption smoothing continue to hold if we define assets as the level of checking account balances rather than as a buffer ratio (see Tables A-9 and A-11).

[39] Table A-9 shows a similarly sharp and statistically precise negative relationship between assets and consumption sensitivity when assets are measured in levels of checking account balances rather than as a buffer, while Table A-6 shows that this relationship between assets and consumption sensitivity holds for a variety spending definitions.

[40] Tables A-10 and A-12 show the regression results and scaling factors by asset quartile and decile.

[41] For U.S. stimulus payments 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 only to expenditure categories captured in the Nielsen retail panel. More recently Lewis, Melcangi, and Pilossoph (2025) re-examined the 2008 stimulus payments using the same data as Parker et al. (2013) but employing an alternative approach to estimating MPC heterogeneity based on clustering regressions. The paper also finds a statistically imprecise (and, if anything, positive) relationship between liquidity and MPCs.

at all.

A main contribution of our paper is that we make progress towards resolving this uncertainty. By studying a typical source of income variation, we are able to construct our instrument for every household in every month, generating tight statistical precision. Figure 6a compares our estimates to those in the prior literature and demonstrates that the consumption response to firm pay shocks declines steeply as liquid wealth rises. Figure 6b moves beyond the cutoffs in the prior literature and shows even more striking heterogeneity by decile of liquid assets. Moving from the bottom decile to the top decile of the liquid wealth distribution, the MPC falls from 0.41 to 0.02. Moreover, this gradient is precisely estimated. If the spending response to labor income is the same as the response to the windfalls studied in the prior literature, this result resolves much of the statistical uncertainty in the prior literature. In work subsequent to ours, Baker et al. (2023) similarly documents a precise negative MPC gradient out of windfall pandemic stimulus checks.

The finding that MPCs decline sharply with liquid wealth provides a test between alternative theoretical models of the consumption-savings decision. In many benchmark models, liquid assets are a key state variable and sometimes the only state variable (Aiyagari 1994; Carroll 1997; Kaplan and Violante 2014; Laibson et al. 2024). A central prediction of these models is that there is a tight negative correlation between MPCs and liquid assets. However, a number of recent theoretical models—motivated in part by statistical uncertainty in the estimates for MPCs of high-wealth households—have proposed explanations for why even households with substantial liquid wealth might have high MPCs. These include an awareness of future behavioral biases (Lian 2023), bounded rationality (Ilut and Valchev 2023), diagnostic expectations (Bianchi, Ilut, and Saijo 2023), mental accounting (Mijakovic 2024), and anticipatory savings (Campbell and Hercowitz 2019).

Although our paper substantially improves precision regarding the degree to which consumption sensitivity to exogenous temporary income shocks varies with liquidity, there is prior empirical evidence looking outside natural experiments or outside temporary shocks which also points towards a downward-sloping gradient. This evidence can be organized into three groups. First, looking beyond natural experiments, Blundell, Pistaferri, and Preston (2008) and Crawley and Kuchler (2023) find that temporary income changes (which include both endogenous and exogenous changes to income) have larger consumption effects on households with low assets. Second, looking at clearly-identifiable but potentially endogenous income changes, Gelman (2021), Graham and McDowall (2025), and Souleles (1999) find that the consumption effects of predictable tax refunds, whose timing can be chosen by households themselves, are largest among households with low assets. Relatedly, two papers, Hamilton et al. (2024) and Kreiner, Lassen, and Leth-Petersen (2019), show that consumption responses to (endogenously chosen) retirement account withdrawals are largest for households with low liquid assets. Third, Baker (2018) and Higgins (2025) show that low-asset households have a larger consumption response to *permanent* shocks.

One limitation of both our estimates and this prior work is that differences in assets can arise for many reasons. Thus, it may not be the case that exogenous variation in assets will generate the differences in consumption smoothing depicted in Figure 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. For example some unobserved factor (like patience) could cause some households to have both high assets and low MPCs. The best way to assess the role of potential omitted variables in driving the strong liquidity-MPC relationship we observe is by

Figure 6: Heterogeneity in MPCs by Liquid Assets

(a) Comparison to Prior Literature

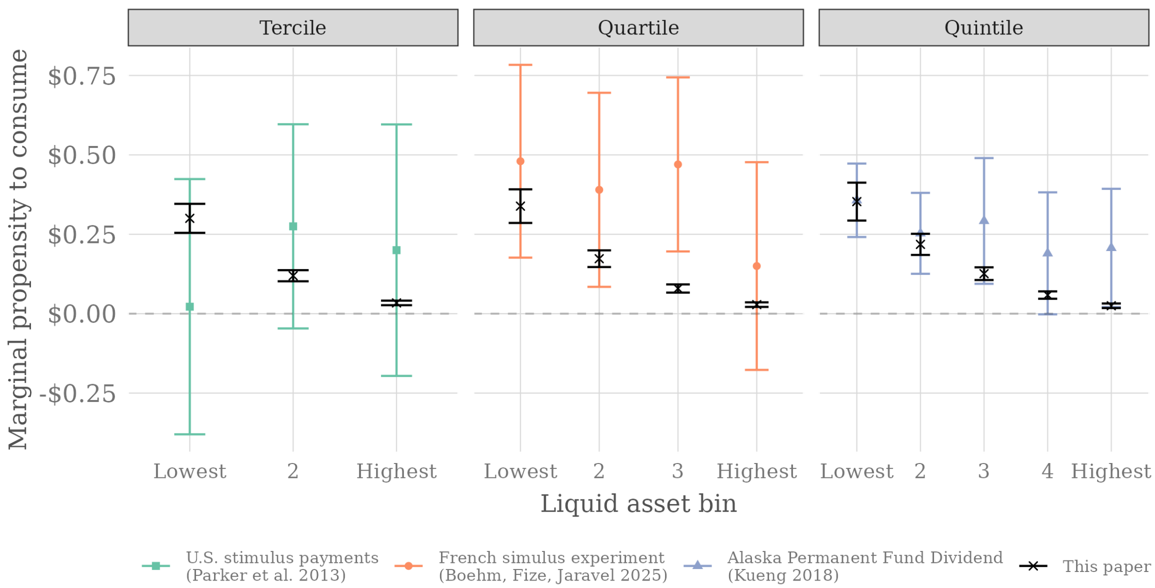

(b) MPC by Decile of Liquid Assets

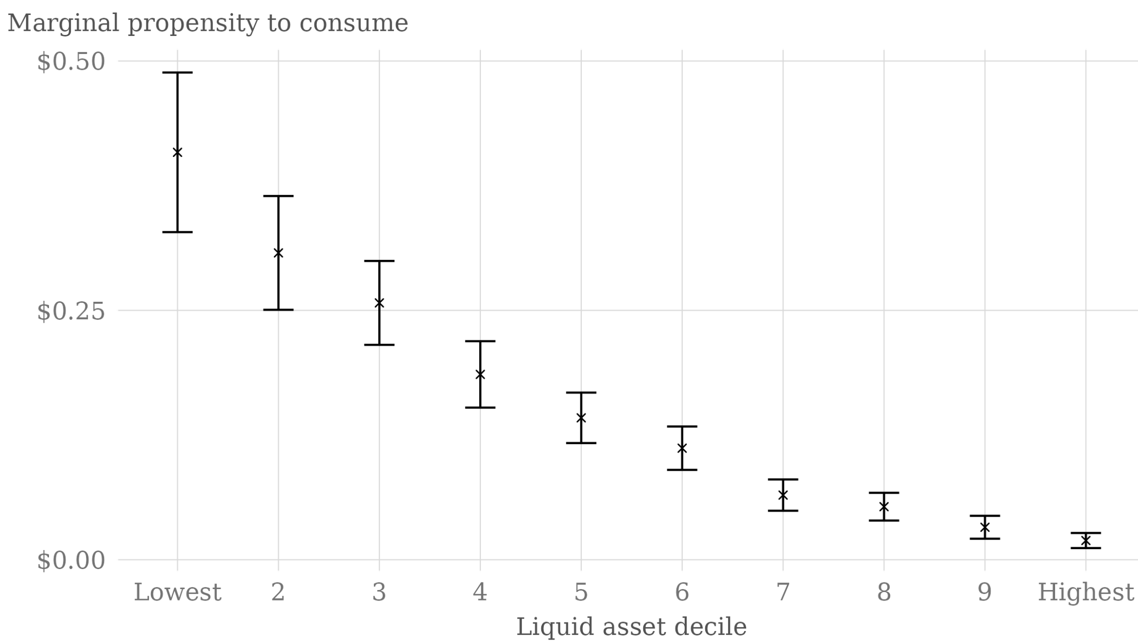

Notes: This figure shows the relationship between MPCs and liquid assets. It reports both point estimates and 95 percent confidence intervals. Panel (a) compares the estimates from our research design to several estimates in the prior literature. Parker et al. (2013) groups households by terciles of liquid assets in levels, Boehm, Fize, and Jaravel (2025) groups by quartiles of net liquid assets in levels, and Kueng (2018) groups by quintiles of liquid asset buffer. For comparison, we report the MPC from our research design by tercile, quartile, and quintile of liquid asset buffer. We use the main specification in each paper: Kueng (2018) and Parker et al. (2013) report heterogeneity in quarterly nondurable MPCs, Boehm, Fize, and Jaravel (2025) reports heterogeneity in monthly card and cash MPCs, and our paper reports heterogeneity in monthly nondurable MPCs. Figure A-4 compares heterogeneity using quarterly MPC specifications from all four papers. Panel (b) shows heterogeneity in our estimates of the monthly nondurable spending MPC by decile of liquid asset buffer. Table A-13 reports point estimates and standard errors for the monthly and quarterly MPCs by liquidity quantile using our research design.

leveraging an instrument for liquidity. Unfortunately we do not have such an instrument in this paper. One prior paper that uses such an instrument is Ganong et al. (2024). Based on a combination of reduced-form empirical evidence and a structural model, that paper concludes that about two-thirds of the correlation between MPCs and a household's current circumstances is caused by liquidity, and one-third is caused by selection.

The evidence in this paper is agnostic about the source of the liquidity-MPC correlation. It is consistent with models such as Kaplan and Violante (2014) and Laibson et al. (2024) in which low liquidity arising in part from negative income shocks *causes* high sensitivity. It is also consistent with models where an alternative channel that is correlated with liquidity (e.g., inter-household transfers in Chiteji and Hamilton 2002, or permanent heterogeneity as in Aguiar, Bils, and Boar 2021, or Epper et al. 2020) contributes to both low liquidity and high MPCs. In either case, the key moment that a model must match in order to be consistent with the empirical evidence in this paper is that the households who typically find themselves with low liquidity are much more sensitive than the households who typically find themselves with high liquidity.

4.5 Consumption Smoothing for Alternative Broader Instrument Definitions and the Role of Predictability

Our preferred specification discussed above takes several steps to isolate exogenous, unpredictable, and transitory shocks to income. In this section, we show how increasingly broad definitions of the instrument—for which these three conditions might not be satisfied—are correlated with changes in consumption. To do so, we re-estimate equations (18) and (19) using four different definitions of the coworker instrument. The benefit of looking at broader definitions of the instrument is to assess the extent to which any of our key results hinge on these steps, and whether workers appear to respond differently to different types of income volatility.

Before looking at the results for consumption, Figure 7 shows the dynamic relationship between household labor income and these broader definitions of the coworker instrument. We re-estimate equation (17) using total coworker monthly pay, coworker pay per paycheck, and coworker seasonally adjusted pay per paycheck as the instrument. In all three cases, the key regressor captures a contemporaneous change from $t^* - 1$ to t^*. Although the definition of the instrument on the right hand side of equation (17) changes with each specification, the left-hand-side endogenous regressor is always total labor income at the household level.

The top-left panel of Figure 7 shows the response of worker income to the broadest definition of the instrument: total coworker monthly pay. Three patterns stand out in the figure. First, there is a sharp decrease in pay from $t^* - 2$ to $t^* - 1$. This is what we would expect in a model where there are transitory shocks to income: an increase from $t^* - 1$ to t^* will often be the predictable reversion of a negative temporary shock at $t^* - 1$. Second, there is a spike in pay at $t^* + 12$ and a smaller but still obvious spike at $t^* - 12$, indicating that some of the coworker pay shock is attributable to annually recurring seasonality. In further analysis using PayrollCompany data, we find that this reflects—in order of decreasing importance—annual bonuses, annual wage increases, and hours fluctuations that occur in the same month each year.

Third, there is evidence of periodicity of coworker pay changes at roughly a quarterly frequency. This pattern arises because about three-quarters of U.S. workers are paid either every week or every fortnight. Workers who are paid every week get four checks in eight months of the year and an extra

Figure 7: Income Event Study Around Realizations of Different Coworker Instruments

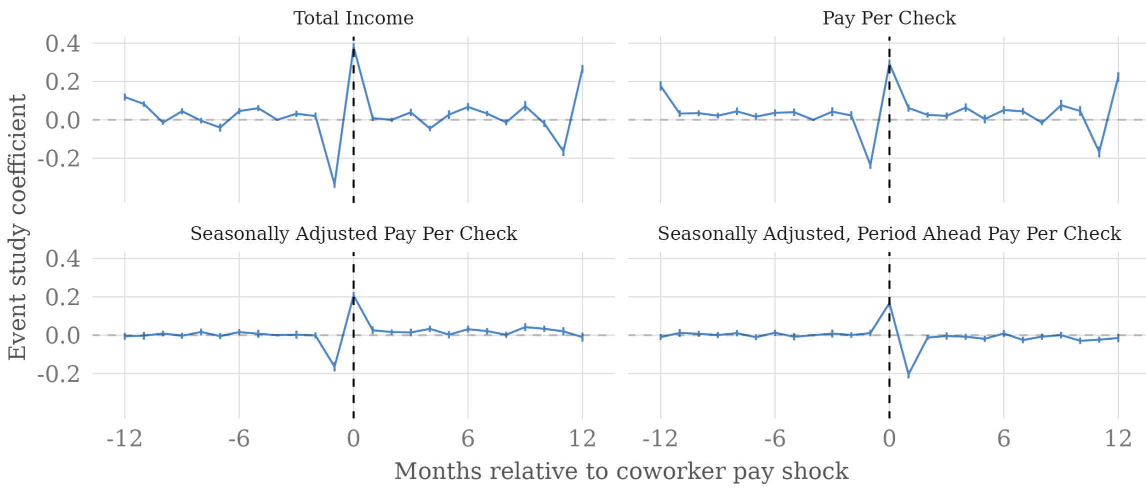

Notes: This figure shows the dynamic response of worker income to broader definitions of the coworker instrument. Each panel shows the time path of event study-like coefficients ρ_k from estimating equation (17) using various specifications for the continuous instrument D_{i,t^*}. The top-left panel presents the response to total coworker monthly pay. This corresponds to the specification in column (5) of Table 1. The top-right panel refines the instrument by dividing total monthly pay by the number of checks received, corresponding to column (4) of Table 1. The bottom-left panel further adjusts the instrument and the endogenous regressor for predictable annual seasonality, corresponding to column (3) of Table 1. Finally, the bottom-right panel isolates an unpredictable transitory shock using the seasonally adjusted pay per paycheck for one period ahead. This corresponds to columns (1) and (2) of Table 1, and is our main specification. The black dashed line at time 0 marks the timing of coworker pay instrument realization and is included to highlight the reversal of the pattern when moving to the period-ahead specification. All panels include uniform 95 percent confidence intervals.

fifth check in four months of the year; a similar pattern occurs for workers who are paid every fortnight. At least three prior consumption-smoothing papers have specifically studied the effects of variation induced by the timing of paychecks (Baugh and Wang 2021; Baugh and Correia 2022; Zhang 2022).

This broadest version of the instrument therefore mixes several different sources of income variation. In addition to the unpredictable transitory shocks we are seeking to isolate, it also includes predictable variation from the pay calendar, predictable annual seasonal variation, and the predictable reversion of the prior period's transitory shock.

Relative to this broad definition of the instrument, we take three steps to isolate an unpredictable and temporary shock to workers' income in our main specification. To eliminate changes in coworker income driven by pay schedules, we re-define the instrument as coworker total monthly pay divided by the number of checks received by coworkers. The top-right panel shows that this procedure eliminates the serial correlation of changes in coworker pay in most months, but the changes at $t^* + 12$ and $t^* - 12$ still remain.

Next, to address recurring annual seasonality, we re-define the instrument as the instrument from the prior step minus the predicted change based on coworker pay twelve months beforehand. This seasonal adjustment procedure is described in more detail in Appendix B.1. The bottom-left panel shows that this procedure eliminates the correlation between the change in coworker pay from $t^* - 1$ to t^* and the changes at $t^* + 12$ and $t^* - 12$. However, there is still a clear dip in pay from $t^* - 2$ to $t^* - 1$, leading to a predictable reversion of that transitory shock between periods $t^* - 1$ and t^*.

Furthermore, there is a subtle indication in the post-period that this instrument also picks up a small amount of permanent shocks, since the average level in the post-period is slightly above the average level in the pre-period.

Finally, to isolate a fully unpredictable, transitory shock to income, the bottom-right panel shows the results for using the seasonally adjusted pay per paycheck for one period ahead as the instrument. This is the instrument that we use in our main specifications in the paper, reflected in columns (1) and (2) of Table 1. Relative to the bottom-left panel, this change simply shifts the coefficients one period forward and multiplies them by negative one. The increase from period $t^* - 1$ to t^* no longer includes the reversion of the prior period's shock. The fact that the patterns do not change beyond this is consistent with what we would expect in a stationary environment. Further, the patterns here mirror what we would expect to see based on the stylized example in Figure 1a.

Columns (3), (4), and (5) of Table 1 show two key results from estimating consumption sensitivity using the broader definitions of the instrument. First, the results in the prior sections of the paper—that consumption is sensitive to coworker income shocks (row one) and that there is a strong asset gradient (row two)—continue to hold with these broader definitions of the instrument. Second, the elasticity of consumption to income is consistently smaller across the three broader definitions of the instrument than it is under our main specification in column (2).[42] All three broader definitions rely on variation in coworker income which is partially predictable, for the various reasons discussed above. Because our methodology is able to capture the consumption response to exogenous income shocks with varying degrees of predictability for the same households, it offers more direct evidence on the role of advance information than has been available in the prior literature. We find larger responses to unpredictable shocks than predictable ones, consistent with many theoretical models of consumption.

It is also possible to ignore the coworkers completely and instead ask how consumption is correlated with *worker* income changes using each of the above regressions. Of course, such an analysis will not capture the causal impact of income on consumption because worker-level income changes include an individual worker's endogenous labor supply decisions. However, this exercise is nevertheless informative about the dynamics of typical monthly income changes. Figure A-5 repeats Figure 7 but substituting a worker's monthly income, a worker's monthly pay per paycheck, and a worker's seasonally adjust pay per paycheck for the coworker components. The figure shows that the dynamic patterns of income with respect to a worker's own change in income from $t^* - 1$ to t^* are very similar to the dynamics with respect to coworker income in Figure 7.[43] This bolsters the claim that the coworker instrument is indeed capturing "typical" income fluctuations from the worker's perspective. Coworker shocks have the same dynamics as worker-level shocks. The key distinction between the two is that the coworker shocks are more plausibly exogenous. This suggests that any difference in the consumption response to worker shocks relative to coworker shocks is likely due to endogeneity rather than to any differences in the other characteristics of the underlying shocks.

Table A-14 shows that the elasticities estimated using worker income are generally smaller than the corresponding estimates using changes in coworker income. This suggests that workers' endogenous changes in income are associated with a smaller change in consumption than the exogenous changes in income captured by the coworker instrument. The table also shows that assets continue to strongly

[42] Row (1) of columns (2) through (5) reports elasticity estimates for the median-asset household using each specification. Results are similar if we drop the checking asset control and report the average elasticity instead.

[43] Permanent income shocks seem to play a slightly bigger role for individual workers than they do for firm-wide shocks. This can be seen in, for example, the bottom-left panel of Figure A-5 where the estimates are more clearly above zero, albeit slightly, for periods $t^* + 1$ and forward.

predict the MPC, consistent with the findings in Blundell, Pistaferri, and Preston (2008) and Crawley and Kuchler (2023) which use annual data and a similar worker-based approach that mixes exogenous and endogenous income changes.

5 Welfare Costs of Volatility

We have documented that consumption is sensitive to typical labor income fluctuations. In this section, we ask whether this sensitivity is economically meaningful. To evaluate this question, we calculate the cost of temporary income volatility implied by our reduced-form estimates. We combine our empirical estimates with a simple, widely used framework based on Lucas (1987). We find that temporary income volatility has a meaningful welfare cost. We then discuss how our finding that households are sensitive even to frequently occurring income fluctuations with meaningful welfare costs helps to further distinguish between alternative explanations for consumption sensitivity.

5.1 Framework for Calculating Welfare Cost of Income Volatility

We extend our model of income and consumption from Section 2, and evaluate the welfare cost of income shocks in the spirit of Lucas (1987). We model log income in each period as the sum of three orthogonal components: a permanent idiosyncratic component (z_t) and a transitory idiosyncratic component (ε_t), as in equation (3), and, also, a predictable, life-cycle/seasonal component (Γ_t):

$$\begin{aligned} y_t &= \Gamma_t + z_t + \varepsilon_t \\ z_t &= z_{t-1} + \eta_t, \end{aligned} \tag{20}$$

where ε_t and η_t are i.i.d. and normally distributed random variables with variances σ_ε^2 and σ_η^2. These shocks are inclusive of the firm-level shocks discussed above. We assume $\mathbb{E}(e^\varepsilon) = \mathbb{E}(e^\eta) = 1$.

Generalizing equation (4), let consumption growth be a function of the different components of income:

$$\Delta c_t = \Delta \tilde{c}_t \left(\Gamma_t, z_t, \zeta_t\right) + \beta_\varepsilon \varepsilon_t - \beta_\varepsilon \varepsilon_{t-1} \tag{21}$$

where the error term ζ_t is orthogonal to (Γ, z, ε). This implies that the *log* of consumption is:

$$c_t = \tilde{c}_t(\Gamma_t, z_t, \zeta_t) + \beta_\varepsilon \varepsilon_t, \tag{22}$$

Kaplan and Violante (2010) shows that a function where consumption is linear in temporary income shocks is a good approximation to predictions from a structural life-cycle model.[44]

We now derive a formula, similar to that of Lucas (1987), that captures the welfare gain from eliminating the variation in income due to transitory shocks. From equation (22) above, we can write the *level* of consumption as:

$$C_t = e^{\tilde{c}_t} e^{\beta_\varepsilon \varepsilon_t} = \tilde{C}_t e^{\beta_\varepsilon \varepsilon_t}, \tag{23}$$

[44] Relative to the consumption growth in equation (4), the additional term $-\beta_\varepsilon \varepsilon_{t-1}$ limits responses to transitory shocks to one period. The transitory component of this consumption function is a special case of that in our dynamic model in Appendix J, equation (J.1), where $\beta_{\varepsilon,1} = -\beta_{\varepsilon,0}$ and $\beta_{\varepsilon,2} = 0$.

where $\tilde{C}_t = e^{\tilde{c}_t}$ is the level of consumption, net of the impact of 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 the variation in income due to those shocks shut off, while preserving the mean of income.[45] Assuming constant relative risk aversion (CRRA) preferences with risk parameter γ, we have:

$$\mathbb{E}\left(\sum_{t=1}^{T}\delta^t\frac{((1+\lambda)C_t)^{1-\gamma}}{1-\gamma}\right) \equiv \mathbb{E}\left(\sum_{t=1}^{T}\delta^t\frac{\tilde{C}_t^{1-\gamma}}{1-\gamma}\right) \tag{24}$$

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

$$\lambda \simeq \left(\beta_\varepsilon + (\gamma-1)\,\beta_\varepsilon^2\right)\frac{\sigma_\varepsilon^2}{2} \tag{25}$$

When $\beta_\varepsilon = 1$, this result reduces to the Lucas (1987) result that $\lambda \simeq \gamma\frac{\sigma_\varepsilon^2}{2}$.

5.2 Welfare Cost Estimates

In order to apply this formula in our empirical setting, we need values of β_ε, σ_ε^2, and γ. For the elasticity of consumption to transitory shocks β_ε we draw on our main estimate from the first column of Table 1. For the variance of transitory shocks σ_ε^2, we show in Appendix K.4 how our general estimation method can also yield estimates of this variance. Importantly, we only include the variance of *unpredictable* shocks for this welfare cost estimate, since the transitory shock in equation (20) is unpredictable. Hence, we measure this variance after applying our adjustments to remove predictable income changes arising from paycheck timing and recurring annual seasonality.[46] Conceptually, we would like to capture a measure of the variance of unpredictable transitory income fluctuations arising from both firmwide shocks and worker-specific shocks. We therefore use all worker-level income fluctuations, rather than just firmwide fluctuations, to measure this variance.[47] This procedure yields a variance of 0.05.

Table 2 reports the welfare gain of shutting off variance due to transitory shocks for different combinations of parameters. We produce two sets of results, which allow for varying degrees of dynamics in the response to a transitory shock. In the first row of Table 2, we assume that the consumption response to a transitory shock only lasts one month, as modeled above in equation (21). In this "Monthly" model, we get a range of welfare costs between 0.55 and 0.91 percent of lifetime consumption. Moving across columns, we see that this cost is increasing in the risk aversion parameter γ, which we vary from one to four. In Appendix K.2 we adapt our welfare measure to a setting with a more dynamic consumption response that may last up to three months, consistent with our findings in Section 4.3. In this "Quarterly" or "Dynamic" model, we find even larger welfare costs, ranging from 1.05 to 1.57 percent of lifetime consumption, allaying concerns over high-frequency intertemporal

[45] This is equivalent to setting $\varepsilon_t = 0$ in every period.

[46] In Appendix K.3, we show how one might extend our welfare measure in the case where *predictable* transitory shocks are also welfare relevant. This would result in even larger welfare costs.

[47] Worker-level idiosyncratic fluctuations can arise from sources that we would want to include in this calculation—like a health shock or a shock to the firm's demand for one worker's labor—as well as sources that we would not want to include, like leisure. Ganong et al. (2025) calculates that unpaid vacation accounts for only a very small share of month-to-month earnings volatility, suggesting the bias from simply using all worker-level fluctuations is small.

substitution driving these results. This is because we find empirically that cumulative consumption impacts grow, rather than shrink, over subsequent months.

Table 2: Welfare Gain of Eliminating Transitory Shocks

	Welfare gain: λ			
	$\gamma = 1$	$\gamma = 2$	$\gamma = 3$	$\gamma = 4$
Monthly	0.55%	0.67%	0.79%	0.91%
Quarterly (Dynamic)	1.05%	1.23%	1.40%	1.57%

Notes: This table reports the results of an exercise in the spirit of Lucas (1987) by calculating the welfare gain from eliminating unpredictable transitory income shocks. The welfare gain in the first row is calculated using the approximation in equation (25) and is expressed as a fraction of lifetime consumption. This row only incorporates the first month's consumption impact of a transitory income shock. The second row calculates the cost allowing the consumption impact to last up to three months. See Appendix K.2 for details on this dynamic calculation.

In all cases, our estimates of the welfare cost exceed 0.5 percent of lifetime consumption, and are consistently above one percent when the full three-month consumption response is incorporated into the calculation. These estimates may be an upper bound on the true cost of transitory fluctuations for two reasons. First, as is the case with any calculations within the standard Lucas framework, they assume that temporary income fluctuations do not directly affect utility (e.g., through changes in hours worked). Second, they assume that the elasticity of consumption with respect to income for all unpredictable temporary labor income changes is the same as the Local Average Treatment Effect we estimate using the firm component of unpredictable temporary changes. Nevertheless, relative to the benchmark view in Lucas that a cost of 0.5 percent of lifetime consumption is "large," these estimates suggest there could be a substantial welfare cost from transitory income volatility. Although Lucas conjectured that the welfare cost of transitory income volatility might be large if his formula were applied to individual volatility rather than aggregate volatility, we are unaware of any prior estimate of these costs.[48]

In Appendix K.5 we complement the calculations from our statistical model using the kind of modern two-asset structural model underlying the low-liquidity interpretation of consumption sensitivity to windfalls (e.g., Kaplan and Violante 2022). We show that such a model can indeed match the sensitivity to typical labor income volatility that we document and also features a negative gradient in this sensitivity with respect to liquid assets. With a CRRA parameter of $\gamma = 1$ and a transitory shock variance of $\sigma_\varepsilon^2 = 0.05$, the welfare cost of transitory income volatility in this model is 0.41 percent, similar to the welfare cost we calculate using the simple statistical model in Table 2 and using similar values for γ and σ_ε^2.

5.3 Implications for Near-Rational Interpretations of Consumption Sensitivity

While our finding of a sharp liquidity gradient provided one test between alternative models of consumption, our finding that households are sensitive even to frequently occurring labor income fluctuations with meaningful welfare costs provides another such test. In particular, this finding is inconsistent with some models that treat consumption-sensitivity as "near-rational" behavior.

[48]De Nardi, Fella, and Paz-Pardo (2020) and Hai, Krueger, and Postlewaite (2020) document a significant welfare gain from eliminating all income volatility (both temporary shocks and permanent shocks). In complementary work, Constantinides (2025) documents a significant welfare gain from eliminating all permanent idiosyncratic consumption shocks.

The logic for the near-rationality argument can be seen by examining the expression for welfare costs in equation (25). This cost is proportional to the product of sensitivity β and the variance of the shocks σ_ε^2. A line of argument dating back to Cochrane (1989) argues that when this product—which captures the upper bound on the welfare loss from setting consumption sub-optimally—is small, it is hard to learn much about individuals' true consumption model. Instead, to learn about the true model of consumption it is necessary to look at welfare-relevant shocks. By definition, σ_ε^2 is small for unusual windfalls; even if $\beta = 1$, $\gamma\sigma_\varepsilon^2$ is still small. This line of reasoning suggests that sensitivity to windfalls may simply reflect a low-cost deviation from the true "rational" model. Under one common version of this reasoning, households would smooth frequently occurring income shocks, in line with Friedman's Permanent Income Hypothesis, because the cost of failing to smooth such shocks is large (see e.g., Fuchs-Schündeln and Hassan 2016).

In contrast, we find that households are sensitive even to frequently occurring labor income volatility. Although our study is similar to prior work in finding a sensitivity estimate (β_ε) of around 0.2, our study differs in that it analyzes a type of income variation whose variance σ_ε^2 is much larger.[49] Finding a high degree of consumption sensitivity even in this setting thus favors the low-liquidity interpretation of consumption sensitivity over this version of the near-rationality interpretation.

While our evidence is inconsistent with near-rationality theories that conjecture that households would smooth consumption when welfare costs are *high*, our evidence is not informative about two other related types of near-rationality theories. The first is described in the study by Kueng (2018) of households receiving predictable annual payouts from the Alaska Permanent Fund. That paper finds that households with a *low* welfare cost from failing to smooth are especially sensitive. Because the coworker pay shock that we study results in similar proportional shocks across most groups of households, we have little power to study heterogeneity in the response by the size of the welfare loss. We can only document that when the shocks are welfare-relevant, households are still sensitive; sensitivity isn't confined to low-stakes environments.

Second, a new and intriguing strand of the near-rationality literature focuses on the emergence of heuristics or "quick-fixes" in explaining within-person heterogeneity in MPCs. Andre et al. (2025) finds evidence that households are more likely to have extreme responses (MPCs of zero or one) to small shocks than to large shocks. Unfortunately, we are unable to measure individual-level MPCs with our method and so we are unable to investigate this type of behavior.

6 Conclusion

In this paper, we develop a methodology to estimate the effect of exogenous, transitory, and unpredictable shocks to labor income on consumption. Our approach combines elements from two distinct consumption smoothing literatures—one that looks at quasi-experiments and another that relies on

[49] One other study which analyzes the consumption response to large pay fluctuations is Browning and Collado (2001, henceforth BC). Our finding that consumption is sensitive to frequently occurring labor income fluctuations is superficially at odds with the finding of no sensitivity in that paper; however, further inspection suggests the estimates may be mutually consistent. BC compare the quarterly consumption of n = 341 workers who are paid in 12 monthly installments to n = 1877 workers who are paid in 14 installments, with double installments in June and December. The standard deviation of quarterly consumption in that survey is 29 percent, and the change in quarterly income from the installments is about 25 percent, so the design would only detect consumption fluctuations if the elasticity is 0.13 or greater (assuming a 95% confidence interval). The specification which most relies on predictable income variation in this paper yields an elasticity of 0.11 for a household at the median of the liquidity buffer distribution. Thus, the design in BC is both unable to detect the elasticity to predictable income variation that we estimate in this paper and also unable to reject the null hypothesis of full consumption smoothing.

covariance restrictions. We show that it is possible to get close to the "best of both worlds," using the covariance restrictions to extract transitory shocks, using firm pay for exogenous variation, and using diagnostic tests of identification assumptions similar to those used in event-study literature.

We find that household consumption is highly sensitive to monthly labor income shocks. Furthermore, consumption is most sensitive for households with low liquidity and almost unchanged for households with high liquidity. Our findings bolster the assumptions of liquidity-focused models of the consumption-savings decision and are at odds with the predictions of some behavioral models of consumption. Finally, embedding our findings in a simple framework implies that temporary income volatility has a substantial welfare cost.

These findings suggest at least two avenues for future research. First, although our key finding that consumption sensitivity varies sharply with liquid assets bolsters low-liquidity interpretations of consumption sensitivity, it would be useful to further investigate *why* low-liquidity households are especially sensitive. Our evidence is consistent both with models where low liquidity arising in part from negative income shocks causes high sensitivity, and also with models where an alternative channel that is correlated with liquidity contributes to both low liquidity and high MPCs. Finding more instruments for liquidity and more tags for persistent household characteristics would help distinguish between these competing explanations.

Second, while this paper focuses on the consumption response to *transitory* shocks, it would be interesting to further investigate the pass-through of *permanent* shocks. Many theories suggest that consumption should be more responsive to permanent shocks than to transitory shocks. The methodology we develop in this paper to isolate exogenous transitory shocks can be adapted to isolate exogenous permanent shocks. We think this is an exciting avenue for future research.

References

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

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

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

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

Andre, Peter, Joel P. Flynn, George Nikolakoudis, and Karthik A. Sastry. 2025. "Quick-Fixing: Near-Rationality in Consumption and Savings Behavior." Working Paper 33464, National Bureau of Economic Research.

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

Attanasio, Orazio P. and Guglielmo Weber. 1995. "Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey." *Journal of Political Economy*, 103(6): 1121–1157.

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.

Auclert, Adrien, Matthew Rognlie, and Ludwig Straub. 2024. "The Intertemporal Keynesian Cross." *Journal of Political Economy*, 132(12): 4068–4121.

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., Robert A. Farrokhnia, Steffen Meyer, Michaela Pagel, and Constantine Yannelis. 2023. "Income, Liquidity, and the Consumption Response to the 2020 Economic Stimulus Payments." *Review of Finance*, 27(6): 2271–2304.

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

Baugh, Brian and Filipe Correia. 2022. "Does Paycheck Frequency Matter? Evidence from Micro Data." *Journal of Financial Economics*, 143(3): 1026–1042.

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

Bianchi, Francesco, Cosmin Ilut, and Hikaru Saijo. 2023. "Diagnostic Business Cycles." *Review of Economic Studies*, 91(1): 129–162.

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

Blundell, Richard, Hamish Low, and Ian Preston. 2013. "Decomposing Changes in Income Risk using Consumption Data." *Quantitative Economics*, 4(1): 1–37.

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

Board of Governors of the Federal Reserve System. 2019. "Survey of Consumer Finances, 2019." *https://www.federalreserve.gov/econres/scf_2019.htm*, Accessed: March 7, 2025.

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

Boehm, Johannes, Etienne Fize, and Xavier Jaravel. 2023. "Five Facts about MPCs: Evidence from a Randomized Experiment." *Working Paper, HAL Id: hal-04305929.*

Boehm, Johannes, Etienne Fize, and Xavier Jaravel. 2025. "Five Facts about MPCs: Evidence from a Randomized Experiment." *American Economic Review*, 115(1): 1–42.

Brewer, Mike, Nye Cominetti, and Stephen P. Jenkins. 2025. "What Do We Know About Income and Earnings Volatility?" *Review of Income and Wealth*, 71(2): e70013.

Broda, Christian and Jonathan A. Parker. 2014. "The Economic Stimulus Payments of 2008 and the Aggregate Demand for Consumption." *Journal of Monetary Economics*, 68: S20–S36.

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

Bureau of Labor Statistics. 2016–2023. "Consumer Price Index for All Urban Consumers: All Items in U.S. City Average (CPIAUCSL)." *https://fred.stlouisfed.org/series/CPIAUCSL*, Accessed: March 7, 2025.

Campbell, Jeffrey R. and Zvi Hercowitz. 2019. "Liquidity Constraints of the Middle Class." *American Economic Journal: Economic Policy*, 11(3): 130–155.

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

CDC, National Center for Health Statistics. 1999–2019. "Underlying Cause of Death, 1999–2019." *https://wonder.cdc.gov/ucd-icd10.html*, Accessed: March 7, 2025.

Chiteji, Ngina S. and Darrick Hamilton. 2002. "Family Connections and the Black-White Wealth Gap among Middle-Class Families." *Review of Black Political Economy*, 30(1): 9–28.

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

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

Commault, Jeanne. 2022. "Does Consumption Respond to Transitory Shocks? Reconciling Natural Experiments and Semistructural Methods." *American Economic Journal: Macroeconomics*, 14(2): 96–122.

Constantinides, George M. 2025. "Welfare Costs of Idiosyncratic and Aggregate Consumption Shocks." *Review of Asset Pricing Studies*, 15(2): 103–120.

Crawley, Edmund and Andreas Kuchler. 2023. "Consumption Heterogeneity: Micro Drivers and Macro Implications." *American Economic Journal: Macroeconomics*, 15(1): 314–41.

Crossley, Thomas F., Paul Fischer, Peter Levell, and Hamish Low. 2025. "Eliciting the Marginal Propensity to Consume in Surveys." *Working Paper.*

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

De Nardi, Mariacristina, Giulio Fella, and Gonzalo Paz-Pardo. 2020. "Nonlinear Household Earnings Dynamics, Self-Insurance, and Welfare." *Journal of the European Economic Association*, 18(2): 890–926.

Druedahl, Jeppe. 2021. "A Guide on Solving Non-convex Consumption-Saving Models." *Computational Economics*, 58(3): 747–775.

Druedahl, Jeppe, Michael Graber, and Thomas H. Jørgensen. 2023. "High Frequency Income Dynamics." *Working Paper.*

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

Etheridge, Ben. 2015. "A Test of the Household Income Process Using Consumption and Wealth Data." *European Economic Review*, 78: 129–157.

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

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

Fuchs-Schündeln, N. and Tarek A. Hassan. 2016. "Natural Experiments in Macroeconomics." In *Handbook of Macroeconomics*. Vol. 2, ed. John B. Taylor and Harald Uhlig, 923–1012, Amsterdam: Elsevier.

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

Ganong, Peter, Fiona Greig, Pascal Noel, Daniel M. Sullivan, and Joseph Vavra. 2024. "Spending and Job-Finding Impacts of Expanded Unemployment Benefits: Evidence from Administrative Micro Data." *American Economic Review*, 114(9): 2898–2939.

Ganong, Peter, Pascal Noel, Christina Patterson, Joseph Vavra, and Alexander Weinberg. 2025. "Earnings Instability." *Working Paper.*

Gelman, Michael. 2021. "What Drives Heterogeneity in the Marginal Propensity to Consume? Temporary Shocks vs Persistent Characteristics." *Journal of Monetary Economics*, 117: 521–542.

Graham, James and Robert A. McDowall. 2025. "Mental Accounts and Consumption Sensitivity across the Distribution of Liquid Assets." *American Economic Journal: Macroeconomics*, forthcoming.

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

Guvenen, Fatih, Fatih Karahan, Serdar Ozkan, and Jae Song. 2014. "Inferring Income Processes from Earnings and Consumption Data." *Econometrica*, 82(6): 2085–2129.

Hai, Rong, Dirk Krueger, and Andrew Postlewaite. 2020. "On the Welfare Cost of Consumption Fluctuations in the Presence of Memorable Goods." *Quantitative Economics*, 11(4): 1177–1214.

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

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

Hamilton, Steven, Geoffrey Liu, Jorge Miranda-Pinto, and Tristram Sainsbury. 2024. "A $100,000 Marshmallow Experiment: Withdrawal and Spending Responses to Early Retirement-Savings Access." *Working Paper.*

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.

Higgins, Brian. 2025. "Differential Effects of LTV and LTI Limits on Mortgage Choices." *Working Paper.*

Ilut, Cosmin and Rosen Valchev. 2023. "Economic Agents as Imperfect Problem Solvers." *Quarterly Journal of Economics*, 138(1): 313–362.

Jappelli, Tullio and Luigi Pistaferri. 2010. "The Consumption Response to Income Changes." *Annual Review of Economics*, 2: 479–506.

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

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

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

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

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

Kreiner, Claus Thustrup, David D. Lassen, and Søren Leth-Petersen. 2019. "Liquidity Constraint Tightness and Consumer Responses to Fiscal Stimulus Policy." *American Economic Journal: Economic Policy*, 11(1): 351–79.

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

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

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

Lachowska, Marta, Alexandre Mas, Raffaele Saggio, and Stephen A. Woodbury. 2023. "Do Firm Effects Drift? Evidence from Washington Administrative Data." *Journal of Econometrics*, 233(2): 375–395.

Laibson, David, Peter Maxted, and Benjamin Moll. 2022. "A Simple Mapping from MPCs to MPXs." Working Paper 29664, National Bureau of Economic Research.

Laibson, David, Sean Chanwook Lee, Peter Maxted, Andrea Repetto, and Jeremy Tobacman. 2024. "Estimating Discount Functions with Consumption Choices over the Lifecycle." *Review of Financial Studies*, forthcoming.

Lamadon, Thibaut, Magne Mogstad, and Bradley Setzler. 2022. "Imperfect Competition, Compensating Differentials, and Rent Sharing in the US Labor Market." *American Economic Review*, 112(1): 169–212.

Lewis, Daniel, Davide Melcangi, and Laura Pilossoph. 2025. "Latent Heterogeneity in the Marginal Propensity to Consume." Working Paper 32523, National Bureau of Economic Research.

Lian, Chen. 2023. "Mistakes in Future Consumption, High MPCs Now." *American Economic Review: Insights*, 5(4): 563–81.

Lucas, Robert E. 1987. *Models of Business Cycles.* Oxford: Wiley-Blackwell.

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

Massenot, Baptiste and Thierry Foucault. 2022. "The Pain of Paying: Effects on Consumption, Borrowing, and Welfare." *Review of Financial Studies*, 36(3): 975–1003.

Mijakovic, Andrej. 2024. "Marginal Propensities to Consume with Mental Accounting." *Working Paper.*

Mitchell, Olivia S. and John W. R. Phillips. 2006. "Social Security Replacement Rates for Alternative Earnings Benchmarks." *Working Paper.*

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

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.

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

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

Stephens, Jr., Melvin and Takashi Unayama. 2019. "Estimating the Impacts of Program Benefits: Using Instrumental Variables with Underreported and Imputed Data." *Review of Economics and Statistics*, 101(3): 468–475.

TD Bank. 2014. "TD Bank Survey Finds Many Couples Maintain Separate Bank Accounts." TD Bank Technical Report.

US Census Bureau. 2017–2023. "Current Population Survey Annual Social and Economic Supplement, 2017–2023." *https://www.census.gov/data/datasets/time-series/demo/cps/cps-asec.html*, Accessed: March 7, 2025.

US Census Bureau and FDIC. 2017–2023. "Current Population Survey: Unbanked/Underbanked Supplement, 2017–2023." *https://www.census.gov/data/datasets/time-series/demo/cps/cps-supp_cps-repwgt/cps-banked.html*, Accessed: March 7, 2025.

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

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

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

Appendix A Appendix Figures and Tables

Figure A-1: Distribution of Income and Checking Balance in Public Data vs Analysis Sample

(a) Income

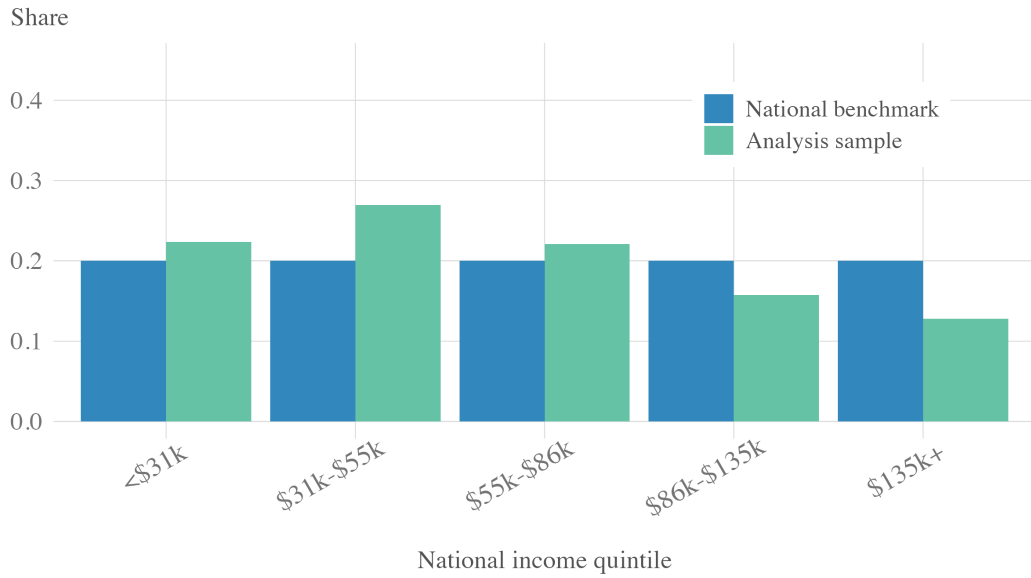

(b) Checking Account Balance

Notes: This figure shows the distribution of household labor income and checking account balance in the U.S. as a whole (blue), and in our analysis sample (green). The bins are chosen to capture the quintiles of the national distribution. Thus, by construction, the blue bars are all equal to 20 percent. In panel (a) the national income cutoffs are calculated using the CPS Annual Social and Economic Supplement from 2018 to 2023, where income is measured for the prior calendar year (2017 to 2022). We restrict to CPS households with positive household labor income. In the bank sample we increase measured labor income by 15% to account for missed income in the form of paper checks and prepaid debit cards, as in Ganong and Noel (2019), and we also follow the method from Ganong and Noel (2019) to rescale the net labor income we observe in bank data into pre-tax dollars for comparison to gross labor income reported in the CPS. Income from both the bank sample and the CPS is inflated to 2019 dollars using the Consumer Price Index for All Urban Consumers (CPI-U). In panel (b) the national checking account cutoffs are calculated using the 2019 and 2022 SCF, restricted to households with positive labor income. We define primary checking account balance as the sum of balances at the institution that respondents "use the most" for checking. We focus on observations having more than $100 and less than $100,000 in their checking account to eliminate outliers.

Figure A-2: Liquid Assets and Checking Account Balances

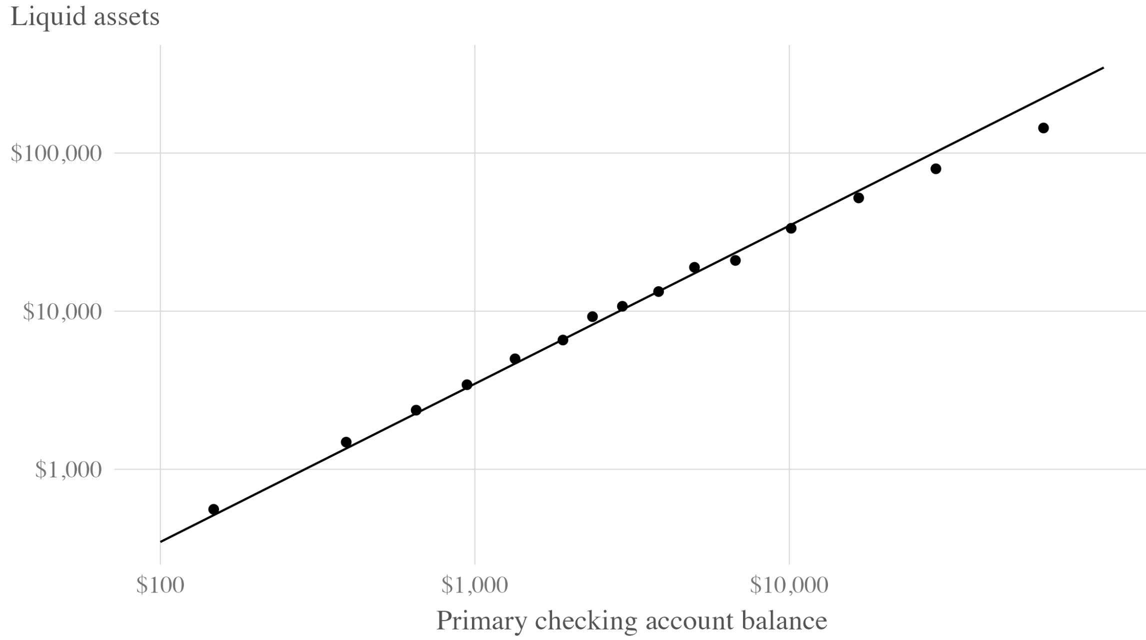

Notes: This figure shows the relationship between a household's liquid assets and the balance in their primary checking accounts, using data from the 2019 and 2022 SCF. We define primary checking account balance as the sum of balances at the institution that respondents "use the most" for checking. We define liquid assets using the liquid assets variable in the SCF Summary File, which includes balances in checking accounts, money market accounts, savings accounts, call accounts, and prepaid cards. We focus on observations having more than $100 and less than $100,000 in their checking account to eliminate outliers. The axes are both in log scale (base 10). The x-axis shows a dot for each vigntile of a household's balance in their primary checking account. The y-axis shows conditional means of liquid assets within that vigntile of checking account balance. The figure includes a best fit line.

Figure A-3: Binscatters

(a) First Stage

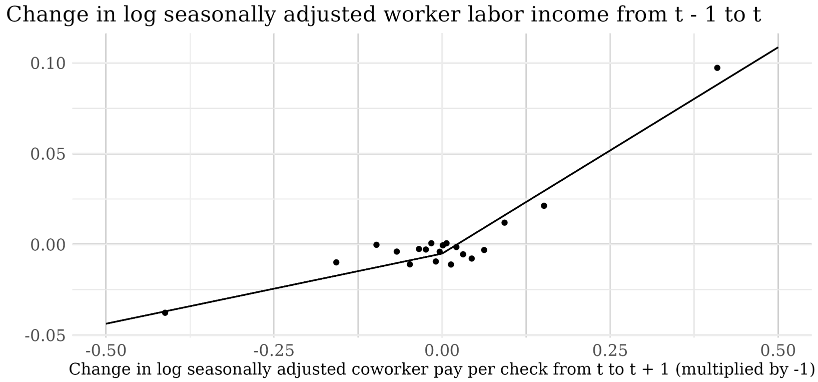

(b) Reduced Form

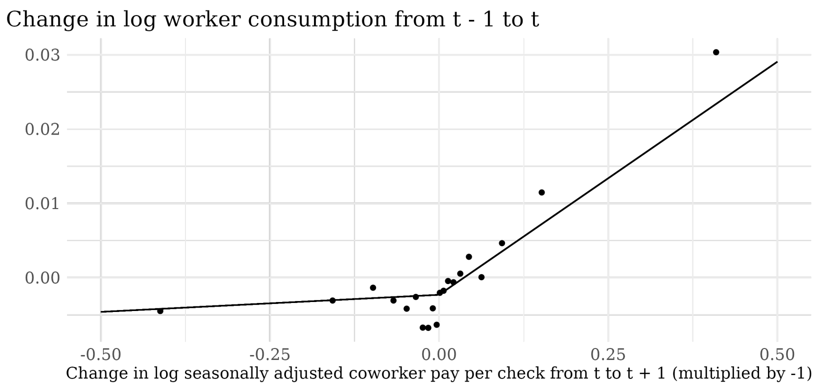

(c) Second Stage

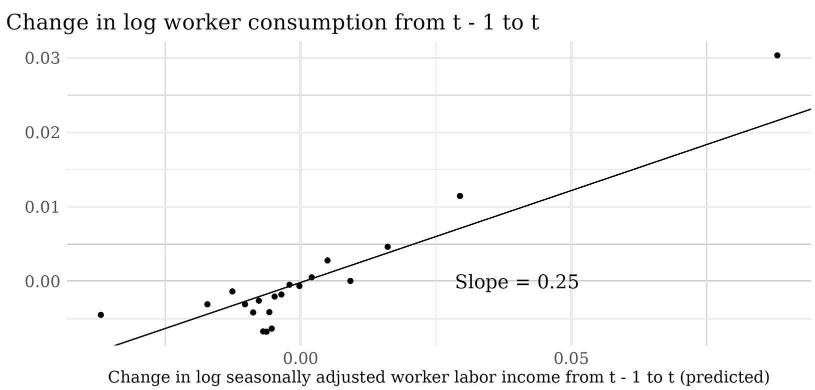

Notes: The top two panels in this figure show the effect of changes in firm pay on labor income (panel (a)) and nondurable consumption (panel (b)) in the JPMCI data. The x-axis shows a dot for each vingtile of the period-ahead change in log mean pay of all other workers at the same firm. The y-axis shows conditional means of the change in log labor income and the change in log nondurable consumption. Both panels include a best fit linear spline with a single knot at zero. Panel (c) shows an instrumental variables interpretation of the first two panels, where the x-axis is the change in labor income (instrumented using a linear spline with a single knot at zero of 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 comparable to $\hat{\beta}_{\varepsilon}$ from equation (18). See Section 4.3 for details.

Figure A-4: Heterogeneity in Quarterly MPCs by Liquid Assets

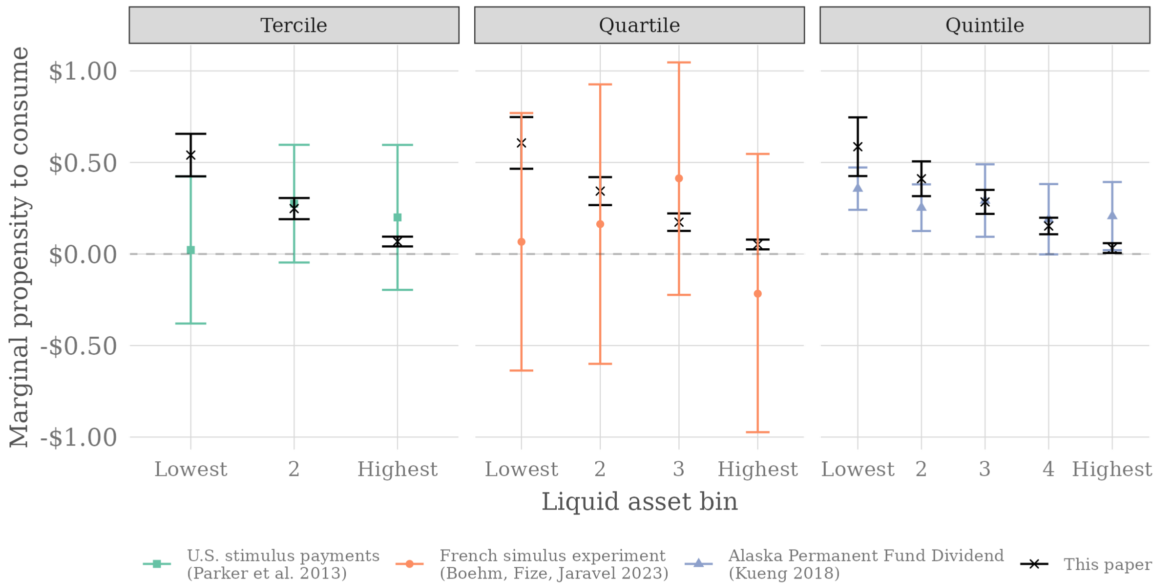

Notes: This figure shows the relationship between quarterly MPCs and liquid assets. It reports both point estimates and 95 percent confidence intervals. It compares the estimates from our research design to several estimates in the prior literature, using a quarterly horizon for all estimates. Parker et al. (2013) groups households by terciles of liquid assets in levels, Boehm, Fize, and Jaravel (2023) groups by quartiles of net liquid assets in levels, and Kueng (2018) groups by quintiles of liquid asset buffer. For comparison, we report the quarterly MPC from our research design by tercile, quartile, and quintile of liquid asset buffer.

Figure A-5: Income Event Study Around Realizations of Different Worker Instruments

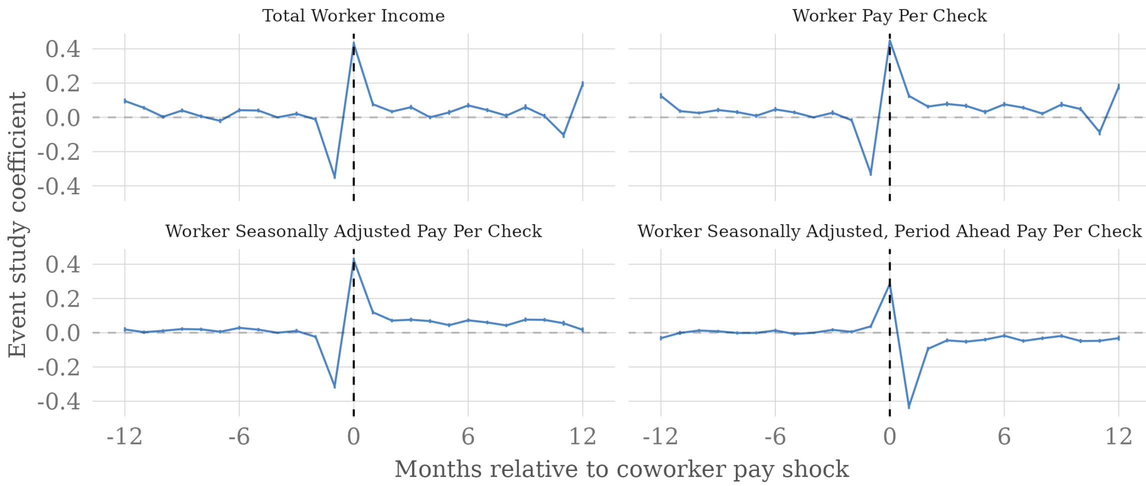

Notes: This figure shows the dynamic response of worker income to broader definitions of the *worker* instrument. It is equivalent to Figure 7, except that it uses worker income instead of coworker income for each version of the instrument. The top-left panel presents the response to total worker monthly pay. This corresponds to the specification in column (5) of Table A-14. The top-right panel refines the instrument by dividing total monthly pay by the number of checks received, corresponding to column (4) of Table A-14. The bottom-left panel further adjusts the instrument and the endogenous regressor for predictable annual seasonality, corresponding to column (3) of Table A-14. Finally, the bottom-right panel uses seasonally adjusted pay per paycheck for one period ahead. This corresponds to columns (1) and (2) of Table A-14. The black dashed line at time 0 marks the timing of coworker pay instrument realization and is included to highlight the reversal of the pattern when moving to the period-ahead specification. All panels include uniform 95 percent confidence intervals.

Table A-1: Summary Statistics for Main Analysis Sample

Panel A: Household-Level Variables	**Raw Data**			**Winsorized Data**		
	Mean	Median	Std. Dev	Mean	Median	Std. Dev
Household Labor Income	$5,281	$3,918	$10,177	$4,801	$3,918	$3,277
Nondurable Consumption	$2,389	$1,740	$7,800	$2,094	$1,740	$1,521
Checking Account Balance	$9,220	$2,415	$30,073	$6,548	$2,415	$9,790
Checking Buffer Ratio	11.7	1.3	1,560.7	5.1	1.3	9.0
Panel B: Job-Level Variables						
Job Labor Income	$4,117	$3,198	$9,142	$3,715	$3,198	$2,482
Coworker Labor Income	$3,685	$3,360	$2,616	$3,524	$3,360	$1,416
Number of Households: 1,327,214						
Number of Household-Employer-Months: 27,881,033						

Notes: This table reports summary statistics for our main analysis sample. Panel A reports variables aggregated at the household level. We construct the checking buffer ratio as the ratio of lagged checking account balance to lagged nondurable consumption. Panel B reports variables aggregated at the job level. Thus, coworker labor income reflects the average of peers within the same job for each household-employer. Although we perform our analysis using the raw data, we also report summary statistics for winsorized data. Winsorization is applied at the 95th percentile for the winsorized summary statistics. Winsorization makes a notable difference for the checking buffer summary statistics, since this is a ratio and therefore summary statistics are especially sensitive to outliers.

Table A-2: First Stage for Various Specifications

	Dependent Variable: Δ Log Income			
	(1)	(2)	(3)	(4)
Δ Log Instrument	0.153 (0.007)	0.369 (0.011)	0.508 (0.007)	0.713 (0.006)
Coworker Instrument	Period-Ahead Pay Per Check	Pay Per Check	Pay Per Check	Total Pay
Seasonal Adjustment	Yes	Yes	No	No
Type of Income Variation Captured by Instrument				
New Transitory Shock	Y	Y	Y	Y
Predictable Reversion of Transitory Shock		Y	Y	Y
Permanent Shock		Y	Y	Y
Predictable Recurring Annual Changes			Y	Y
Predictable Pay Schedule Variation				Y

Notes: This table reports the first stage coefficients from equation (19) for different versions of our instrument. Column (1) uses a one-period-ahead coworker pay per paycheck instrument, and is the first stage specification for column (1) of Table 1. Column (2) uses a contemporaneous coworker pay per paycheck instrument. Column (3) uses a contemporaneous coworker pay per paycheck instrument without seasonal adjustment. Column (4) uses a contemporaneous coworker total pay instrument, without seasonal adjustment. All specifications use data from 27,881,033 observations. Standard errors, which are clustered two ways at the household-job level, are reported in parentheses.

Table A-3: Summary Statistics for Distribution of Shock Sizes in First Stage Regression

Statistic	Coworker	Instrumented Worker
Mean	-0.002	0.001
Median	-0.001	0.001
5th Percentile	-0.209	-0.031
25th Percentile	-0.041	-0.005
75th Percentile	0.037	0.007
95th Percentile	0.202	0.032
Standard Deviation	0.171	0.026
Mean (Absolute Value)	0.085	0.013
Median (Absolute Value)	0.039	0.006

Notes: This table reports summary statistics on the instrument and fitted values from the first stage regression described in equation (19). Column (1) describes the one-period-ahead coworker pay per paycheck instrument used in column (1) of Table A-2. Column (2) describes the fitted values from the regression in column (1) of Table A-2, the corresponding first stage regression to column (1) in Table 1. All specifications use data from 27,881,033 observations.

Table A-4: First-Stage for Multiple Datasets

Periodicity	Data Source	First stage	
		Coef	(SE)
Paycheck	Bank account	0.15	(0.007)
	Payroll	0.17	(0.01)
	Time clock	0.14	(0.005)
Monthly	Bank account	0.71	(0.006)
	Payroll	0.63	(0.01)
	Time clock	0.80	(0.002)
Quarterly	Bank account	0.58	(0.008)
	Payroll	0.90	(0.01)
	Time clock	0.81	(0.004)
	Tax, WA, Lachowska et al. (2022)	0.56 to 0.65	(0.01 to 0.02)
	Tax, 7 states, CWBH	0.42 to 0.61	(0.01 to 0.03)

Notes: This table shows the results of the first stage regression on multiple datasets and with multiple levels of time aggregation. The rows labeled "bank account" correspond to the JPMCI data and regress change in log worker pay on change in log coworker pay (contemporaneous total pay for monthly and quarterly, and minus the period-ahead change in seasonally adjusted pay per paycheck for the paycheck level). The rows labeled "payroll" correspond to data from an anonymous payroll processor (PayrollCompany), and use the same specifications as with the JPMCI data. The rows labeled "time clock" correspond to the Homebase data and regress change in log worker hours on change in log coworker hours. The row labeled "CWBH" corresponds to the Continuous Wage and Benefit History data and regresses change in log worker pay on change in log coworker pay. Standard errors are clustered by firm (and by account for the JPMCI data). For the bank account results, the paycheck-level regression corresponds to column (1) of Table A-2. The monthly-level regression corresponds to column (4) in Table A-2, and the quarterly-level regression estimates the same specification on data rolled up to the quarterly level. The specification for Lachowska et al. (2022) is reported in Table C.1 of the paper.

Table A-5: Monthly and Cumulative MPCs

	Month 1	Month 2	Month 3
Monthly MPC	0.100	0.057	0.040
	(0.007)	(0.007)	(0.009)
Cumulative MPC	0.100	0.159	0.200
	(0.007)	(0.013)	(0.020)

Notes: This table reports the dynamic response of nondurable consumption to a transitory income shock. The first row shows month-by-month MPCs for the month of the shock and the two subsequent months, while the second row shows the cumulative MPC. Standard errors are clustered two ways at the household-job level. For more details, see Appendix J.

Table A-6: Consumption Elasticity Using Various Expenditure Categories

Dependent Variable: Δ *Log*	(1)	(2)
Nondurable Consumption		
Δ Log Income	0.218	0.273
	(0.015)	(0.019)
(Δ Log Income) x Checking		−0.509
		(0.050)
Strict Nondurables		
Δ Log Income	0.151	0.191
	(0.011)	(0.014)
(Δ Log Income) x Checking		−0.381
		(0.040)
Non-Work Related		
Δ Log Income	0.262	0.305
	(0.020)	(0.022)
(Δ Log Income) x Checking		−0.542
		(0.061)
Groceries		
Δ Log Income	0.221	0.252
	(0.016)	(0.019)
(Δ Log Income) x Checking		−0.624
		(0.060)

Notes: This table reports the elasticity of consumption with respect to income separately for all nondurable spending (as in our baseline specification in Table 1, columns (1) and (2)), for the subset of strict nondurables spending, for non-work-related nondurable spending, and spending on groceries. We define a spending category as work-related if it exhibits a larger-than-median drop at retirement (Aguiar and Hurst 2013). We define retirement as the first receipt of a Social Security check, as in Ganong and Noel (2019). The disaggregated analysis is run on a slightly different sample than the main analysis, which explains the differences in nondurable estimates to Table 1. All specifications use data from 27,860,495 household-employer-months. Standard errors, which are clustered two ways at the household-job level, are reported in parentheses.

Table A-7: Quarterly MPX and MPC Estimates Using Method of Laibson, Maxted, and Moll (2022)

	Estimate
Estimated Nondurable MPX	0.20
	(0.02)
Implied Total MPX	0.72
	(0.07)
Implied Notional MPC	0.23
	(0.02)

Notes: This table uses the method developed in Laibson, Maxted, and Moll (2022) to translate between various MPX and MPC estimates. The first row is the one-quarter nondurables MPX that we estimate in this paper (from Table A-5). The second row maps this nondurables MPX into a total spending MPX by multiplying our estimate by the Laibson, Maxted, and Moll (2022) adjustment factor of 3.6. The third row maps our nondurables MPX into the model-consistent, or "notional," MPC by multiplying our estimate by the Laibson, Maxted, and Moll (2022) adjustment factor of 1.14. Both adjustment factors are described in more detail on page three of Laibson, Maxted, and Moll (2022). Standard errors are reported in parentheses.

Table A-8: Impact of Income on Consumption - With Time-by-Location Fixed Effects

	Dependent Variable: Δ Log Non-Durable Consumption					
	(1)	(2)	(3)	(4)	(5)	(6)
Δ Log Income	0.221	0.274	0.218	0.267	0.211	0.260
	(0.015)	(0.018)	(0.014)	(0.006)	(0.014)	(0.017)
(Δ Log Income) x Checking		−0.475		-0.475		−0.475
		(0.047)		(0.021)		(0.046)
Coworker Instrument	Period-Ahead Pay Per Check	Period-Ahead Pay Per Check	Period-Ahead Pay Per Check	Period-Ahead Pay Per Check	Period-Ahead Pay Per Check	Period-Ahead Pay Per Check
Seasonal Adjustment	Yes	Yes	Yes	Yes	Yes	Yes
Fixed Effects	None	None	State-by-year	State-by-year	State-by-quarter	State-by-quarter
FEs Interacted with Assets	No	No	No	Yes	No	No

Notes: This table shows estimates of the elasticity of consumption with respect to income ($\hat{\beta}$) as in Table 1, but with additional time-location fixed effects. To accommodate computational constraints from a large number of fixed effects, we include the 10 largest states in our sample individually and combine observations from all other states into an 11th location category. This 11th location category contains 17% of observations. Columns (1) and (2) are the same as in Table 1. Columns (3) and (4) include state-by-year fixed effects, which are interacted with checking asset rank in column (4). Columns (5) and (6) include state-by-quarter fixed effects. Due to computational constraints, we are not able to interact the more granular state-by-quarter fixed effects with the assets control in column (6), and we are not able to report cluster-robust standard errors for column (4). Standard errors for all other columns are clustered two ways at the firm and household level.

Table A-9: Impact of Income on Consumption - Asset Level Control

	Dependent Variable: Δ Log Non-Durable Consumption				
	(1)	(2)	(3)	(4)	(5)
Δ Log Income	0.221	0.289	0.193	0.205	0.116
	(0.015)	(0.022)	(0.010)	(0.009)	(0.003)
(Δ Log Income) x Checking level		−0.526	−0.343	−0.367	−0.150
		(0.059)	(0.026)	(0.023)	(0.006)
Coworker Instrument	Period-Ahead Pay Per Check	Period-Ahead Pay Per Check	Pay Per Check	Pay Per Check	Total Pay
Seasonal Adjustment	Yes	Yes	Yes	No	No
Type of Income Variation Captured by Instrument					
New Transitory Shock	Y	Y	Y	Y	Y
Predictable Reversion of Transitory Shock			Y	Y	Y
Permanent Shock			Y	Y	Y
Predictable Recurring Annual Changes				Y	Y
Predictable Pay Schedule Variation					Y

Notes: This table shows estimates of the elasticity of consumption with respect to income ($\hat{\beta}$) as in Table 1. The difference is that, instead of interacting income changes with asset buffer, we interact it with the level of average checking account balance in the last 6 months. As in Table 1, asset variables are parameterized as $AssetRank/(N-0.5)$, so the variable is scaled from −0.5 for the lowest checking balance household to 0.5 for the highest checking balance household. All specifications use data from 27,881,033 worker-months. Standard errors, which are clustered two ways at the household-job level, are reported in parentheses.

Table A-10: MPC by Asset Buffer Quartile and Average MPC

	Quartile 1	Quartile 2	Quartile 3	Quartile 4	All
MPC	0.339	0.173	0.079	0.029	0.100
	(0.027)	(0.013)	(0.007)	(0.004)	(0.007)
Inputs					
Elasticity	0.482	0.317	0.183	0.105	0.221
	(0.038)	(0.025)	(0.015)	(0.014)	(0.015)
Mean Nondurable Consumption	$2,633	$2,579	$2,434	$1,911	$2,389
Mean Labor Income	$3,750	$4,724	$5,612	$7,037	$5,281
Liquid Asset Statistics					
Median Checking Account Balance	$485	$1,536	$3,739	$13,353	$2,415
Median Checking Account Buffer	0.28	0.82	2.21	12.13	1.31

Notes: This table reports MPCs separately by quartile of liquid asset buffer in columns (1) to (4), and the average MPC for all households in column (5). MPCs are calculated by estimating the elasticity separately for each group, and then scaling this estimate by the ratio of average nondurable consumption to average income in each bin. The table also reports the elasticity and the scaling factors for each group, as well as the median level of checking account balances and buffers for each group. One notable feature of this table is that high buffer households have lower average consumption than low buffer households. This pattern arises because we construct our asset buffer measure as the ratio of lagged assets to lagged nondurable consumption. Thus, on average high buffer households have relatively higher levels of checking account assets and relatively lower levels of consumption. Our MPC estimates are nearly identical if we group by asset *level* rather than asset *buffer* (see Table A-11), so this pattern does not drive our results. Another notable feature is that the average MPC does not equal the average of the MPCs for the four quartiles. This is because IV estimates are not generally a sample-weighted average of subsample coefficients. Rather, they are a weighted average of each subsample IV plus an additional IV estimated using the subsample means. The weights on the subsample IV estimates are increasing in the within-group variance of the instrument and strength of the subsample first-stage (Stephens and Unayama 2019). Each quartile contains 6,985,699 worker-months. Standard errors, which are clustered two ways at the household-job level, are reported in parentheses.

Table A-11: MPC by Asset Level Quartile

	Quartile 1	Quartile 2	Quartile 3	Quartile 4	All
MPC	0.335	0.177	0.086	0.035	0.100
	(0.034)	(0.015)	(0.007)	(0.004)	(0.007)
Inputs					
Elasticity	0.520	0.328	0.196	0.101	0.221
	(0.053)	(0.028)	(0.016)	(0.012)	(0.015)
Mean Nondurable Consumption	$1,873	$2,286	$2,507	$2,890	$2,389
Mean Labor Income	$2,911	$4,251	$5,715	$8,246	$5,281
Liquid Asset Statistics					
Median Checking Account Balance	$404	$1,471	$3,802	$15,340	$2,415
Median Checking Account Buffer	0.30	0.82	2.00	9.58	1.31

Notes: This figure repeats Table A-10 but groups households by the level of average checking account balance in the last six months rather than by their liquid asset buffer. Each quartile contains 6,985,699 worker-months. Standard errors, which are clustered two ways at the household-job level, are reported in parentheses.

Table A-12: MPC by Asset Buffer Decile

	Decile 1	Decile 2	Decile 3	Decile 4	Decile 5	Decile 6	Decile 7	Decile 8	Decile 9	Decile 10
MPC	0.408	0.308	0.257	0.186	0.142	0.112	0.065	0.053	0.033	0.019
	(0.041)	(0.029)	(0.021)	(0.017)	(0.013)	(0.011)	(0.008)	(0.007)	(0.006)	(0.004)
Inputs										
Elasticity	0.519	0.459	0.423	0.333	0.277	0.237	0.153	0.144	0.107	0.092
	(0.052)	(0.043)	(0.035)	(0.030)	(0.025)	(0.024)	(0.019)	(0.019)	(0.019)	(0.018)
Mean Nondurable Consumption	$2,670	$2,610	$2,599	$2,591	$2,560	$2,517	$2,417	$2,262	$2,033	$1,635
Mean Labor Income	$3,395	$3,891	$4,271	$4,640	$4,989	$5,334	$5,698	$6,114	$6,680	$7,795
Liquid Asset Statistics										
Median Checking Account Balance	$260	$600	$956	$1,405	$1,994	$2,813	$4,104	$6,442	$11,164	$25,566
Median Checking Account Buffer	0.15	0.32	0.51	0.74	1.08	1.59	2.49	4.38	9.47	35.80

Notes: This table reports MPCs separately by decile of liquid asset buffer. MPCs are calculated by estimating the elasticity separately for each liquid asset bin, and then scaling this estimate by the ratio of average nondurable consumption to average income in each bin. The table also reports the elasticity and the scaling factors for each group, as well as the median level of checking account balances and buffers for each group. One notable feature of this table is that high buffer households have lower average consumption than low buffer households. This pattern arises because we construct our asset buffer measure as the ratio of lagged assets to lagged nondurable consumption. Thus, on average high buffer households have relatively higher levels of checking account assets and relatively lower levels of consumption. Each decile contains 2,796,378 worker-months. Standard errors, which are clustered two ways at the household-job level, are reported in parentheses.

Table A-13: Monthly and Quarterly MPCs by Asset Buffer Quantiles

	Monthly	Quarterly
Tercile 1	0.300	0.540
	(0.023)	(0.059)
Tercile 2	0.120	0.248
	(0.009)	(0.029)
Tercile 3	0.034	0.068
	(0.004)	(0.014)
Quartile 1	0.339	0.607
	(0.027)	(0.072)
Quartile 2	0.173	0.344
	(0.013)	(0.039)
Quartile 3	0.079	0.174
	(0.007)	(0.024)
Quartile 4	0.029	0.052
	(0.004)	(0.014)
Quintile 1	0.353	0.586
	(0.030)	(0.082)
Quintile 2	0.218	0.411
	(0.017)	(0.048)
Quintile 3	0.126	0.285
	(0.010)	(0.033)
Quintile 4	0.059	0.153
	(0.006)	(0.023)
Quintile 5	0.025	0.032
	(0.004)	(0.014)
Decile 1	0.408	0.670
	(0.041)	(0.120)
Decile 2	0.308	0.516
	(0.029)	(0.087)
Decile 3	0.257	0.482
	(0.021)	(0.064)
Decile 4	0.186	0.352
	(0.017)	(0.056)
Decile 5	0.142	0.336
	(0.013)	(0.045)
Decile 6	0.112	0.240
	(0.011)	(0.040)
Decile 7	0.065	0.139
	(0.008)	(0.033)
Decile 8	0.053	0.162
	(0.007)	(0.028)
Decile 9	0.033	0.045
	(0.006)	(0.023)
Decile 10	0.019	0.023
	(0.004)	(0.015)

Notes: This table reports MPCs separately by quantiles of liquid asset buffer. MPCs are calculated by estimating the elasticity separately for each liquid asset bin, and then scaling this estimate by the ratio of average nondurable consumption to average income in each bin. We estimate the MPCs by terciles, quartiles, quintiles and deciles of liquid asset buffer. The monthly estimates are obtained using the specification from equation (18), whereas the quarterly estimates are the three-month cumulative estimates obtained using the method in Appendix J. Standard errors, which are clustered two ways at the household-job level, are reported in parentheses.

Table A-14: Impact of Income on Consumption Using Worker Rather than Coworker Instrument

	Dependent Variable: Δ Log Non-Durable Consumption					
	(1)	(2)	(3)	(4)	(5)	(6)
Δ Log Income	0.155	0.169	0.114	0.126	0.112	0.120
	(0.004)	(0.004)	(0.002)	(0.002)	(0.001)	(0.001)
(Δ Log Income) × Checking		−0.273	−0.244	−0.271	−0.232	−0.237
		(0.010)	(0.005)	(0.006)	(0.004)	(0.004)
Worker Instrument	Period-Ahead Pay Per Check	Period-Ahead Pay Per Check	Pay Per Check	Pay Per Check	Total Pay	OLS: Worker all Labor Income
Seasonal Adjustment	Yes	Yes	Yes	No	No	No
Type of Income Variation Captured by Instrument						
New Transitory Shock	Y	Y	Y	Y	Y	Y
Predictable Reversion of Transitory Shock			Y	Y	Y	Y
Permanent Shock			Y	Y	Y	Y
Predictable Recurring Annual Changes				Y	Y	Y
Predictable Pay Schedule Variation					Y	Y

Notes: This table reports estimates of the elasticity of consumption with respect to income ($\hat{\beta}_\varepsilon$). It is similar to Table 1 except these specifications use a worker's own job-level income change, rather than their coworkers' average income changes, as an instrument for changes in household labor income. Columns (1) and (2) show IV estimates of the effect of income on consumption using worker period-ahead pay per paycheck as an instrument for income changes. Column (2) controls for checking account buffer. The checking buffer variable is parameterized as $BufferRank/N$ - 0.5, so that the variable is scaled from −0.5 for the lowest asset buffer household to 0.5 for the highest asset buffer household. Column (3) shows IV estimates using a contemporaneous worker pay per paycheck instrument, rather than a period ahead one. Column (4) shows IV estimates using a contemporaneous worker pay per paycheck instrument without seasonal adjustment. Column (5) shows IV estimates using a contemporaneous worker total pay instrument without seasonal adjustment. Column (6) shows an OLS version of column (5). While column (5) uses a worker's change in total pay from one job as an instrument for that household's total labor income change in a given month, column (6) simply regresses a household's total labor income change on a household's change in consumption (without any instrument). All specifications use data from 27,881,033 observations. Standard errors, which are clustered two ways at the household-job level, are reported in parentheses.

Appendix B Generalizing the Income Process

Until this point, we have assumed a standard income process with a completely transitory shock, and a completely permanent one, i.e. a unit root shock. In addition, we have abstracted from nearly all predictable components of income and consumption growth. These are typically addressed prior to analysis by a detrending process—consumption and income are regressed on a set of age or time controls, and the resulting residuals are used. In our setting, however, the income process may feature more complicated serial correlation structures that have to be addressed. This is particularly an issue in our context of monthly time aggregation, where some forms of seasonality that occur at sub-annual frequencies have not been of concern to studies that use annual data.

We discuss here the key deviations for the income process in our sample and choice of frequency and adjustments to the data that allow us to retain our empirical results.

B.1 Firm-Specific Annual Seasonality

We find evidence in our data of a serial correlation between income growth in month t and income growth in month $t-12$. Moreover, the correlation remains if we instead use average coworker income growth in period t and relate it to own-income growth in period $t-12$. We can see this, for example, in Figure 7, in the top two panels. The instrument is correlated with a spike in income in period 0, and there are similar deviations at event times -12 and 12. This is a violation of our Prediction E2.

We note that these correlations are not simply secular trends for specific calendar months—that is, it is not the case that pay is simply higher in December for all firms. If we control for a set of calendar month fixed effects, these patterns persist. This means there is serial correlation at a 12-month lag that is firm-specific. This could happen, for example, if different firms have a particular month where operations shut down, or when bonuses are paid.

Our solution is to apply a seasonal adjustment to all of our data before implementing our IV estimator. We use information from income growth 12 months prior to create a linear prediction of current income growth, $\tau_s \Delta y_{i,t-12}$. The residual variation is therefore not well-explained by lagged income growth. Formally, we choose the seasonal adjustment factor, τ_s, to minimize the residual correlation between our instrument and growth 12 months ago:

$$\begin{aligned} \tau_s \equiv \underset{\tau}{\operatorname{argmin}} & \left(\frac{\operatorname{cov}(\breve{\Delta} y_{i,t-12}, -\breve{\Delta} y_{j(-i,t),t+1})}{\operatorname{var}(-\breve{\Delta} y_{j(-i,t),t+1})} \right)^2 \\ \text{where} & \\ \breve{\Delta} y_{i,t} &\equiv \Delta y_{i,t} - \tau \Delta y_{i,t-12} \\ \breve{\Delta} y_{j(-i,t),t+1} &\equiv \Delta y_{j(-i,t),t+1} - \tau \Delta y_{j(-i,t),t-11} \end{aligned} \tag{B.1}$$

In other words, τ_s is chosen to minimize the reduced-form coefficient from a regression of *adjusted* income growth in period $t-12$ on an *adjusted* instrument, i.e. *adjusted* coworker income growth in period $t+1$, with its sign inverted.

We can see, in Figure 7, as we move from the top panels to the bottom panels, that this seasonal adjustment effectively removes the correlation between the instrument and normalized income growth at event times -12 and 12.

B.2 Spurious Correlation in Pay Frequency

At the monthly frequency, a particular type of correlation arises due to the discrete nature of pay periods. For workers who receive pay every two weeks (which is the modal pay frequency in the U.S.), for example, some months feature two pay days, and others will feature three. However, this variation in monthly pay is relatively predictable. To the extent that household consumption is not responsive to this predictable change, our estimates of consumption sensitivity may be biased toward zero. It is as if we have measurement error in the monthly variation. Our coworker instrument does not address this issue, since such cycles generate mechanical correlation in monthly pay growth between coworkers, which will be captured in our first-stage regressions. Finally, if there is any type of consumption, especially more durable consumption, that happens mechanically on pay-day, our estimate of consumption sensitivity will pick up this correlation, potentially overstating households' sensitivity.

We can model this directly within our income process. In what follows, we assume that all coworkers at the same firm have the same pay frequency. Let Y_{it} be monthly income in *levels* and $M_{j(i,t),t}$ the number of paychecks received for household i, who works at firm j in month t. Define monthly pay per paycheck as:

$$Y_{it}^p \equiv \frac{Y_{it}}{M_{j(i,t),t}} \tag{B.2}$$

We next take logs and decompose the income process:

$$\begin{aligned} y_{it} &\equiv \log Y_{it} \\ &= \log Y_{it}^{p} + \log M_{j(i,t),t} \\ &= y_{it}^{p} + m_{j(i,t),t} \\ &= z_{it} + \varepsilon_{it}^{w} + \varepsilon_{j(i,t),t}^{f} + m_{j(i,t),t} \end{aligned} \tag{B.3}$$

where $m_{j(i,t),t}$ is the log of the number of paychecks. Since our previous shocks were proportional to total monthly pay, we model them now as shocks to monthly pay per paycheck. We also extend the consumption function to include a potential response to the number of paychecks received in a month:

$$\Delta c_{i,t} = \beta_{\eta}\left(\eta_{i,t}^{w} + \eta_{j(i,t),t}^{f}\right) + \beta_{\varepsilon}\left(\varepsilon_{it}^{w} + \varepsilon_{j(i,t),t}^{f}\right) + \beta_{m} m_{j(i,t),t} + \zeta_{it} \tag{B.4}$$

We can show, in this case, that our estimator will suffer from bias, driven by the variation in number of pay checks from month to month:

Result FREQ (Pay Frequency Bias)

If Assumptions IID2 and CE2 hold, and there is variation in the number of payments in each period, the estimate of β_{ε} from Result ID is biased:

$$\hat{\beta}_{\varepsilon,HM-BPP-Co} \equiv \frac{\operatorname{cov}\left(\Delta c_{it}, -\Delta y_{j(-i,t),t+1}\right)}{\operatorname{cov}\left(\Delta y_{it}, -\Delta y_{j(-i,t),t+1}\right)} = \pi\beta_{\varepsilon} + (1-\pi)\beta_{m}$$

where:

$$\pi = \frac{\sigma_{f\varepsilon}^{2}}{\sigma_{f\varepsilon}^{2} + \sigma_{m}^{2}}$$

and σ_{m}^{2} is the variance of log of the number of pay checks in a month.

Proof. First, note that the coworker instrument can be simplified as follows:

$$
\begin{aligned}
-\Delta y_{j(-i,t),t+1} &\equiv -\Delta \ln\left(\frac{1}{N_{j(-i,t),t+1}} \sum_{i'|j(i',t+1)=j(i,t),i'\neq i} Y_{i',t+1}\right) \\
&= -\Delta \ln\left(\frac{1}{N_{j(-i,t),t+1}} \sum_{i'|j(i',t+1)=j(i,t),i'\neq i} e^{y_{i',t+1}}\right) \\
&= -\ln\left(\frac{1}{N_{j(-i,t),t+1}} \sum_{i'|j(i',t+1)=j(i,t),i'\neq i} e^{y_{i',t+1}}\right) \\
&\quad + \ln\left(\frac{1}{N_{j(-i,t),t}} \sum_{i'|j(i',t)=j(i,t),i'\neq i} e^{y_{i',t}}\right) \\
&= -\ln\left(\frac{1}{N_{j(-i,t),t+1}} \sum_{i'|j(i',t+1)=j(i,t),i'\neq i} \left(e^{\eta^f_{j(i',t+1),t+1}+\eta^f_{j(i',t+1),t}+\varepsilon^f_{j(i',t+1),t+1}+m_{j(i',t+1),t+1}}\right.\right. \\
&\qquad\qquad \left.\left.\times\, e^{z_{i',t-1}+\eta^w_{i',t}+\eta^w_{i',t+1}+\varepsilon^w_{i',t+1}}\right)\right) \\
&\quad + \ln\left(\frac{1}{N_{j(-i,t),t}} \sum_{i'|j(i',t)=j(i,t),i'\neq i} e^{\eta^f_{j(i',t),t}+\varepsilon^f_{j(i',t),t}+m_{j(i',t),t}} e^{z_{i',t-1}+\eta^w_{i',t}+\varepsilon^w_{i',t}}\right) \\
&= -\ln\left(e^{\eta^f_{j(i,t),t+1}+\eta^f_{j(i,t),t}+\varepsilon^f_{j(i,t),t+1}+m_{j(i,t),t+1}} \frac{1}{N_{j(-i,t),t+1}} \sum_{i'|j(i',t+1)=j(i,t),i'\neq i} e^{z_{i',t-1}+\eta^w_{i',t}+\eta^w_{i',t+1}+\varepsilon^w_{i',t+1}}\right) \\
&\quad + \ln\left(e^{\eta^f_{j(i,t),t}+\varepsilon^f_{j(i,t),t}+m_{j(i,t),t}} \frac{1}{N_{j(-i,t),t}} \sum_{i'|j(i',t)=j(i,t),i'\neq i} e^{z_{i',t-1}+\eta^w_{i',t}+\varepsilon^w_{i',t}}\right) \\
&= -\eta^f_{j(i,t),t+1} - \eta^f_{j(i,t),t} - \varepsilon^f_{j(i,t),t+1} - m_{j(i,t),t+1} + \eta^f_{j(i,t),t} + \varepsilon^f_{j(i,t),t} + m_{j(i,t),t} \\
&\quad - \ln\left(\frac{1}{N_{j(-i,t),t+1}} \sum_{i'|j(i',t+1)=j(i,t),i'\neq i} e^{z_{i',t-1}+\eta^w_{i',t}+\eta^w_{i',t+1}+\varepsilon^w_{i',t+1}}\right) \\
&\quad + \ln\left(\frac{1}{N_{j(-i,t),t}} \sum_{i'|j(i',t)=j(i,t),i'\neq i} e^{z_{i',t-1}+\eta^w_{i',t}+\varepsilon^w_{i',t}}\right) \\
&= -\eta^f_{j(i,t),t+1} - \Delta\varepsilon^f_{j(i,t),t+1} - \Delta m_{j(i,t),t+1} + g_{j(-i,t+1),t+1}
\end{aligned}
$$

where the function $g(\cdot)$ is shorthand for the final two terms and is a function of a combination of shocks that are independent of $(\Delta y_{i,t}, \Delta c_{i,t})$.

Next, we can directly calculate the numerator and denominator of our main expression:

$$
\begin{aligned}
\operatorname{cov}(\Delta c_{it}, -\Delta y_{j(-i,t),t+1}) &= \operatorname{cov}\Big(\beta_\eta\left(\eta^w_{i,t}+\eta^f_{j(i,t),t}\right)+\beta_\varepsilon\left(\varepsilon^w_{it}+\varepsilon^f_{j(i,t),t}\right)+\beta_m m_{j(i,t),t}+\zeta_{it},\\
&\qquad -\eta^f_{j(i,t),t+1}-\Delta\varepsilon^f_{j(i,t),t+1}-\Delta m_{j(i,t),t+1}+g_{j(-i,t+1),t+1}\Big)\\
&= \operatorname{cov}\Big(\beta_\eta\left(\eta^w_{i,t}+\eta^f_{j(i,t),t}\right)+\beta_\varepsilon\left(\varepsilon^w_{it}+\varepsilon^f_{j(i,t),t}\right)+\beta_m m_{j(i,t),t}+\zeta_{it},\\
&\qquad -\eta^f_{j(i,t),t+1}-\varepsilon^f_{j(i,t),t+1}+\varepsilon^f_{j(i,t),t}-m_{j(i,t),t+1}+m_{j(i,t),t}+g_{j(-i,t+1),t+1}\Big)\\
&= \beta_\varepsilon\sigma^2_{f\varepsilon}+\beta_m\sigma^2_m\\
\operatorname{cov}(\Delta y_{it}, -\Delta y_{j(-i,t),t+1}) &= \operatorname{cov}\Big(\eta^w_{i,t}+\eta^f_{j(i,t),t}+\Delta\varepsilon^w_{i,t}+\Delta\varepsilon^f_{j(i,t),t}+\Delta m_{j(i,t),t},\\
&\qquad -\eta^f_{j(i,t),t+1}-\Delta\varepsilon^f_{j(i,t),t+1}-\Delta m_{j(i,t),t+1}+g_{j(-i,t+1),t+1}\Big)\\
&= \operatorname{cov}\Big(\eta^w_{i,t}+\eta^f_{j(i,t),t}+\varepsilon^w_{i,t}-\varepsilon^w_{i,t-1}+\varepsilon^f_{j(i,t),t}-\varepsilon^f_{j(i,t-1),t-1}+m_{j(i,t),t}-m_{j(i,t),t-1},\\
&\qquad -\eta^f_{j(i,t),t+1}-\varepsilon^f_{j(i,t),t+1}+\varepsilon^f_{j(i,t),t}-m_{j(i,t),t+1}+m_{j(i,t),t}+g_{j(-i,t+1),t+1}\Big)\\
&= \sigma^2_{f\varepsilon}+\sigma^2_m
\end{aligned}
$$

Taking the ratio of the two covariances yields our result.

□

There is bias from two sources. First, through the parameter β_m, if consumption also changes due to the number of paychecks in a month, above and beyond the growth in monthly pay, this will be captured in our estimator. This pattern, we should note, does not follow from a standard model, but there is some evidence of modest excess consumption on the day one gets paid (Olafsson and Pagel 2018). Second, even if the parameter β_m were zero, there is attenuation bias of β_ε captured by $\pi \equiv \sigma^2_{f\varepsilon} \Big/ (\sigma^2_{f\varepsilon}+\sigma^2_m)$, where σ^2_m is the variance of the log of number of paychecks in a month. Intuitively, part of the variation in log pay is due to predictable changes in the number of pay checks, overstating the variation in unpredictable and transitory shocks to pay and leading to downward bias in the regression coefficient.

A straightforward solution is available for this issue. We can instead use as an instrument the average pay per paycheck of coworkers:

$$
-\Delta y^p_{j(-i,t),t+1} \equiv -\Delta \ln\left(\frac{1}{N_{j(-i,t),t+1}}\sum_{i'|j(i',t+1)=j(i,t),i'\neq i} Y^p_{i',t+1}\right) \tag{B.5}
$$

With this instrument, we will again identify our parameter of interest:

Result PER (Identification via Pay Per Paycheck)

If Assumptions IID2 and CE2 hold, and there is variation in the number of payments in each period, β_ε is identified if coworker pay per paycheck, $\Delta y^p_{j(-i,t+),t+1}$, is used as an instrument:

$$
\hat{\beta}_{\varepsilon,PPP} \equiv \frac{\operatorname{cov}\left(\Delta c_{it}, -\Delta y^p_{j(-i,t),t+1}\right)}{\operatorname{cov}\left(\Delta y_{it}, -\Delta y^p_{j(-i,t),t+1}\right)} = \frac{\operatorname{cov}\left(\Delta c_{it}, \varepsilon^f_{j(i,t),t}\right)}{\operatorname{var}\left(\varepsilon^f_{j(i,t),t}\right)} = \frac{\beta_\varepsilon\sigma^2_{f\varepsilon}}{\sigma^2_{f\varepsilon}} = \beta_\varepsilon
$$

Proof. Note that the new coworker instrument, under our assumptions, can be simplified as follows:

$$
\begin{aligned}
-\Delta y^p_{j(-i,t),t+1} &\equiv -\Delta \ln\left(\frac{1}{N_{j(-i,t),t+1}} \sum_{i'|j(i',t+1)=j(i,t),i'\neq i} Y^p_{i',t+1}\right) \\
&= -\Delta \ln\left(\frac{1}{N_{j(-i,t),t+1}} \sum_{i'|j(i',t+1)=j(i,t),i'\neq i} e^{y^p_{i',t+1}}\right) \\
&= -\ln\left(\frac{1}{N_{j(-i,t),t+1}} \sum_{i'|j(i',t+1)=j(i,t),i'\neq i} e^{y^p_{i',t+1}}\right) \\
&\quad + \ln\left(\frac{1}{N_{j(-i,t),t}} \sum_{i'|j(i',t)=j(i,t),i'\neq i} e^{y^p_{i',t}}\right) \\
&= -\ln\left(\frac{1}{N_{j(-i,t),t+1}} \sum_{i'|j(i',t+1)=j(i,t),i'\neq i} e^{\eta^f_{j(i',t+1),t+1}+\eta^f_{j(i',t+1),t}+\varepsilon^f_{j(i',t+1),t+1}} e^{z_{i',t-1}+\eta^w_{i',t}+\eta^w_{i',t+1}+\varepsilon^w_{i',t+1}}\right) \\
&\quad + \ln\left(\frac{1}{N_{j(-i,t),t}} \sum_{i'|j(i',t)=j(i,t),i'\neq i} e^{\eta^f_{j(i',t),t}+\varepsilon^f_{j(i',t),t}} e^{z_{i',t-1}+\eta^w_{i',t}+\varepsilon^w_{i',t}}\right) \\
&= -\ln\left(e^{\eta^f_{j(i,t),t+1}+\eta^f_{j(i,t),t}+\varepsilon^f_{j(i,t),t+1}} \frac{1}{N_{j(-i,t),t+1}} \sum_{i'|j(i',t+1)=j(i,t),i'\neq i} e^{z_{i',t-1}+\eta^w_{i',t}+\eta^w_{i',t+1}+\varepsilon^w_{i',t+1}}\right) \\
&\quad + \ln\left(e^{\eta^f_{j(i,t),t}+\varepsilon^f_{j(i,t),t}} \frac{1}{N_{j(-i,t),t}} \sum_{i'|j(i',t)=j(i,t),i'\neq i} e^{z_{i',t-1}+\eta^w_{i',t}+\varepsilon^w_{i',t}}\right) \\
&= -\eta^f_{j(i,t),t+1} - \eta^f_{j(i,t),t} - \varepsilon^f_{j(i,t),t+1} + \eta^f_{j(i,t),t} + \varepsilon^f_{j(i,t),t} \\
&\quad - \ln\left(\frac{1}{N_{j(-i,t),t+1}} \sum_{i'|j(i',t+1)=j(i,t),i'\neq i} e^{z_{i',t-1}+\eta^w_{i',t}+\eta^w_{i',t+1}+\varepsilon^w_{i',t+1}}\right) \\
&\quad + \ln\left(\frac{1}{N_{j(-i,t),t}} \sum_{i'|j(i',t)=j(i,t),i'\neq i} e^{z_{i',t-1}+\eta^w_{i',t}+\varepsilon^w_{i',t}}\right) \\
&= -\eta^f_{j(i,t),t+1} - \Delta\varepsilon^f_{j(i,t),t+1} + g_{j(-i,t+1),t+1}
\end{aligned}
$$

where the function $g(\cdot)$ is shorthand for the final two terms and is a function of a combination of shocks that are independent of $(\Delta y_{i,t}, \Delta c_{i,t})$.

We can then show our result directly by simplifying the numerator and denominator:

$$
\begin{aligned}
\text{cov}\left(\Delta c_{it}, -\Delta y^p_{j(-i,t),t+1}\right) &= \text{cov}\Big(\beta_\eta(\eta^w_{it} + \eta^f_{j(i,t),t}) + \beta_\varepsilon(\varepsilon^w_{it} + \varepsilon^f_{j(i,t),t}) + \beta_m m_{j(i,t),t} + \zeta_{it}, \\
&\qquad -\eta^f_{j(i,t),t+1} - \Delta\varepsilon^f_{j(i,t),t+1} + g_{j(-i,t),t+1}\Big) \\
&= \text{cov}\Big(\beta_\eta(\eta^w_{it} + \eta^f_{j(i,t),t}) + \beta_\varepsilon(\varepsilon^w_{it} + \varepsilon^f_{j(i,t),t}) + \beta_m m_{j(i,t),t} + \zeta_{it}, \\
&\qquad -\eta^f_{j(i,t),t+1} - \varepsilon^f_{j(i,t),t+1} + \varepsilon^f_{j(i,t),t} + g_{j(-i,t),t+1}\Big) \\
&= \beta_\varepsilon \sigma^2_{f\varepsilon} \\
\text{cov}\left(\Delta y_{it}, -\Delta y^p_{j(-i,t),t+1}\right) &= \text{cov}\Big(\eta^w_{i,t} + \eta^f_{j(i,t),t} + \Delta\varepsilon^w_{i,t} + \Delta\varepsilon^f_{j(i,t),t} + \Delta m_{j(i,t),t}, \\
&\qquad -\eta^f_{j(i,t),t+1} - \Delta\varepsilon^f_{j(i,t),t+1} + g_{j(-i,t),t+1}\Big) \\
&= \text{cov}\Big(\eta^w_{i,t} + \eta^f_{j(i,t),t} + \varepsilon^w_{i,t} - \varepsilon^w_{i,t-1} + \varepsilon^f_{j(i,t),t} - \varepsilon^f_{j(i,t-1),t-1} + \Delta m_{j(i,t),t}, \\
&\qquad -\eta^f_{j(i,t),t+1} - \varepsilon^f_{j(i,t),t+1} + \varepsilon^f_{j(i,t),t} + g_{j(-i,t),t+1}\Big) \\
&= \sigma^2_{f\varepsilon}
\end{aligned}
$$

The result then follows by taking the ratio of the two covariances. □

B.3 Less-Than-Full Persistence in the Permanent Shock

We have assumed thus far that permanent shocks were completely permanent, i.e. they follow a unit root process. It is possible that they may instead have high, but not complete, persistence. We show in this section that our estimate of β_ε is largely robust to this more general income process. Suppose that the permanent component instead follows an autoregressive, AR(1), process as follows:

$$z_{it} = \rho z_{i,t-1} + \eta_{it}, \tag{B.6}$$

with $\rho < 1$. In this case, the first difference of log income in period t will now also be a function of $z_{i,t-1}$, confounding our above results. However, as shown in Kaplan and Violante (2010), there is negligible bias from assuming a unit root *permanent* component when estimating responses to *transitory* shocks, when the *permanent* component of income is instead only highly persistent.

Nevertheless, if we wish to directly account for this feature of the income process and we have an estimate of the parameter ρ, Kaplan and Violante (2010) shows that we can use a quasi-difference in log income that will restore the properties of the income process that we have leveraged above. Define the quasi-difference in log income as:

$$\grave{\Delta} y_{it} = y_{it} - \rho y_{i,t-1} \tag{B.7}$$

Using this quasi-difference, we can still recover the parameter β_ε:

Result AR (Identification with Persistent Shocks)

If Assumptions IID2 and CE2 hold, and the permanent component of income follows equation (B.6), β_ε is identified using a one-period lead, quasi-differenced coworker income instrument for the quasi-difference of endogenous log income, as in equation (B.7):

$$\hat{\beta}_{\varepsilon,\rho} \equiv \frac{\text{cov}\left(\Delta c_{it}, -\grave{\Delta} y_{j(-i,t),t+1}\right)}{\text{cov}\left(\grave{\Delta} y_{it}, -\grave{\Delta} y_{j(-i,t),t+1}\right)} = \frac{\rho\text{cov}\left(\Delta c_{it}, \varepsilon^f_{j(i,t),t}\right)}{\rho\text{var}\left(\varepsilon^f_{j(i,t),t}\right)} = \frac{\rho\beta_\varepsilon\sigma^2_{f\varepsilon}}{\rho\sigma^2_{f\varepsilon}} = \beta_\varepsilon$$

Proof. First, note that in the presence of a persistent shock, as in equation (B.6), the coworker instrument can be simplified as follows:

$$\begin{aligned}
-\grave{\Delta} y_{j(-i,t),t+1} &\equiv -\grave{\Delta} \ln\left(\frac{1}{N_{j(-i,t),t+1}} \sum_{i'|j(i',t+1)=j(i,t),i'\neq i} Y_{i',t+1}\right) \\
&= -\grave{\Delta} \ln\left(\frac{1}{N_{j(-i,t),t+1}} \sum_{i'|j(i',t+1)=j(i,t),i'\neq i} e^{y_{i',t+1}}\right) \\
&= -\ln\left(\frac{1}{N_{j(-i,t),t+1}} \sum_{i'|j(i',t+1)=j(i,t),i'\neq i} e^{y_{i',t+1}}\right) \\
&\quad + \rho\ln\left(\frac{1}{N_{j(-i,t),t}} \sum_{i'|j(i',t)=j(i,t),i'\neq i} e^{y_{i',t}}\right) \\
&= -\ln\left(\frac{1}{N_{j(-i,t),t+1}} \sum_{i'|j(i',t+1)=j(i,t),i'\neq i} e^{\eta^f_{j(i',t+1),t+1}+\rho\eta^f_{j(i',t+1),t}+\varepsilon^f_{j(i',t+1),t+1}} e^{\rho^2 z_{i',t-1}+\rho\eta^w_{i',t}+\eta^w_{i',t+1}+\varepsilon^w_{i',t+1}}\right) \\
&\quad + \rho\ln\left(\frac{1}{N_{j(-i,t),t}} \sum_{i'|j(i',t)=j(i,t),i'\neq i} e^{\eta^f_{j(i',t),t}+\varepsilon^f_{j(i',t),t}} e^{\rho z_{i',t-1}+\eta^w_{i',t}+\varepsilon^w_{i',t}}\right) \\
&= -\ln\left(e^{\eta^f_{j(i,t),t+1}+\rho\eta^f_{j(i,t),t}+\varepsilon^f_{j(i,t),t+1}} \frac{1}{N_{j(-i,t),t+1}} \sum_{i'|j(i',t+1)=j(i,t),i'\neq i} e^{\rho^2 z_{i',t-1}+\rho\eta^w_{i',t}+\eta^w_{i',t+1}+\varepsilon^w_{i',t+1}}\right) \\
&\quad + \rho\ln\left(e^{\eta^f_{j(i,t),t}+\varepsilon^f_{j(i,t),t}} \frac{1}{N_{j(-i,t),t}} \sum_{i'|j(i',t)=j(i,t),i'\neq i} e^{\rho z_{i',t-1}+\eta^w_{i',t}+\varepsilon^w_{i',t}}\right) \\
&= -\eta^f_{j(i,t),t+1} - \rho\eta^f_{j(i,t),t} - \varepsilon^f_{j(i,t),t+1} + \rho\eta^f_{j(i,t),t} + \rho\varepsilon^f_{j(i,t),t} \\
&\quad - \ln\left(\frac{1}{N_{j(-i,t),t+1}} \sum_{i'|j(i',t+1)=j(i,t),i'\neq i} e^{\rho^2 z_{i',t-1}+\rho\eta^w_{i',t}+\eta^w_{i',t+1}+\varepsilon^w_{i',t+1}}\right) \\
&\quad + \rho\ln\left(\frac{1}{N_{j(-i,t),t}} \sum_{i'|j(i',t)=j(i,t),i'\neq i} e^{\rho z_{i',t-1}+\eta^w_{i',t}+\varepsilon^w_{i',t}}\right) \\
&= -\eta^f_{j(i,t),t+1} - \grave{\Delta}\varepsilon^f_{j(i,t),t+1} + g_{j(-i,t+1),t+1}
\end{aligned}$$

where the function $g(\cdot)$ is shorthand for the final two terms and is a function of a combination of shocks that are independent of $(\grave{\Delta} y_{i,t}, \Delta c_{i,t})$.

We can then show our result directly by simplifying the numerator and denominator:

$$
\begin{aligned}
\text{cov}\left(\Delta c_{it}, -\grave{\Delta} y_{j(-i,t),t+1}\right) &= \text{cov}\left(\beta_\eta(\eta^w_{it} + \eta^f_{j(i,t),t}) + \beta_\varepsilon(\varepsilon^w_{it} + \varepsilon^f_{j(i,t),t}) + \zeta_{it},\right. \\
&\qquad \left. -\eta^f_{j(i,t),t+1} - \grave{\Delta}\varepsilon^f_{j(i,t),t+1} + g_{j(-i,t),t+1}\right) \\
&= \text{cov}\left(\beta_\eta(\eta^w_{it} + \eta^f_{j(i,t),t}) + \beta_\varepsilon(\varepsilon^w_{it} + \varepsilon^f_{j(i,t),t}) + \zeta_{it},\right. \\
&\qquad \left. -\eta^f_{j(i,t),t+1} - \varepsilon^f_{j(i,t),t+1} + \rho\varepsilon^f_{j(i,t),t} + g_{j(-i,t),t+1}\right) \\
&= \rho\beta_\varepsilon\sigma^2_{f\varepsilon} \\
\text{cov}\left(\grave{\Delta} y_{it}, -\grave{\Delta} y_{j(-i,t),t+1}\right) &= \text{cov}\left(\eta^w_{i,t} + \eta^f_{j(i,t),t} + \grave{\Delta}\varepsilon^w_{i,t} + \grave{\Delta}\varepsilon^f_{j(i,t),t},\right. \\
&\qquad \left. -\eta^f_{j(i,t),t+1} - \grave{\Delta}\varepsilon^f_{j(i,t),t+1} + g_{j(-i,t),t+1}\right) \\
&= \text{cov}\left(\eta^w_{i,t} + \eta^f_{j(i,t),t} + \varepsilon^w_{i,t} - \rho\varepsilon^w_{i,t-1} + \varepsilon^f_{j(i,t),t} - \rho\varepsilon^f_{j(i,t-1),t-1},\right. \\
&\qquad \left. -\eta^f_{j(i,t),t+1} - \varepsilon^f_{j(i,t),t+1} + \rho\varepsilon^f_{j(i,t),t} + g_{j(-i,t),t+1}\right) \\
&= \rho\sigma^2_{f\varepsilon}
\end{aligned}
$$

The result then follows by taking the ratio of the two covariances.

□

In order to implement this version of our estimator, we need an estimate of ρ. We can directly estimate this parameter in our sample by noting that, under equation (B.6), we have the following:

$$
\begin{aligned}
\hat{\rho}_k &= \frac{\text{cov}(y_{i,t+1}, y_{i,t-1-k})}{\text{cov}(y_{i,t}, y_{i,t-1-k})} \\
&= \frac{\text{cov}(z_{i,t+1} + \varepsilon_{i,t+1}, z_{i,t-1-k} + \varepsilon_{i,t-1-k})}{\text{cov}(z_{i,t} + \varepsilon_{i,t}, z_{i,t-1-k} + \varepsilon_{i,t-1-k})} \\
&= \frac{\text{cov}(\rho^{2+k} z_{i,t-1-k} + \sum_{s=0}^{1+k} \rho^s \eta_{i,t+1-s} + \varepsilon_{i,t+1}, z_{i,t-1-k} + \varepsilon_{i,t-1-k})}{\text{cov}(\rho^{1+k} z_{i,t-1-k} + \sum_{s=0}^{k} \rho^s \eta_{i,t-s} + \varepsilon_{i,t}, z_{i,t-1-k} + \varepsilon_{i,t-1-k})} \\
&= \frac{\rho^{2+k}\text{var}(z_{i,t-1-k})}{\rho^{1+k}\text{var}(z_{i,t-1-k})} \\
&= \rho
\end{aligned}
\tag{B.8}
$$

We estimate this parameter for $k = 0$, separately for seasonally adjusted, log-worker total income ($\rho^w = 0.898$) and for seasonally adjusted coworker pay per paycheck ($\rho^{co,ppp} = 0.954$).

As suggested by equation (B.8), the model implies that we should get similar results for different lags, i.e. values of k. However, when varying k between zero and three, our estimates range between 0.898 and 0.959 for worker total pay and 0.954 and 0.968 for coworker, pay per paycheck. For this reason, we also take an alternative approach of drawing on an external value, following Commault (2022), which uses an *annual* value of 0.95. We adjust to our monthly context by setting a *monthly* value of $0.95^{1/12} = 0.996$.

Table B-1 compares three estimates. In the first column, we repeat our main elasticity estimate from column (1) of Table 1. Columns 2 and 3 present our estimate, after accounting for a highly persistent, but not fully permanent, shocks to income. Consistent with Kaplan and Violante (2010), our estimate of the elasticity is largely robust to adjustments for a less-than-fully permanent income shock.

Table B-1: Impact of Income on Consumption

	(1)	(2)	(3)
Δ Log Income	0.221 (0.015)	0.198 (0.013)	0.196 (0.012)
Coworker Instrument	Period-Ahead Pay Per Paycheck	Period-Ahead Pay Per Paycheck	Period-Ahead Pay Per Paycheck
Seasonal Adjustment Differencing (Δ)	Yes First difference	Yes Quasi-difference	Yes Quasi-difference
Permanent Income Process			
	Unit Root $\rho = 1$	AR(1) Worker: $\rho = 0.898$ (Estimated) Co-Worker: $\rho = 0.954$ (Estimated)	AR(1) $\rho = 0.996$ (Commault 2022)

Notes: This table reports estimates of the elasticity of consumption with respect to income ($\hat{\beta}$). Column (1) repeats the estimate from Table 1, Column (1). Columns (2) and (3) account for less-than-fully persistent shocks to income. All columns use a one-period-ahead coworker instrument and are adjusted for annual seasonality, as in Table 1. Standard errors, which are clustered two ways at the household-job level, are reported in parentheses.

B.4 Serial Correlation in the Transitory Shock

Another possibility is that log income is a function of not only the current period's transitory shock, but also additional, recent, transitory shocks. For example, if our transitory component has an MA(q) structure, we will have:

$$y_{it} = z_{it} + \sum_{k=0}^{q} \theta_k \varepsilon_{i,t-k} \tag{B.9}$$

As explained in Commault (2022), HM-BPP style estimators may suffer from bias if they make use of instruments at leads shorter than $q+1$ periods ahead. In our main specifications, we assume that the transitory shock follows an MA(0) process. To assess the plausibility of that assumption, we calculate the sample covariance of current income growth of worker total monthly pay, $\Delta y_{i,t}$, and leads of the coworker, pay per paycheck growth. In column (1) of Table B-2, we plot these covariances at different horizons ranging from the concurrent growth in coworker pay per paycheck to the four-periods-ahead lead. Under an MA(0) process, we would expect covariances to be present up until the one-period-ahead lead, and then to equal zero thereafter. Our covariances do not technically reach zero, but they become an order of magnitude smaller beyond the one-period-ahead lead, giving us reassurance in our baseline approach. Because these covariances could also be affected by permanent shocks that are not unit root, we adjust for an AR(1) process, as detailed in Appendix B.3, using a quasi-difference in log income, either with an AR(1) parameter estimated in our data, or based on Commault (2022). The pattern, qualitatively consistent with an MA(0) process, continues to hold.

Table B-2: Covariance between $\Delta y_{i,t}$ and Leads of the Instrument

k	$\text{cov}(\Delta y_{i,t}, -\Delta y_{j(-i,t),t+k})$	$\text{cov}(\grave{\Delta} y_{i,t}, -\grave{\Delta} y_{j(-i,t),t+k})$	$\text{cov}(\grave{\Delta} y_{i,t}, -\grave{\Delta} y_{j(-i,t),t+k})$
0	-0.01118	-0.01109	-0.01182
1	0.00472	0.00531	0.00560
2	0.00007	0.00007	0.00014
3	-0.00003	-0.00015	-0.00010
4	0.00020	0.00006	0.00013
Permanent Income Process			
	Unit Root	AR(1)	AR(1)
	$\rho = 1$	Worker: $\rho = 0.898$ (Estimated)	$\rho = 0.996$ (Commault 2022)
		Co-Worker: $\rho = 0.954(Estimated)$	

Notes: This table reports estimates of the covariance between $\Delta y_{i,t}$, growth in worker total monthly pay, and various leads of our coworker instrument, $-\Delta y_{j(-i,t),t+k}$, including the concurrent instrument. Column (1) uses a first-difference, while Columns (2) and (3) use quasi-differences, based on adjustments for deviations from a unit root permanent shock process, as detailed in Appendix B.3. All columns are adjusted for annual seasonality, as in Table 1.

Appendix C Proofs

Result RCT (Simple RCT)

Under random assignment of a shock to log income in period t^*, we are able to identify β_ε using the following ratio:

$$\hat{\beta}_{\varepsilon,RCT} \equiv \frac{\mathbb{E}\left(c_{i,t^*}\left|D_i^\varepsilon = 1\right.\right) - \mathbb{E}\left(c_{i,t^*}\left|D_i^\varepsilon = 0\right.\right)}{\mathbb{E}\left(y_{i,t^*}^{RCT}\left|D_i^\varepsilon = 1\right.\right) - \mathbb{E}\left(y_{i,t^*}^{RCT}\left|D_i^\varepsilon = 0\right.\right)} = \beta_\varepsilon$$

where y_{it}^{RCT} is log income, inclusive of the treatment payment.

Proof. We assume that the transfer in period t^* is treated as a transitory shock , i.e:

$$\Delta c_{it} = \beta_\eta \eta_{it} + \beta_\varepsilon(\varepsilon_{it} + D_i^\varepsilon \mu_\varepsilon \mathbf{1}(t = t^*)) + \zeta_{it}$$

Note that we can write the log of consumption in period t^*, as:

$$\begin{aligned} c_{it^*} &= c_{i,t^*-1} + \Delta c_{it^*} \\ &= c_{i,t^*-1} + \beta_\eta \eta_{it^*} + \beta_\varepsilon(\varepsilon_{it^*} + D_i^\varepsilon \mu_\varepsilon) + \zeta_{it^*} \end{aligned}$$

It follows that

$$\begin{aligned} \mathbb{E}\left(c_{it^*}\left|D_i^\varepsilon = d\right.\right) &= \mathbb{E}\left(c_{i,t^*-1} + \beta_\eta \eta_{it^*} + \beta_\varepsilon \varepsilon_{it^*} + \zeta_{it^*}\left|D_i^\varepsilon = d\right.\right) + \beta_\varepsilon \mathbb{E}\left(D_i^\varepsilon\left|D_i^\varepsilon = d\right.\right)\mu_\varepsilon \\ &= \mathbb{E}\left(c_{i,t^*-1} + \beta_\eta \eta_{it^*} + \beta_\varepsilon \varepsilon_{it^*} + \zeta_{it^*}\right) + \beta_\varepsilon d \mu_\varepsilon \end{aligned}$$

where the second line follows from the independence between the treatment indicator (D_i^ε) and the initial log of consumption and all shocks, by virtue of randomization. We also have, for the log of income, inclusive of the experimental transfer:

$$\begin{aligned}\mathbb{E}\left(y_{i,t^*}^{RCT}\left|D_i^\varepsilon = d\right.\right) &= \mathbb{E}\left(z_{it^*} + \varepsilon_{it^*} + D_i^\varepsilon \mu_\varepsilon \left|D_i^\varepsilon = d\right.\right) \\ &= \mathbb{E}\left(z_{it^*} + \varepsilon_{it^*}\right) + d\mu_\varepsilon\end{aligned}$$

where the second line again follows from the randomization of D_i^ε. It then follows that:

$$\begin{aligned}\beta_{\varepsilon,RCT} &\equiv \frac{\mathbb{E}\left(c_{i,t^*}\left|D_i^\varepsilon = 1\right.\right) - \mathbb{E}\left(c_{i,t^*}\left|D_i^\varepsilon = 0\right.\right)}{\mathbb{E}\left(y_{i,t^*}^{RCT}\left|D_i^\varepsilon = 1\right.\right) - \mathbb{E}\left(y_{i,t^*}^{RCT}\left|D_i^\varepsilon = 0\right.\right)} \\ &= \frac{\mathbb{E}\left(c_{i,t^*-1} + \beta_\eta \eta_{it^*} + \beta_\varepsilon \varepsilon_{it^*} + \zeta_{it^*}\right) + \beta_\varepsilon \mu_\varepsilon - \mathbb{E}\left(c_{i,t^*-1} + \beta_\eta \eta_{it^*} + \beta_\varepsilon \varepsilon_{it^*} + \zeta_{it^*}\right)}{\mathbb{E}\left(z_{it^*} + \varepsilon_{it^*}\right) + \mu_\varepsilon - \mathbb{E}\left(z_{it^*} + \varepsilon_{it^*}\right)} \\ &= \frac{\beta_\varepsilon \mu_\varepsilon}{\mu_\varepsilon} \\ &= \beta_\varepsilon\end{aligned}$$

□

Result HM-BPP (Identification of β_ε)

If Assumptions IID and CE hold, the ratio of the covariance between consumption growth and the (negative) one-period income growth lead, and the covariance between income growth and the (negative) one-period income growth lead is equal to β_ε (Blundell, Pistaferri, and Preston 2008; Hall and Mishkin 1982):

$$\hat{\beta}_{\varepsilon,HM-BPP} \equiv \frac{\operatorname{cov}\left(\Delta c_{it}, -\Delta y_{i,t+1}\right)}{\operatorname{cov}\left(\Delta y_{it}, -\Delta y_{i,t+1}\right)} = \frac{\operatorname{cov}\left(\Delta c_{it}, \varepsilon_{it}\right)}{\operatorname{var}\left(\varepsilon_{it}\right)} = \frac{\beta_\varepsilon \sigma_\varepsilon^2}{\sigma_\varepsilon^2} = \beta_\varepsilon$$

Proof. We can show this directly by solving for the numerator and denominator:

$$\begin{aligned}\operatorname{cov}\left(\Delta c_{it}, -\Delta y_{i,t+1}\right) &= \operatorname{cov}\left(\beta_\eta \eta_{it} + \beta_\varepsilon \varepsilon_{it} + \zeta_{it}, -\eta_{i,t+1} - \varepsilon_{i,t+1} + \varepsilon_{it}\right) \\ &= \beta_\varepsilon \sigma_\varepsilon^2 \\ \operatorname{cov}\left(\Delta y_{it}, -\Delta y_{i,t+1}\right) &= \operatorname{cov}\left(\eta_{it} + \varepsilon_{it} - \varepsilon_{i,t-1}, -\eta_{i,t+1} - \varepsilon_{i,t+1} + \varepsilon_{it}\right) \\ &= \sigma_\varepsilon^2\end{aligned}$$

The result then follows by taking the ratio of the two. □

Lemma CEI (Conditional Expectation Inequality)

If X and Y are independent random variables, with strictly positive variances, and $\Pr(X > Y) \in (0, 1)$*, then:*

$$\mathbb{E}(X|X > Y) > \mathbb{E}(X) > \mathbb{E}(X|X < Y)$$

Proof. Note that:

$$\begin{aligned}
\mathbb{E}\left(X \left| X > Y\right.\right) &= \frac{\mathbb{E}\left(X\mathbf{1}\left(X > Y\right)\right)}{\Pr\left(X > Y\right)} \\
&= \frac{\mathbb{E}\left(X\mathbb{E}\left(\mathbf{1}\left(X > Y\right)\left| X\right.\right)\right)}{\Pr\left(X > Y\right)} \\
&= \frac{\mathbb{E}\left(XF_Y^-\left(X\right)\right)}{\Pr\left(X > Y\right)} \\
&= \frac{\mathbb{E}\left(X\right)\mathbb{E}\left(F_Y^-\left(X\right)\right)}{\Pr\left(X > Y\right)} + \frac{\operatorname{cov}\left(X, F_Y^-\left(X\right)\right)}{\Pr\left(X > Y\right)} \\
&= \frac{\mathbb{E}\left(X\right)\Pr\left(X > Y\right)}{\Pr\left(X > Y\right)} + \frac{\operatorname{cov}\left(X, F_Y^-\left(X\right)\right)}{\Pr\left(X > Y\right)} \\
&= \mathbb{E}\left(X\right) + \frac{\operatorname{cov}\left(X, F_Y^-\left(X\right)\right)}{\Pr\left(X > Y\right)} \\
&> \mathbb{E}\left(X\right)
\end{aligned}$$

where $F_Y^-\left(a\right) \equiv \Pr\left(Y < a\right)$, the third line follows from the independence of X and Y, the fourth line follows from the fact that $\operatorname{cov}(X,Y) = \mathbb{E}(XY) - \mathbb{E}(X)\mathbb{E}(Y)$, and the last line uses the fact that since F_Y^- is an increasing function, then the covariance term is positive, and furthermore, the covariance will be strictly positive if the variance of X is strictly positive. We can likewise show that:

$$\mathbb{E}\left(X \left| X < Y\right.\right) = \mathbb{E}\left(X\right) + \frac{\operatorname{cov}\left(X, 1 - F_Y^-\left(X\right)\right)}{\Pr\left(X < Y\right)} < \mathbb{E}\left(X\right)$$

The proof concludes by combining these results. Note, if $\Pr(X > Y) = 0$, then the first conditional expectation is not well-defined and second inequality becomes weak, and if $\Pr(X > Y) = 1$, the first inequality becomes weak and the second conditional expectation is not well-defined. □

Testable Prediction 1 (First Stage: Positive Income Difference at t^*)

If Assumption IID holds, the average difference in (normalized) log income between the treated and control group in period t^* is equal to the average difference in the transitory shock in period t^*, and is positive:

$$\mathbb{E}\left(\tilde{y}_{i,t^*} \left| D_{i,t^*} = 1\right.\right) - \mathbb{E}\left(\tilde{y}_{i,t^*} \left| D_{i,t^*} = 0\right.\right) = \mathbb{E}\left(\varepsilon_{i,t^*} \left| D_{i,t^*} = 1\right.\right) - \mathbb{E}\left(\varepsilon_{i,t^*} \left| D_{i,t^*} = 0\right.\right) > 0$$

Proof. To show the first part of the result, note that we can write normalized log income in period t^* as:

$$\begin{aligned}
\tilde{y}_{it^*} &\equiv y_{it^*} - y_{i,t^*+r} \\
&= z_{it^*} + \varepsilon_{it^*} - z_{i,t^*+r} - \varepsilon_{i,t^*+r} \\
&= \sum_{s=t^*+r+1}^{t^*} \eta_{i,s} + \varepsilon_{it^*} - \varepsilon_{i,t^*+r}
\end{aligned}$$

Taking expectations conditional on $D_{i,t^*} = 1$ we have:

$$\begin{aligned}\mathbb{E}\left(\tilde{y}_{it^*} \middle| D_{i,t^*} = 1\right) &= \mathbb{E}\left(\sum_{s=t^*+r+1}^{t^*} \eta_{i,s} + \varepsilon_{it^*} - \varepsilon_{i,t^*+r} \middle| D_{i,t^*} = 1\right) \\ &= \mathbb{E}\left(\sum_{s=t^*+r+1}^{t^*} \eta_{i,s} - \varepsilon_{i,t^*+r}\right) + \mathbb{E}\left(\varepsilon_{it^*} \middle| D_{i,t^*} = 1\right)\end{aligned}$$

where the last line follows from Assumption IID and the fact that $r < 0$. Intuitively, the lagged values of the shocks are independent of the shocks involved in the definition of D_{i,t^*}, $-\Delta y_{i,t^*+1} = -\eta_{i,t^*+1} - \varepsilon_{i,t^*+1} + \varepsilon_{i,t^*}$. Likewise, we have:

$$\mathbb{E}\left(\tilde{y}_{it^*} \middle| D_{i,t^*} = 0\right) = \mathbb{E}\left(\sum_{s=t^*+r+1}^{t^*} \eta_{i,s} - \varepsilon_{i,t^*+r}\right) + \mathbb{E}\left(\varepsilon_{it^*} \middle| D_{i,t^*} = 0\right)$$

Substituting for $\mathbb{E}\left(\tilde{y}_{it^*} \middle| D_{i,t^*} = 1\right)$ and $\mathbb{E}\left(\tilde{y}_{it^*} \middle| D_{i,t^*} = 0\right)$, the first part of the result then follows.

We next rely on Lemma CEI, defining $X = \varepsilon_{it^*}$ and $Y = M_{-\Delta y} + \eta_{i,t^*+1} + \varepsilon_{i,t^*+1}$. Note that by Assumption IID, these are independent variables, and thus:

$$\begin{aligned}\mathbb{E}\left(\varepsilon_{i,t^*} \middle| D_{i,t^*} = 1\right) - \mathbb{E}\left(\varepsilon_{i,t^*} \middle| D_{i,t^*} = 0\right) &= \mathbb{E}\left(\varepsilon_{i,t^*} \middle| \varepsilon_{it^*} > M_{-\Delta y} + \eta_{i,t^*+1} + \varepsilon_{i,t^*+1}\right) \\ &\quad - \mathbb{E}\left(\varepsilon_{i,t^*} \middle| \varepsilon_{it^*} < M_{-\Delta y} + \eta_{i,t^*+1} + \varepsilon_{i,t^*+1}\right) \\ &> 0\end{aligned}$$

where the last inequality is strict when the shocks have strictly positive variances and that by the definition of the median, $\Pr(D_{i,t^*} = 1) = \Pr(-\Delta y_{i,t^*+1} > M_{-\Delta y}) \in (0, 1)$. □

Testable Prediction 2 (Parallel Pre-event Income Trends)

If Assumption IID holds, the average difference in (normalized) log income between the treated and control group during periods before t^* is zero, i.e. there are parallel trends:

$$\mathbb{E}\left(\tilde{y}_{i,t^*+k} \middle| D_{i,t^*} = 1\right) - \mathbb{E}\left(\tilde{y}_{i,t^*+k} \middle| D_{i,t^*} = 0\right) = 0 \text{ for } k < 0$$

Proof. This follows immediately from the IID assumption. Note that normalized income in period $t^* + k$, $k < 0$ can be written as:

$$\begin{aligned}\tilde{y}_{i,t^*+k} &\equiv y_{i,t^*+k} - y_{i,t^*+r} \\ &= z_{i,t^*+k} + \varepsilon_{i,t^*+k} - z_{i,t^*+r} - \varepsilon_{i,t^*+r} \\ &= (-1)^{\mathbf{1}(r>k)} \sum_{s=t^*+\min(k,r)+1}^{t^*+\max(k,r)} \eta_{i,s} + \varepsilon_{i,t^*+k} - \varepsilon_{i,t^*+r}\end{aligned}$$

Taking expectations conditional on $D_{i,t^*} = 1$, we have:

$$\begin{aligned}\mathbb{E}\left(\tilde{y}_{i,t^*+k} \,\middle|\, D_{i,t^*} = 1\right) &= \mathbb{E}\left((-1)^{\mathbf{1}(r>k)} \sum_{s=t^*+\min(k,r)+1}^{t^*+\max(k,r)} \eta_{i,s} + \varepsilon_{i,t^*+k} - \varepsilon_{i,t^*+r} \,\middle|\, D_{i,t^*} = 1\right) \\ &= \mathbb{E}\left((-1)^{\mathbf{1}(r>k)} \sum_{s=t^*+\min(k,r)+1}^{t^*+\max(k,r)} \eta_{i,s} + \varepsilon_{i,t^*+k} - \varepsilon_{i,t^*+r}\right)\end{aligned}$$

where the second line follows from the IID assumption and the fact that $k, r < 0$. Likewise, we have:

$$\mathbb{E}\left(\tilde{y}_{i,t^*+k} \,\middle|\, D_{i,t^*} = 0\right) = \mathbb{E}\left((-1)^{\mathbf{1}(r>k)} \sum_{s=t^*+\min(k,r)+1}^{t^*+\max(k,r)} \eta_{i,s} + \varepsilon_{i,t^*+k} - \varepsilon_{i,t^*+r}\right)$$

Taking the difference between these two yields the desired result. □

Testable Prediction 3 (Negative Income Difference at $t^* + 1$)

If Assumption IID holds, the average difference in (normalized) log income between the treated and control group in period $t^* + 1$ is equal to the average difference in the permanent and transitory shocks in period $t^* + 1$, and both average differences are negative.

$$\begin{aligned}\mathbb{E}\left(\tilde{y}_{i,t^*+1} \,\middle|\, D_{i,t^*} = 1\right) - \mathbb{E}\left(\tilde{y}_{i,t^*+1} \,\middle|\, D_{i,t^*} = 0\right) &= \underbrace{\mathbb{E}\left(\eta_{i,t^*+1} \,\middle|\, D_{i,t^*} = 1\right) - \mathbb{E}\left(\eta_{i,t^*+1} \,\middle|\, D_{i,t^*} = 0\right)}_{<0} \\ &\quad + \underbrace{\mathbb{E}\left(\varepsilon_{i,t^*+1} \,\middle|\, D_{i,t^*} = 1\right) - \mathbb{E}\left(\varepsilon_{i,t^*+1} \,\middle|\, D_{i,t^*} = 0\right)}_{<0} \\ &< 0\end{aligned}$$

Proof. Recall, that we can write normalized income in period $t^* + 1$ as:

$$\begin{aligned}\tilde{y}_{i,t^*+1} &\equiv y_{i,t^*+1} - y_{i,t^*+r} \\ &= z_{i,t^*+1} + \varepsilon_{i,t^*+1} - z_{i,t^*+r} - \varepsilon_{i,t^*+r} \\ &= \sum_{s=t^*+r+1}^{t^*+1} \eta_{i,s} + \varepsilon_{i,t^*+1} - \varepsilon_{i,t^*+r}\end{aligned}$$

Taking expectations conditional on $D_{i,t^*} = 1$, we have:

$$\begin{aligned}\mathbb{E}\left(\tilde{y}_{i,t^*+1} \,\middle|\, D_{i,t^*} = 1\right) &= \mathbb{E}\left(\sum_{s=t^*+r+1}^{t^*+1} \eta_{i,s} + \varepsilon_{i,t^*+1} - \varepsilon_{i,t^*+r} \,\middle|\, D_{i,t^*} = 1\right) \\ &= \mathbb{E}\left(\sum_{s=t^*+r+1}^{t^*} \eta_{i,s} - \varepsilon_{i,t^*+r}\right) + \mathbb{E}\left(\eta_{i,t^*+1} + \varepsilon_{i,t^*+1} \,\middle|\, D_{i,t^*} = 1\right)\end{aligned}$$

where the last line follows from Assumption IID and the fact that $r < 0$. Intuitively, the lagged values of the shocks are independent of the shocks involved in the definition of D_{i,t^*}, $-\Delta y_{i,t^*+1} =$

$-\eta_{i,t^*+1} - \varepsilon_{i,t^*+1} + \varepsilon_{i,t^*}$. Likewise, we have:

$$\mathbb{E}\left(\tilde{y}_{i,t^*+1} \,\middle|\, D_{i,t^*} = 0\right) = \mathbb{E}\left(\sum_{s=t^*+r+1}^{t^*} \eta_{i,s} - \varepsilon_{i,t^*+r}\right) + \mathbb{E}\left(\eta_{i,t^*+1} + \varepsilon_{i,t^*+1} \,\middle|\, D_{i,t^*} = 0\right)$$

Taking the difference between the two yields the main expression.

To show that each component is less than zero, we can again appeal to Lemma CEI. First, we define $X = \eta_{i,t^*+1}$ and $Y = \varepsilon_{i,t^*} - \varepsilon_{i,t^*+1} - M_{-\Delta y}$, and note that the two are independent. Second, we redefine $X = \varepsilon_{i,t^*+1}$ and $Y = \varepsilon_{i,t^*} - \eta_{i,t^*+1} - M_{-\Delta y}$, and again note that the two are independent. In both cases, $D_{i,t^*} = 1 \Rightarrow X < Y$. The inequalities then follow from the Lemma and the fact that by the definition of the median, $\Pr(D_{i,t^*} = 1) = \Pr(-\Delta y_{i,t^*+1} > M_{-\Delta y}) \in (0, 1)$. □

Testable Prediction 4 (Negative Income Difference at $t^* + 2$ Onward)

If Assumption IID holds, the difference in (normalized) log income between the treated and control group in periods $t > t^* + 1$ is equal to the average difference in the permanent shock in period $t^* + 1$, and this difference is negative.

$$\begin{aligned}\mathbb{E}\left(\tilde{y}_{i,t^*+k} \,\middle|\, D_{i,t^*} = 1\right) - \mathbb{E}\left(\tilde{y}_{i,t^*+k} \,\middle|\, D_{i,t^*} = 0\right) =& \mathbb{E}\left(\eta_{i,t^*+1} \,\middle|\, D_{i,t^*} = 1\right) - \mathbb{E}\left(\eta_{i,t^*+1} \,\middle|\, D_{i,t^*} = 0\right) \\ <& 0 \quad \textit{for } k > 1\end{aligned}$$

Proof. Recall, that we can write normalized income in period $t^* + k$, $k > 1$ as:

$$\begin{aligned}\tilde{y}_{i,t^*+k} &\equiv y_{i,t^*+k} - y_{i,t^*+r} \\ &= z_{i,t^*+k} + \varepsilon_{i,t^*+k} - z_{i,t^*+r} - \varepsilon_{i,t^*+r} \\ &= \sum_{s=t^*+r+1}^{t^*+k} \eta_{i,s} + \varepsilon_{i,t^*+k} - \varepsilon_{i,t^*+r}\end{aligned}$$

Taking expectations conditional on $D_{i,t^*} = 1$, we have:

$$\begin{aligned}\mathbb{E}\left(\tilde{y}_{i,t^*+k} \,\middle|\, D_{i,t^*} = 1\right) &= \mathbb{E}\left(\sum_{s=t^*+r+1}^{t^*+k} \eta_{i,s} + \varepsilon_{i,t^*+k} - \varepsilon_{i,t^*+r} \,\middle|\, D_{i,t^*} = 1\right) \\ &= \mathbb{E}\left(\sum_{s=t^*+r+1}^{t^*} \eta_{i,s} + \sum_{s=t^*+2}^{t^*+k} \eta_{i,s} + \varepsilon_{i,t^*+k} - \varepsilon_{i,t^*+r}\right) + \mathbb{E}\left(\eta_{i,t^*+1} \,\middle|\, D_{i,t^*} = 1\right)\end{aligned}$$

where the last line follows from Assumption IID and the fact that $r < 0$. Intuitively, the lagged and lead values of the shocks are independent of the shocks involved in the definition of D_{i,t^*}, $-\Delta y_{i,t^*+1} = -\eta_{i,t^*+1} - \varepsilon_{i,t^*+1} + \varepsilon_{i,t^*}$. Likewise, we have:

$$\mathbb{E}\left(\tilde{y}_{i,t^*+k} \,\middle|\, D_{i,t^*} = 0\right) = \mathbb{E}\left(\sum_{s=t^*+r}^{t^*} \eta_{i,s} + \sum_{s=t^*+2}^{t^*+k} \eta_{i,s} + \varepsilon_{i,t^*+k} - \varepsilon_{i,t^*+r}\right) + \mathbb{E}\left(\eta_{i,t^*+1} \,\middle|\, D_{i,t^*} = 0\right)$$

Taking the difference between the two yields the main expression.

To show that difference is less than zero, we can again appeal to Lemma CEI. First, we define $X = \eta_{i,t^*+1}$ and $Y = \varepsilon_{i,t^*} - \varepsilon_{i,t^*+1} - M_{-\Delta y}$, and note that the two are independent. Furthermore, we have that $D_{i,t^*} = 1 \Rightarrow X < Y$. The inequality then follows from the Lemma and the fact that by the definition of the median, $\Pr(D_{i,t^*} = 1) = \Pr(-\Delta y_{i,t^*+1} > M_{-\Delta y}) \in (0, 1)$. □

Testable Prediction 5 (Parallel Pre-event Consumption Trends)

If Assumptions IID and CE hold, the average difference in (normalized) log consumption between the treated and control group is zero during periods $t < t^*$, i.e. there are parallel trends:

$$\mathbb{E}\left(\tilde{c}_{i,t^*+k} \,\middle|\, D_{i,t^*} = 1\right) - \mathbb{E}\left(\tilde{c}_{i,t^*+k} \,\middle|\, D_{i,t^*} = 0\right) = 0 \text{ for } k < 0$$

Proof. Note that normalized consumption in period $t^* + k$, $k < 0$ can be written as:

$$\begin{aligned}
\tilde{c}_{i,t^*+k} &= c_{i,t^*+k} - c_{i,t^*+r} \\
&= (-1)^{\mathbf{1}(r>k)} \sum_{s=t^*+\min(k,r)+1}^{t^*+\max(k,r)} \Delta c_{is} \\
&= (-1)^{\mathbf{1}(r>k)} \sum_{s=t^*+\min(k,r)+1}^{t^*+\max(k,r)} (\beta_\eta \eta_{is} + \beta_\varepsilon \varepsilon_{is} + \zeta_{is})
\end{aligned}$$

Taking expectations conditional on $D_{i,t^*} = 1$, we have:

$$\begin{aligned}
\mathbb{E}\left(\tilde{c}_{i,t^*+k} \,\middle|\, D_{i,t^*} = 1\right) &= \mathbb{E}\left((-1)^{\mathbf{1}(r>k)} \sum_{s=t^*+\min(k,r)+1}^{t^*+\max(k,r)} (\beta_\eta \eta_{is} + \beta_\varepsilon \varepsilon_{is} + \zeta_{is}) \,\middle|\, D_{i,t^*} = 1\right) \\
&= \mathbb{E}\left((-1)^{\mathbf{1}(r>k)} \sum_{s=t^*+\min(k,r)+1}^{t^*+\max(k,r)} (\beta_\eta \eta_{is} + \beta_\varepsilon \varepsilon_{is} + \zeta_{is})\right)
\end{aligned}$$

where the second line follows from the independence of the lagged shocks and those involved in the definition of D_{i,t^*}, $-\Delta y_{i,t^*+1} = -\eta_{i,t^*+1} - \varepsilon_{i,t^*+1} + \varepsilon_{it^*}$. We can likewise show that:

$$\mathbb{E}\left(\tilde{c}_{i,t^*+k} \,\middle|\, D_{i,t^*} = 0\right) = \mathbb{E}\left((-1)^{\mathbf{1}(r>k)} \sum_{s=t^*+\min(k,r)+1}^{t^*+\max(k,r)} (\beta_\eta \eta_{is} + \beta_\varepsilon \varepsilon_{is} + \zeta_{is})\right)$$

Taking the difference between the two yields the desired result. □

Result Pooled (Parameter Equivalence of Pooled RCT and Simple RCT)

Using the Pooled RCT, we are able to identify β_ε using the following ratio:

$$\begin{aligned}
\hat{\beta}_{\varepsilon,Pooled} &\equiv \frac{\mathbb{E}\left(c_{i,t^*} \,\middle|\, D_i^{\text{Pooled}} = 1\right) - \mathbb{E}\left(c_{i,t^*} \,\middle|\, D_i^{\text{Pooled}} = 0\right)}{\mathbb{E}\left(y_{i,t^*}^{RCT} \,\middle|\, D_i^{\text{Pooled}} = 1\right) - \mathbb{E}\left(y_{i,t^*}^{RCT} \,\middle|\, D_i^{\text{Pooled}} = 0\right)} \\
&= \frac{\frac{1}{3}\left[\mathbb{E}\left(c_{i,t^*} \,\middle|\, \tilde{D}_i^\varepsilon = 1\right) - \mathbb{E}\left(c_{i,t^*} \,\middle|\, \tilde{D}_i^\varepsilon = 0\right)\right]}{\frac{1}{3}\left[\mathbb{E}\left(y_{i,t^*}^{RCT} \,\middle|\, \tilde{D}_i^\varepsilon = 1\right) - \mathbb{E}\left(y_{i,t^*}^{RCT} \,\middle|\, \tilde{D}_i^\varepsilon = 0\right)\right]} \\
&= \beta_\varepsilon
\end{aligned}$$

where y_{it}^{RCT} is log income, inclusive of the treatment payment.

Proof. First note that for the $\beta_{-\varepsilon}$ RCT and the β_η RCTs, since the treatment is randomized and transfers are not revealed until period t^*+1, we have:

$$\begin{aligned}
\mathbb{E}\left(c_{i,t^*}\left|D_i^{-\varepsilon}=1\right.\right)-\mathbb{E}\left(c_{i,t^*}\left|D_i^{-\varepsilon}=0\right.\right)&=\mathbb{E}\left(c_{i,t^*}\right)-\mathbb{E}\left(c_{i,t^*}\right)=0\\
\mathbb{E}\left(c_{i,t^*}\left|D_i^{\eta}=1\right.\right)-\mathbb{E}\left(c_{i,t^*}\left|D_i^{\eta}=0\right.\right)&=\mathbb{E}\left(c_{i,t^*}\right)-\mathbb{E}\left(c_{i,t^*}\right)=0\\
\mathbb{E}\left(y_{i,t^*}^{RCT}\left|D_i^{-\varepsilon}=1\right.\right)-\mathbb{E}\left(y_{i,t^*}^{RCT}\left|D_i^{-\varepsilon}=0\right.\right)&=\mathbb{E}\left(y_{i,t^*}^{RCT}\right)-\mathbb{E}\left(y_{i,t^*}^{RCT}\right)=0\\
\mathbb{E}\left(y_{i,t^*}^{RCT}\left|D_i^{\eta}=1\right.\right)-\mathbb{E}\left(y_{i,t^*}^{RCT}\left|D_i^{\eta}=0\right.\right)&=\mathbb{E}\left(y_{i,t^*}^{RCT}\right)-\mathbb{E}\left(y_{i,t^*}^{RCT}\right)=0
\end{aligned}$$

The numerator of the pooled estimator is therefore:

$$\begin{aligned}
\mathbb{E}\left(c_{i,t^*}\left|D_i^{\text{Pooled}}=1\right.\right)-\mathbb{E}\left(c_{i,t^*}\left|D_i^{\text{Pooled}}=0\right.\right)&=\frac{1}{3}\left[\mathbb{E}\left(c_{i,t^*}\left|D_i^{\varepsilon}=1\right.\right)-\mathbb{E}\left(c_{i,t^*}\left|D_i^{\varepsilon}=0\right.\right)\right]\\
&\quad+\frac{1}{3}\left[\mathbb{E}\left(c_{i,t^*}\left|D_i^{-\varepsilon}=1\right.\right)-\mathbb{E}\left(c_{i,t^*}\left|D_i^{-\varepsilon}=0\right.\right)\right]\\
&\quad+\frac{1}{3}\left[\mathbb{E}\left(c_{i,t^*}\left|D_i^{\eta}=1\right.\right)-\mathbb{E}\left(c_{i,t^*}\left|D_i^{\eta}=0\right.\right)\right]\\
&=\frac{1}{3}\left[\mathbb{E}\left(c_{i,t^*}\left|D_i^{\varepsilon}=1\right.\right)-\mathbb{E}\left(c_{i,t^*}\left|D_i^{\varepsilon}=0\right.\right)\right]
\end{aligned}$$

Similarly, the denominator of the estimator is:

$$\begin{aligned}
\mathbb{E}\left(y_{i,t^*}^{RCT}\left|D_i^{\text{Pooled}}=1\right.\right)-\mathbb{E}\left(y_{i,t^*}^{RCT}\left|D_i^{\text{Pooled}}=0\right.\right)&=\frac{1}{3}\left[\mathbb{E}\left(y_{i,t^*}^{RCT}\left|D_i^{\varepsilon}=1\right.\right)-\mathbb{E}\left(y_{i,t^*}^{RCT}\left|D_i^{\varepsilon}=0\right.\right)\right]\\
&\quad+\frac{1}{3}\left[\mathbb{E}\left(y_{i,t^*}^{RCT}\left|D_i^{-\varepsilon}=1\right.\right)-\mathbb{E}\left(y_{i,t^*}^{RCT}\left|D_i^{-\varepsilon}=0\right.\right)\right]\\
&\quad+\frac{1}{3}\left[\mathbb{E}\left(y_{i,t^*}^{RCT}\left|D_i^{\eta}=1\right.\right)-\mathbb{E}\left(y_{i,t^*}^{RCT}\left|D_i^{\eta}=0\right.\right)\right]\\
&=\frac{1}{3}\left[\mathbb{E}\left(y_{i,t^*}^{RCT}\left|D_i^{\varepsilon}=1\right.\right)-\mathbb{E}\left(y_{i,t^*}^{RCT}\left|D_i^{\varepsilon}=0\right.\right)\right]
\end{aligned}$$

Taking the ratio of the two, we then have:

$$\begin{aligned}
\hat{\beta}_{\varepsilon,Pooled}&\equiv\frac{\mathbb{E}\left(c_{i,t^*}\left|D_i^{\text{Pooled}}=1\right.\right)-\mathbb{E}\left(c_{i,t^*}\left|D_i^{\text{Pooled}}=0\right.\right)}{\mathbb{E}\left(y_{i,t^*}^{RCT}\left|D_i^{\text{Pooled}}=1\right.\right)-\mathbb{E}\left(y_{i,t^*}^{RCT}\left|D_i^{\text{Pooled}}=0\right.\right)}\\
&=\frac{\frac{1}{3}\left[\mathbb{E}\left(c_{i,t^*}\left|D_i^{\varepsilon}=1\right.\right)-\mathbb{E}\left(c_{i,t^*}\left|D_i^{\varepsilon}=0\right.\right)\right]}{\frac{1}{3}\left[\mathbb{E}\left(y_{i,t^*}^{RCT}\left|D_i^{\varepsilon}=1\right.\right)-\mathbb{E}\left(y_{i,t^*}^{RCT}\left|D_i^{\varepsilon}=0\right.\right)\right]}\\
&=\beta_\varepsilon
\end{aligned}$$

where the last line follows from Result RCT above. □

Result Pooled RCT-HMBPP (Income Equivalence of Pooled RCT and HM-BPP)

If Assumption IID holds, the average difference in (normalized) log income based on the binary instrument, $D_{i,t^*}\equiv\mathbf{1}(-\Delta y_{i,t^*+1}<M_{-\Delta y})$*, is equivalent to the average difference in log income between the treatment and control groups of the pooled experiment, in all time periods:*

$$\mathbb{E}\left(\tilde{y}_{i,t}\left|D_{i,t^*}=1\right.\right)-\mathbb{E}\left(\tilde{y}_{i,t}\left|D_{i,t^*}=0\right.\right)=\mathbb{E}\left(y_{i,t}^{RCT}\left|D_i^{\text{Pooled}}=1\right.\right)-\mathbb{E}\left(y_{i,t}^{RCT}\left|D_i^{\text{Pooled}}=0\right.\right)\forall t$$

where the treatments in each of the three sub-RCTs are defined as:

$$
\begin{aligned}
\mu_\varepsilon &\equiv 3\left(\mathbb{E}\left(\varepsilon_{i,t^*}\left|D_{i,t^*}=1\right.\right)-\mathbb{E}\left(\varepsilon_{i,t^*}\left|D_{i,t^*}=0\right.\right)\right)\\
\mu_{-\varepsilon} &\equiv 3\left(\mathbb{E}\left(\varepsilon_{i,t^*+1}\left|D_{i,t^*}=1\right.\right)-\mathbb{E}\left(\varepsilon_{i,t^*+1}\left|D_{i,t^*}=0\right.\right)\right)\\
\mu_\eta &\equiv 3\left(\mathbb{E}\left(\eta_{i,t^*+1}\left|D_{i,t^*}=1\right.\right)-\mathbb{E}\left(\eta_{i,t^*+1}\left|D_{i,t^*}=0\right.\right)\right)
\end{aligned}
$$

and y_{it}^{RCT} is log income, inclusive of the treatment payment.

Proof. Note, that for periods prior to t^*, by way of randomization and the fact that none of the experimental transfers have been announced, income is independent of treatment status in each of the sub-RCTs. That is, for $k<0$, we have:

$$
\begin{aligned}
\mathbb{E}\left(y_{i,t^*+k}^{RCT}\left|D_i^\varepsilon=1\right.\right)-\mathbb{E}\left(y_{i,t^*+k}^{RCT}\left|D_i^\varepsilon=0\right.\right)&=\mathbb{E}\left(y_{i,t^*+k}^{RCT}\right)-\mathbb{E}\left(y_{i,t^*+k}^{RCT}\right)=0\\
\mathbb{E}\left(y_{i,t^*+k}^{RCT}\left|D_i^{-\varepsilon}=1\right.\right)-\mathbb{E}\left(y_{i,t^*+k}^{RCT}\left|D_i^{-\varepsilon}=0\right.\right)&=\mathbb{E}\left(y_{i,t^*+k}^{RCT}\right)-\mathbb{E}\left(y_{i,t^*+k}^{RCT}\right)=0\\
\mathbb{E}\left(y_{i,t^*+k}^{RCT}\left|D_i^\eta=1\right.\right)-\mathbb{E}\left(y_{i,t^*+k}^{RCT}\left|D_i^\eta=0\right.\right)&=\mathbb{E}\left(y_{i,t^*+k}^{RCT}\right)-\mathbb{E}\left(y_{i,t^*+k}^{RCT}\right)=0
\end{aligned}
$$

It follows that the pooled RCT, which is a simple average of these three RCTs, exhibits the same pattern for $k<0$:

$$
\begin{aligned}
\mathbb{E}\left(y_{i,t^*+k}^{RCT}\left|D_i^{\text{Pooled}}=1\right.\right)-\mathbb{E}\left(y_{i,t^*+k}^{RCT}\left|D_i^{\text{Pooled}}=0\right.\right)&=\frac{1}{3}\left[\mathbb{E}\left(y_{i,t^*+k}^{RCT}\left|D_i^\varepsilon=1\right.\right)-\mathbb{E}\left(y_{i,t^*+k}^{RCT}\left|D_i^\varepsilon=0\right.\right)\right]\\
&\quad+\frac{1}{3}\left[\mathbb{E}\left(y_{i,t^*+k}^{RCT}\left|D_i^{-\varepsilon}=1\right.\right)-\mathbb{E}\left(y_{i,t^*+k}^{RCT}\left|D_i^{-\varepsilon}=0\right.\right)\right]\\
&\quad+\frac{1}{3}\left[\mathbb{E}\left(y_{i,t^*+k}^{RCT}\left|D_i^\eta=1\right.\right)-\mathbb{E}\left(y_{i,t^*+k}^{RCT}\left|D_i^\eta=0\right.\right)\right]\\
&=0\\
&=\mathbb{E}\left(\tilde{y}_{i,t^*+k}\left|D_{i,t^*}=1\right.\right)-\mathbb{E}\left(\tilde{y}_{i,t^*+k}\left|D_{i,t^*}=0\right.\right)
\end{aligned}
$$

where the last line follows from Testable Prediction 2, thus establishing our desired result for periods t^*+k when $k<0$.

For period t^*, note that, again, because the $\beta_{-\varepsilon}$ and the β_η RCTs are randomized and transfers are not revealed until period t^*+1, we have:

$$
\begin{aligned}
\mathbb{E}\left(y_{i,t^*}^{RCT}\left|D_i^{-\varepsilon}=1\right.\right)-\mathbb{E}\left(y_{i,t^*}^{RCT}\left|D_i^{-\varepsilon}=0\right.\right)&=\mathbb{E}\left(y_{i,t^*}^{RCT}\right)-\mathbb{E}\left(y_{i,t^*}^{RCT}\right)=0\\
\mathbb{E}\left(y_{i,t^*}^{RCT}\left|D_i^\eta=1\right.\right)-\mathbb{E}\left(y_{i,t^*}^{RCT}\left|D_i^\eta=0\right.\right)&=\mathbb{E}\left(y_{i,t^*}^{RCT}\right)-\mathbb{E}\left(y_{i,t^*}^{RCT}\right)=0
\end{aligned}
$$

In addition, for the β_ε RCT, we can write log income, inclusive of the experimental transfer, as:

$$
\begin{aligned}
\mathbb{E}\left(y_{it}^{RCT}\left|D_i^\varepsilon=d\right.\right)&=\mathbb{E}\left(z_{it}+\varepsilon_{it}+D_i^\varepsilon\mu_\varepsilon\mathbf{1}(t=t^*)\left|D_i^\varepsilon=d\right.\right)\\
&=\mathbb{E}\left(z_{it}+\varepsilon_{it}\right)+d\mu_\varepsilon\mathbf{1}(t=t^*)
\end{aligned}
$$

It follows that:

$$\mathbb{E}\left(y_{i,t^*}^{RCT}\left|D_i^\varepsilon=1\right.\right)-\mathbb{E}\left(y_{i,t^*}^{RCT}\left|D_i^\varepsilon=0\right.\right)=\mathbb{E}\left(z_{it^*}+\varepsilon_{it^*}\right)+\mu_\varepsilon-\mathbb{E}\left(z_{it^*}+\varepsilon_{it^*}\right)$$
$$=\mu_\varepsilon$$

We can therefore show that:

$$\begin{aligned}
\mathbb{E}\left(y_{i,t^*}^{RCT}\left|D_i^{\text{Pooled}}=1\right.\right)-\mathbb{E}\left(y_{i,t^*}^{RCT}\left|D_i^{\text{Pooled}}=0\right.\right) &= \frac{1}{3}\left[\mathbb{E}\left(y_{i,t^*}^{RCT}\left|D_i^\varepsilon=1\right.\right)-\mathbb{E}\left(y_{i,t^*}^{RCT}\left|D_i^\varepsilon=0\right.\right)\right]\\
&\quad+\frac{1}{3}\left[\mathbb{E}\left(y_{i,t^*}^{RCT}\left|D_i^{-\varepsilon}=1\right.\right)-\mathbb{E}\left(y_{i,t^*}^{RCT}\left|D_i^{-\varepsilon}=0\right.\right)\right]\\
&\quad+\frac{1}{3}\left[\mathbb{E}\left(y_{i,t^*}^{RCT}\left|D_i^\eta=1\right.\right)-\mathbb{E}\left(y_{i,t^*}^{RCT}\left|D_i^\eta=0\right.\right)\right]\\
&= \frac{1}{3}\left[\mathbb{E}\left(y_{i,t^*}^{RCT}\left|D_i^\varepsilon=1\right.\right)-\mathbb{E}\left(y_{i,t^*}^{RCT}\left|D_i^\varepsilon=0\right.\right)\right]\\
&= \frac{1}{3}\left[\mu_\varepsilon\right]\\
&= \frac{1}{3}\left[3\left(\mathbb{E}\left(\varepsilon_{i,t^*}\left|D_{i,t^*}=1\right.\right)-\mathbb{E}\left(\varepsilon_{i,t^*}\left|D_{i,t^*}=0\right.\right)\right)\right]\\
&= \mathbb{E}\left(\varepsilon_{i,t}\left|D_{i,t^*}=1\right.\right)-\mathbb{E}\left(\varepsilon_{i,t}\left|D_{i,t^*}=0\right.\right)\\
&= \mathbb{E}\left(\tilde{y}_{i,t^*}\left|D_{i,t^*}=1\right.\right)-\mathbb{E}\left(\tilde{y}_{i,t^*}\left|D_{i,t^*}=0\right.\right)
\end{aligned}$$

where the second-to-last line follows from the specified value for μ_ε in the result. The last line follows from Testable Prediction 1, establishing the desired result for period t^*.

For period t^*+1, note that, because the β_ε treatment is randomized, and no transfers are exchanged beyond period t^*, we have:

$$\mathbb{E}\left(y_{i,t^*+1}^{RCT}\left|D_i^\varepsilon=1\right.\right)-\mathbb{E}\left(y_{i,t^*+1}^{RCT}\left|D_i^\varepsilon=0\right.\right)=\mathbb{E}\left(y_{i,t^*+1}^{RCT}\right)-\mathbb{E}\left(y_{i,t^*+1}^{RCT}\right)=0$$

Next, note that income for the $\beta_{-\varepsilon}$ and β_η RCTs can be written, respectively, as:

$$\begin{aligned}
\mathbb{E}\left(y_{it}^{RCT}\left|D_i^{-\varepsilon}=d\right.\right) &= \mathbb{E}\left(z_{it}+\varepsilon_{it}+D_i^{-\varepsilon}\mu_{-\varepsilon}\mathbf{1}(t=t^*+1)\left|D_i^{-\varepsilon}=d\right.\right)\\
&= \mathbb{E}\left(z_{it}+\varepsilon_{it}\right)+d\mu_{-\varepsilon}\mathbf{1}(t=t^*+1)\\
\mathbb{E}\left(y_{it}^{RCT}\left|D_i^{\eta}=d\right.\right) &= \mathbb{E}\left(z_{it}+\varepsilon_{it}+D_i^{\eta}\mu_{\eta}\mathbf{1}(t\geq t^*+1)\left|D_i^{\eta}=d\right.\right)\\
&= \mathbb{E}\left(z_{it}+\varepsilon_{it}\right)+d\mu_{\eta}\mathbf{1}(t\geq t^*+1)
\end{aligned}$$

We therefore have, for the $\beta_{-\varepsilon}$ and β_η RCTs, in period t^*+1:

$$\begin{aligned}
\mathbb{E}\left(y_{i,t^*+1}^{RCT}\left|D_i^{-\varepsilon}=1\right.\right)-\mathbb{E}\left(y_{i,t^*+1}^{RCT}\left|D_i^{-\varepsilon}=0\right.\right) &= \mathbb{E}\left(z_{i,t^*+1}+\varepsilon_{i,t^*+1}\right)+\mu_{-\varepsilon}-\mathbb{E}\left(z_{i,t^*+1}+\varepsilon_{i,t^*+1}\right)\\
&= \mu_{-\varepsilon}\\
\mathbb{E}\left(y_{i,t^*+1}^{RCT}\left|D_i^{\eta}=1\right.\right)-\mathbb{E}\left(y_{i,t^*+1}^{RCT}\left|D_i^{\eta}=0\right.\right) &= \mathbb{E}\left(z_{i,t^*+1}+\varepsilon_{i,t^*+1}\right)+\mu_{\eta}-\mathbb{E}\left(z_{i,t^*+1}+\varepsilon_{i,t^*+1}\right)\\
&= \mu_{\eta}
\end{aligned}$$

It follows that, for the pooled RCT in period $t^* + 1$, we have:

$$\begin{aligned}
\mathbb{E}\left(y_{i,t^*+1}^{RCT} \middle| D_i^{\text{Pooled}} = 1\right) - \mathbb{E}\left(y_{i,t^*+1}^{RCT} \middle| D_i^{\text{Pooled}} = 0\right) &= \frac{1}{3}\left[\mathbb{E}\left(y_{i,t^*+1}^{RCT} \left| D_i^{\varepsilon} = 1\right.\right) - \mathbb{E}\left(y_{i,t^*+1}^{RCT} \left| D_i^{\varepsilon} = 0\right.\right)\right] \\
&\quad + \frac{1}{3}\left[\mathbb{E}\left(y_{i,t^*+1}^{RCT} \left| D_i^{-\varepsilon} = 1\right.\right) - \mathbb{E}\left(y_{i,t^*+1}^{RCT} \left| D_i^{-\varepsilon} = 0\right.\right)\right] \\
&\quad + \frac{1}{3}\left[\mathbb{E}\left(y_{i,t^*+1}^{RCT} \left| D_i^{\eta} = 1\right.\right) - \mathbb{E}\left(y_{i,t^*+1}^{RCT} \left| D_i^{\eta} = 0\right.\right)\right] \\
&= \frac{1}{3}[0] + \frac{1}{3}[\mu_{-\varepsilon}] + \frac{1}{3}[\mu_{\eta}] \\
&= \frac{1}{3}\left[3\left(\mathbb{E}\left(\varepsilon_{i,t^*+1} \left| D_{i,t^*} = 1\right.\right) - \mathbb{E}\left(\varepsilon_{i,t^*+1} \left| D_{i,t^*} = 0\right.\right)\right)\right] \\
&\quad + \frac{1}{3}\left[3\left(\mathbb{E}\left(\eta_{i,t^*+1} \left| D_{i,t^*} = 1\right.\right) - \mathbb{E}\left(\eta_{i,t^*+1} \left| D_{i,t^*} = 0\right.\right)\right)\right] \\
&= \mathbb{E}\left(\varepsilon_{i,t^*+1} \left| D_{i,t^*} = 1\right.\right) - \mathbb{E}\left(\varepsilon_{i,t^*+1} \left| D_{i,t^*} = 0\right.\right) \\
&\quad + \mathbb{E}\left(\eta_{i,t^*+1} \left| D_{i,t^*} = 1\right.\right) - \mathbb{E}\left(\eta_{i,t^*+1} \left| D_{i,t^*} = 0\right.\right) \\
&= \mathbb{E}\left(\tilde{y}_{i,t^*+1} \left| D_{i,t^*} = 1\right.\right) - \mathbb{E}\left(\tilde{y}_{i,t^*+1} \left| D_{i,t^*} = 0\right.\right)
\end{aligned}$$

where the second to last line follows from the values of $\mu_{-\varepsilon}$ and μ_{η} set in the statement of the result. The last line follows from Testable Prediction 3, yielding the desired result for period $t^* + 1$.

Finally, for all periods $t^* + k$ where $k > 1$, we have the following, by randomization of treatment assignment and a lack of any further transfer for the β_{ε} and $\beta_{-\varepsilon}$ RCTs:

$$\begin{aligned}
\mathbb{E}\left(y_{i,t^*+k}^{RCT} \left| D_i^{\varepsilon} = 1\right.\right) - \mathbb{E}\left(y_{i,t^*+k}^{RCT} \left| D_i^{\varepsilon} = 0\right.\right) &= \mathbb{E}\left(y_{i,t^*+k}^{RCT}\right) - \mathbb{E}\left(y_{i,t^*+k}^{RCT}\right) = 0 \\
\mathbb{E}\left(y_{i,t^*+k}^{RCT} \left| D_i^{-\varepsilon} = 1\right.\right) - \mathbb{E}\left(y_{i,t^*+k}^{RCT} \left| D_i^{-\varepsilon} = 0\right.\right) &= \mathbb{E}\left(y_{i,t^*+k}^{RCT}\right) - \mathbb{E}\left(y_{i,t^*+k}^{RCT}\right) = 0
\end{aligned}$$

Using our expression for y_{it}^{RCT} for the β_{η} RCT from above, we have, for $k > 1$:

$$\begin{aligned}
\mathbb{E}\left(y_{i,t^*+k}^{RCT} \left| D_i^{\eta} = 1\right.\right) - \mathbb{E}\left(y_{i,t^*+k}^{RCT} \left| D_i^{\eta} = 0\right.\right) &= \mathbb{E}\left(z_{i,t^*+k} + \varepsilon_{i,t^*+k}\right) + \mu_{\eta} - \mathbb{E}\left(z_{i,t^*+k} + \varepsilon_{i,t^*+k}\right) \\
&= \mu_{\eta}
\end{aligned}$$

It follows, then, for periods $t^* + k$ where $k > 1$, we have the following for the pooled RCT:

$$\begin{aligned}
\mathbb{E}\left(y_{i,t^*+k}^{RCT} \middle| D_i^{\text{Pooled}} = 1\right) - \mathbb{E}\left(y_{i,t^*+k}^{RCT} \middle| D_i^{\text{Pooled}} = 0\right) &= \frac{1}{3}\left[\mathbb{E}\left(y_{i,t^*+k}^{RCT} \left| D_i^{\varepsilon} = 1\right.\right) - \mathbb{E}\left(y_{i,t^*+k}^{RCT} \left| D_i^{\varepsilon} = 0\right.\right)\right] \\
&\quad + \frac{1}{3}\left[\mathbb{E}\left(y_{i,t^*+k}^{RCT} \left| D_i^{-\varepsilon} = 1\right.\right) - \mathbb{E}\left(y_{i,t^*+k}^{RCT} \left| D_i^{-\varepsilon} = 0\right.\right)\right] \\
&\quad + \frac{1}{3}\left[\mathbb{E}\left(y_{i,t^*+k}^{RCT} \left| D_i^{\eta} = 1\right.\right) - \mathbb{E}\left(y_{i,t^*+k}^{RCT} \left| D_i^{\eta} = 0\right.\right)\right] \\
&= \frac{1}{3}[0] + \frac{1}{3}[0] + \frac{1}{3}[\mu_{\eta}] \\
&= \frac{1}{3}\left[3\left(\mathbb{E}\left(\eta_{i,t^*+1} \left| D_{i,t^*} = 1\right.\right) - \mathbb{E}\left(\eta_{i,t^*+1} \left| D_{i,t^*} = 0\right.\right)\right)\right] \\
&= \mathbb{E}\left(\eta_{i,t^*+1} \left| D_{i,t^*} = 1\right.\right) - \mathbb{E}\left(\eta_{i,t^*+1} \left| D_{i,t^*} = 0\right.\right) \\
&= \mathbb{E}\left(\tilde{y}_{i,t^*+k} \left| D_{i,t^*} = 1\right.\right) - \mathbb{E}\left(\tilde{y}_{i,t^*+k} \left| D_{i,t^*} = 0\right.\right)
\end{aligned}$$

The last line follows from Testable Prediction 4, establishing the remaining desired equivalence in periods $t^* + k$ for $k > 1$. □

Result Wald (Identification of β_ε with Wald Estimator)

If Assumptions IID and CE hold, the ratio of the difference in (normalized) log consumption between the treated and control group and the difference in (normalized) log income between the treated and control group in period t^* is equal to β_ε:

$$\hat{\beta}_{\varepsilon,\text{WALD}} \equiv \frac{\mathbb{E}\left(\tilde{c}_{i,t} \left| D_{i,t^*} = 1\right.\right) - \mathbb{E}\left(\tilde{c}_{i,t} \left| D_{i,t^*} = 0\right.\right)}{\mathbb{E}\left(\tilde{y}_{i,t} \left| D_{i,t^*} = 1\right.\right) - \mathbb{E}\left(\tilde{y}_{i,t} \left| D_{i,t^*} = 0\right.\right)} = \beta_\varepsilon$$

Proof. Note that normalized consumption in period t^* can be written as:

$$\begin{aligned}\tilde{c}_{it^*} &\equiv c_{it^*} - c_{i,t^*+r} \\ &= \sum_{s=t^*+r+1}^{t^*} \Delta c_{is} \\ &= \sum_{s=t^*+r+1}^{t^*} \left(\beta_\eta \eta_{is} + \beta_\varepsilon \varepsilon_{is} + \zeta_{is}\right)\end{aligned}$$

Taking expectations conditional on $D_{it^*} = 1$, we have:

$$\begin{aligned}\mathbb{E}\left(\tilde{c}_{it^*} \left| D_{it^*} = 1\right.\right) &= \mathbb{E}\left(\sum_{s=t^*+r+1}^{t^*} \left(\beta_\eta \eta_{is} + \beta_\varepsilon \varepsilon_{is} + \zeta_{is}\right) \left| D_{it^*} = 1\right.\right) \\ &= \mathbb{E}\left(\sum_{s=t^*+r+1}^{t^*-1} \left(\beta_\eta \eta_{is} + \beta_\varepsilon \varepsilon_{is} + \zeta_{is}\right) + \beta_\eta \eta_{it^*} + \zeta_{it^*}\right) + \mathbb{E}\left(\beta_\varepsilon \varepsilon_{it^*} \left| D_{it^*} = 1\right.\right)\end{aligned}$$

where the second line follows from the independence of the nearly all of the shocks contributing to $\tilde{c}_{it^*}$ and those involved in the definition of D_{i,t^*}, $-\Delta y_{i,t^*+1} = -\eta_{i,t^*+1} - \varepsilon_{i,t^*+1} + \varepsilon_{it^*}$. We can likewise show that:

$$\mathbb{E}\left(\tilde{c}_{it^*} \left| D_{it^*} = 0\right.\right) = \mathbb{E}\left(\sum_{s=t^*+r}^{t^*-1} \left(\beta_\eta \eta_{is} + \beta_\varepsilon \varepsilon_{is} + \zeta_{is}\right) + \beta_\eta \eta_{it^*} + \zeta_{it^*}\right) + \mathbb{E}\left(\beta_\varepsilon \varepsilon_{it^*} \left| D_{it^*} = 0\right.\right)$$

Taking the difference between the two yields the following expression for the numerator of the binary estimator:

$$\begin{aligned}\mathbb{E}\left(\tilde{c}_{it^*} \left| D_{it^*} = 1\right.\right) - \mathbb{E}\left(\tilde{c}_{it^*} \left| D_{it^*} = 0\right.\right) &= \mathbb{E}\left(\beta_\varepsilon \varepsilon_{it^*} \left| D_{it^*} = 1\right.\right) - \mathbb{E}\left(\beta_\varepsilon \varepsilon_{it^*} \left| D_{it^*} = 0\right.\right) \\ &= \beta_\varepsilon \left(\mathbb{E}\left(\varepsilon_{it^*} \left| D_{it^*} = 1\right.\right) - \mathbb{E}\left(\varepsilon_{it^*} \left| D_{it^*} = 0\right.\right)\right)\end{aligned}$$

We can likewise simplify the denominator of the binary estimator using Testable Prediction 1:

$$\mathbb{E}\left(\tilde{y}_{i,t^*} \left| D_{i,t^*} = 1\right.\right) - \mathbb{E}\left(\tilde{y}_{i,t^*} \left| D_{i,t^*} = 0\right.\right) = \mathbb{E}\left(\varepsilon_{i,t^*} \left| D_{i,t^*} = 1\right.\right) - \mathbb{E}\left(\varepsilon_{i,t^*} \left| D_{i,t^*} = 0\right.\right)$$

Taking the ratio of the numerator and denominator yields the desired result. □

Result ID (Identification of β_ε II)

If Assumptions IID2 and CE2 hold, β_ε^f is identified by using the one-period lead, coworker income *growth instrument:*

$$\hat{\beta}_{\varepsilon,HM-BPP-Co} \equiv \frac{\text{cov}\left(\Delta c_{it}, -\Delta y_{j(-i,t),t+1}\right)}{\text{cov}\left(\Delta y_{it}, -\Delta y_{j(-i,t),t+1}\right)} = \frac{\text{cov}\left(\Delta c_{it}, \varepsilon^f_{j(i,t),t}\right)}{\text{var}\left(\varepsilon^f_{j(i,t),t}\right)} = \frac{\beta^f_\varepsilon \sigma^2_{f\varepsilon}}{\sigma^2_{f\varepsilon}} = \beta^f_\varepsilon$$

where $\sigma^2_{f\varepsilon}$ is the variance of the transitory firm shock.

Proof. First, note that the coworker instrument can be simplified as follows:

$$\begin{aligned}
-\Delta y_{j(-i,t),t+1} &\equiv -\Delta \ln \left(\frac{1}{N_{j(-i,t),t+1}} \sum_{i'|j(i',t+1)=j(i,t),i'\neq i} Y_{i',t+1} \right) \\
&= -\Delta \ln \left(\frac{1}{N_{j(-i,t),t+1}} \sum_{i'|j(i',t+1)=j(i,t),i'\neq i} e^{y_{i',t+1}} \right) \\
&= -\ln \left(\frac{1}{N_{j(-i,t),t+1}} \sum_{i'|j(i',t+1)=j(i,t),i'\neq i} e^{y_{i',t+1}} \right) \\
&\quad + \ln \left(\frac{1}{N_{j(-i,t),t}} \sum_{i'|j(i',t)=j(i,t),i'\neq i} e^{y_{i',t}} \right) \\
&= -\ln \left(\frac{1}{N_{j(-i,t),t+1}} \sum_{i'|j(i',t+1)=j(i,t),i'\neq i} e^{\eta^f_{j(i',t+1),t+1}+\eta^f_{j(i',t+1),t}+\varepsilon^f_{j(i',t+1),t+1}} e^{z_{i',t-1}+\eta^w_{i',t}+\eta^w_{i',t+1}+\varepsilon^w_{i',t+1}} \right) \\
&\quad + \ln \left(\frac{1}{N_{j(-i,t),t}} \sum_{i'|j(i',t)=j(i,t),i'\neq i} e^{\eta^f_{j(i',t),t}+\varepsilon^f_{j(i',t),t}} e^{z_{i',t-1}+\eta^w_{i',t}+\varepsilon^w_{i',t}} \right) \\
&= -\ln \left(e^{\eta^f_{j(i,t),t+1}+\eta^f_{j(i,t),t}+\varepsilon^f_{j(i,t),t+1}} \frac{1}{N_{j(-i,t),t+1}} \sum_{i'|j(i',t+1)=j(i,t),i'\neq i} e^{z_{i',t-1}+\eta^w_{i',t}+\eta^w_{i',t+1}+\varepsilon^w_{i',t+1}} \right) \\
&\quad + \ln \left(e^{\eta^f_{j(i,t),t}+\varepsilon^f_{j(i,t),t}} \frac{1}{N_{j(-i,t),t}} \sum_{i'|j(i',t)=j(i,t),i'\neq i} e^{z_{i',t-1}+\eta^w_{i',t}+\varepsilon^w_{i',t}} \right) \\
&= -\eta^f_{j(i,t),t+1} - \eta^f_{j(i,t),t} - \varepsilon^f_{j(i,t),t+1} + \eta^f_{j(i,t),t} + \varepsilon^f_{j(i,t),t} \\
&\quad - \ln \left(\frac{1}{N_{j(-i,t),t+1}} \sum_{i'|j(i',t+1)=j(i,t),i'\neq i} e^{z_{i',t-1}+\eta^w_{i',t}+\eta^w_{i',t+1}+\varepsilon^w_{i',t+1}} \right) \\
&\quad + \ln \left(\frac{1}{N_{j(-i,t),t}} \sum_{i'|j(i',t)=j(i,t),i'\neq i} e^{z_{i',t-1}+\eta^w_{i',t}+\varepsilon^w_{i',t}} \right) \\
&= -\eta^f_{j(i,t),t+1} - \Delta\varepsilon^f_{j(i,t),t+1} + g_{j(-i,t+1),t+1}
\end{aligned}$$

where the function $g(\cdot)$ is shorthand for the final two terms and is a function of a combination of shocks that are independent of $(\Delta y_{i,t}, \Delta c_{i,t})$.

We can then show our result directly by simplifying the numerator and denominator:

$$
\begin{aligned}
\text{cov}\left(\Delta c_{it}, -\Delta y_{j(-i,t),t+1}\right) &= \text{cov}\left(\beta_\eta(\eta^w_{it} + \eta^f_{j(i,t),t}) + \beta_\varepsilon(\varepsilon^w_{it} + \varepsilon^f_{j(i,t),t}) + \zeta_{it},\right. \\
&\qquad \left. -\eta^f_{j(i,t),t+1} - \Delta\varepsilon^f_{j(i,t),t+1} + g_{j(-i,t),t+1}\right) \\
&= \text{cov}\left(\beta_\eta(\eta^w_{it} + \eta^f_{j(i,t),t}) + \beta_\varepsilon(\varepsilon^w_{it} + \varepsilon^f_{j(i,t),t}) + \zeta_{it},\right. \\
&\qquad \left. -\eta^f_{j(i,t),t+1} - \varepsilon^f_{j(i,t),t+1} + \varepsilon^f_{j(i,t),t} + g_{j(-i,t),t+1}\right) \\
&= \beta_\varepsilon \sigma^2_{f\varepsilon} \\
\text{cov}\left(\Delta y_{it}, -\Delta y_{j(-i,t),t+1}\right) &= \text{cov}\left(\eta^w_{i,t} + \eta^f_{j(i,t),t} + \Delta\varepsilon^w_{i,t} + \Delta\varepsilon^f_{j(i,t),t},\right. \\
&\qquad \left. -\eta^f_{j(i,t),t+1} - \Delta\varepsilon^f_{j(i,t),t+1} + g_{j(-i,t),t+1}\right) \\
&= \text{cov}\left(\eta^w_{i,t} + \eta^f_{j(i,t),t} + \varepsilon^w_{i,t} - \varepsilon^w_{i,t-1} + \varepsilon^f_{j(i,t),t} - \varepsilon^f_{j(i,t-1),t-1},\right. \\
&\qquad \left. -\eta^f_{j(i,t),t+1} - \varepsilon^f_{j(i,t),t+1} + \varepsilon^f_{j(i,t),t} + g_{j(-i,t),t+1}\right) \\
&= \sigma^2_{f\varepsilon}
\end{aligned}
$$

The result then follows by taking the ratio of the numerator and denominator. □

Appendix D Ordinary Least Squares

Consider the following first-difference OLS specification

$$\Delta c_{it} = \alpha + \beta_{OLS}\Delta y_{it} + \zeta_{it} \tag{D.1}$$

where i indexes households, t indexes time (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_{OLS} can be interpreted as an elasticity of consumption with respect to income, and thus a measure of consumption sensitivity. The parameter β_{OLS} is identified under the following assumption:

Assumption EO (Earnings Orthogonality)

$$\mathbb{E}(\zeta_{it}\,|\Delta y_{it}) = 0$$

Assumption EO asserts that income changes are exogenous with respect to the unobserved determinants of consumption growth.

There are two key challenges to this baseline approach. First, variation in labor income includes unexpected permanent and transitory shocks and mean reversion following transitory shocks, as well as predictable fluctuations, due to, for example, seasonality in earnings. In that case, the parameter β_{OLS} is not readily interpretable as the consumption response to unexpected shocks that is commonly the target of consumption-savings models. To see this, note that we have the following expression that summarizes the different components of this estimator:

Result OLS (OLS Estimator)

If income follows the process in equation (3) and consumption changes are governed by equation

(4), then the OLS regression of consumption changes on income changes is as follows:

$$\hat{\beta}_{\varepsilon,OLS} \equiv \frac{\text{cov}(\Delta c_{it}, \Delta y_{it})}{\text{var}(\Delta y_{it})} = \frac{\sigma_\eta^2}{\sigma_\eta^2 + 2\sigma_\varepsilon^2}\beta_\eta + \frac{\sigma_\varepsilon^2}{\sigma_\eta^2 + 2\sigma_\varepsilon^2}\beta_\varepsilon + \frac{\text{cov}(\zeta_{it}, \Delta y_{it})}{\text{var}(\Delta y_{it})}$$

Proof. This can be directly shown by writing out the full expressions for the consumption and income changes:

$$\begin{aligned}
\text{cov}(\Delta c_{it}, \Delta y_{it}) &= \text{cov}(\beta_\eta \eta_{it} + \beta_\varepsilon \varepsilon + \zeta_{it}, \Delta y_{it}) \\
&= \text{cov}(\beta_\eta \eta_{it} + \beta_\varepsilon \varepsilon, \Delta y_{it}) + \text{cov}(\zeta_{it}, \Delta y_{it}) \\
&= \text{cov}(\beta_\eta \eta_{it} + \beta_\varepsilon \varepsilon, \eta_{it} + \varepsilon_{it} - \varepsilon_{i,t-1}) + \text{cov}(\zeta_{it}, \Delta y_{it}) \\
&= \sigma_\eta^2 \beta_\eta + \sigma_\varepsilon^2 \beta_\varepsilon + \text{cov}(\zeta_{it}, \Delta y_{it}) \\
\text{var}(\Delta y_{it}) &= \sigma_\eta^2 + 2\sigma_\varepsilon^2
\end{aligned}$$

Taking the ratio of the two yields the result. □

If we invoke Assumption EO, we still have that the estimator is a mix of transitory and permanent shocks:

Corollary MIX (Mixture of Responses)

If income follows the process in equation (3) and consumption changes are governed by equation (4), and EO holds, then the OLS estimator is a mixture of permanent and transitory responses:

$$\hat{\beta}_{\varepsilon,OLS} \equiv \frac{\text{cov}(\Delta c_{it}, \Delta y_{it})}{\text{var}(\Delta y_{it})} = \frac{\sigma_\eta^2}{\sigma_\eta^2 + 2\sigma_\varepsilon^2}\beta_\eta + \frac{\sigma_\varepsilon^2}{\sigma_\eta^2 + 2\sigma_\varepsilon^2}\beta_\varepsilon$$

Proof. Assumption EO implies that $\text{cov}(\zeta_{it}, \Delta y_{it}) = 0$. Plugging into Result OLS, the result immediately follows. □

Note, however, that the weights in this mixture do not sum to one. This is because part of the variation in Δy_{it} is due to mean reversion of a transitory shock in period $t-1$. This variation in income that does not affect consumption leads to attenuation. We can see this more clearly if we assume away permanent shocks. Even in that case, the estimator is biased downward due to mean reversion:

Corollary MRA (Mean Reversion Attenuation)

If income follows the process in equation (3) and consumption changes are governed by equation (4), and EO holds, and there are no permanent shocks ($\sigma_\eta^2 = 0$), then the OLS estimator will be an attenuated estimate of β_ε:

$$\hat{\beta}_{\varepsilon,OLS} \equiv \frac{\text{cov}(\Delta c_{it}, \Delta y_{it})}{\text{var}(\Delta y_{it})} = \frac{\sigma_\varepsilon^2}{2\sigma_\varepsilon^2}\beta_\varepsilon = \frac{1}{2}\beta_\varepsilon$$

Proof. Assumption EO implies that $\text{cov}(\zeta_{it}, \Delta y_{it}) = 0$. Using Result OLS, and plugging in $\sigma_\eta^2 = 0$, the result immediately follows. □

Intuitively, the variation in income growth is driven by both the innovation and predictable reversion of the prior shock, while the only variation that affects consumption is the unpredictable innovation half of the change.

Appendix E Testable Predictions for a Continuous Instrument

Assumptions IID and CE support identification of our key parameter of interest, and also yield an additional set of empirical implications regarding earnings that may be inspected with panel data. In the main text, we present a set of testable predictions for normalized log income conditional on a binary version of our primary instrument. In this appendix, we present analogous predictions based on the continuous version of the instrument: $-\Delta y_{i,t^*+1}$.

Testable Prediction E1 (First Stage: Positive Income Covariance at t^*)

If Assumption IID holds, the covariance between the (negative) one-year income growth lead and (normalized) log income in period t^* is positive and equal to the variance of the transitory shock:

$$\text{cov}\left(\tilde{y}_{i,t^*}, -\Delta y_{i,t^*+1}\right) = \sigma_\varepsilon^2 > 0$$

Proof. Note that we can write normalized log income in period t^* as:

$$\begin{aligned}
\tilde{y}_{it^*} &\equiv y_{it^*} - y_{i,t^*+r} \\
&= z_{it^*} + \varepsilon_{it^*} - z_{i,t^*+r} - \varepsilon_{i,t^*+r} \\
&= \sum_{s=t^*+r+1}^{t^*} \eta_{i,s} + \varepsilon_{it^*} - \varepsilon_{i,t^*+r}
\end{aligned}$$

We then have:

$$\begin{aligned}
\text{cov}\left(\tilde{y}_{i,t^*}, -\Delta y_{i,t^*+1}\right) &= \text{cov}\left(\sum_{s=t^*+r+1}^{t^*} \eta_{i,s} + \varepsilon_{it^*} - \varepsilon_{i,t^*+r}, -\eta_{i,t^*+1} - \varepsilon_{i,t^*+1} + \varepsilon_{i,t^*}\right) \\
&= \sigma_\varepsilon^2 \\
&> 0
\end{aligned}$$

□

Note that the covariance in this result is exactly the numerator of our first-stage regression coefficient in income when we set $r = -1$, and continues to equal the numerator of our first stage coefficient for any $r < 0$.

Testable Prediction E2 (Parallel Pre-event Income Trends)

If Assumption IID holds, the covariance between the (negative) one-year income growth lead and lagged (normalized) log income is zero during periods before t^*:

$$\text{cov}\left(\tilde{y}_{i,t^*+k}, -\Delta y_{i,t^*+1}\right) = 0 \quad \text{for} \quad k < 0$$

Proof. Note that normalized income in period $t^* + k$, $k < 0$ can be written as:

$$\begin{aligned}\tilde{y}_{i,t^*+k} &\equiv y_{i,t^*+k} - y_{i,t^*+r} \\ &= z_{i,t^*+k} + \varepsilon_{i,t^*+k} - z_{i,t^*+r} - \varepsilon_{i,t^*+r} \\ &= (-1)^{\mathbf{1}(r>k)} \sum_{s=t^*+\min(k,r)+1}^{t^*+\max(k,r)} \eta_{i,s} + \varepsilon_{i,t^*+k} - \varepsilon_{i,t^*+r}\end{aligned}$$

We then have for $k < 0$:

$$\begin{aligned}\text{cov}\left(\tilde{y}_{i,t^*+k}, -\Delta y_{i,t^*+1}\right) &= \text{cov}\left((-1)^{\mathbf{1}(r>k)} \sum_{s=t^*+\min(k,r)+1}^{t^*+\max(k,r)} \eta_{i,s} + \varepsilon_{i,t^*+k} - \varepsilon_{i,t^*+r}, -\eta_{i,t^*+1} - \varepsilon_{i,t^*+1} + \varepsilon_{i,t^*}\right) \\ &= 0\end{aligned}$$

where the last line follows from Assumption IID and the fact that the shocks contributing to normalized log income do not overlap with the shocks contributing to the instrument. □

As in the case of a difference-in-differences (DD) estimator or event study, our assumptions imply trends that can be visually inspected using panel data. In this way, the semi-structural approach of HM-BPP can be linked to tests that are common in the quasi-experimental literature. Note, as well, that this prediction of parallel trends is related to a result in Blundell, Pistaferri, and Preston (2008), who show that the covariance between lagged first-differences and the same instrument are zero. Those predictions are nested within ours for an appropriately chosen $r < 0$. These predictions are also in the spirit of a prediction by Hall and Mishkin (1982), whose model implies zero correlation between consumption growth and lagged income growth.

We also have predictions for post-event periods:

Testable Prediction E3 (Negative Income Covariance at $t^* + 1$)

If Assumption IID holds, the covariance between the (negative) one-year income growth lead and (normalized) log income in period $t^* + 1$ is negative and equal in absolute magnitude to the sum of the variances of the transitory and permanent shocks:

$$\text{cov}\left(\tilde{y}_{i,t^*+1}, -\Delta y_{i,t^*+1}\right) = -\sigma_\eta^2 - \sigma_\varepsilon^2 < 0$$

Proof. Note that we can write normalized log income in period $t^* + 1$ as:

$$\begin{aligned}\tilde{y}_{it^*+1} &\equiv y_{i,t^*+1} - y_{i,t^*+r} \\ &= z_{i,t^*+1} + \varepsilon_{i,t^*+1} - z_{i,t^*+r} - \varepsilon_{i,t^*+r} \\ &= \sum_{s=t^*+r+1}^{t^*+1} \eta_{i,s} + \varepsilon_{i,t^*+1} - \varepsilon_{i,t^*+r}\end{aligned}$$

We then have:

$$\begin{aligned}\text{cov}\left(\tilde{y}_{i,t^*+1}, -\Delta y_{i,t^*+1}\right) &= \text{cov}\left(\sum_{s=t^*+r+1}^{t^*+1} \eta_{i,s} + \varepsilon_{i,t^*+1} - \varepsilon_{i,t^*+r}, -\eta_{i,t^*+1} - \varepsilon_{i,t^*+1} + \varepsilon_{i,t^*}\right) \\ &= -\sigma_\eta^2 - \sigma_\varepsilon^2 \\ &< 0\end{aligned}$$

□

Testable Prediction E4 (Negative Income Covariance at $t^* + 2$ Onward)

If Assumption IID holds, the covariance between the (negative) one-year income growth lead and (normalized) log income in periods $t > t^* + 1$ is negative and equal in absolute magnitude to the variance of the permanent shock:

$$\text{cov}\left(\tilde{y}_{i,t^*+k}, -\Delta y_{i,t^*+1}\right) = -\sigma_\eta^2 < 0 \quad \text{for} \quad k > 1$$

Proof. Note that we can write normalized log income in period $t^* + k, k > 1$ as:

$$\begin{aligned}\tilde{y}_{it^*+k} &\equiv y_{i,t^*+k} - y_{i,t^*+r} \\ &= z_{i,t^*+k} + \varepsilon_{i,t^*+k} - z_{i,t^*+r} - \varepsilon_{i,t^*+r} \\ &= \sum_{s=t^*+r+1}^{t^*+k} \eta_{i,s} + \varepsilon_{i,t^*+k} - \varepsilon_{i,t^*+r}\end{aligned}$$

We then have for $k > 1$:

$$\begin{aligned}\text{cov}\left(\tilde{y}_{i,t^*+k}, -\Delta y_{i,t^*+1}\right) &= \text{cov}\left(\sum_{s=t^*+r+1}^{t^*+k} \eta_{i,s} + \varepsilon_{i,t^*+k} - \varepsilon_{i,t^*+r}, -\eta_{i,t^*+1} - \varepsilon_{i,t^*+1} + \varepsilon_{i,t^*}\right) \\ &= -\sigma_\eta^2 \\ &< 0\end{aligned}$$

□

These predictions mean that when using a continuous instrument, if we plot the results of an event-study like regression, we end up with similar qualitative patterns as in the case of the binary instrument: parallel trends in the pre-period, a positive spike at time t^*, a negative spike at time $t^* + 1$, and smaller, negative differences thereafter. Note that in our result log income is normalized relative to a pre-period level of income. If we had instead used a simple first-difference in income, we would no longer recover information relevant to the permanent income shock, as the covariance would now equal zero. That is:

$$\text{cov}\left(\Delta y_{i,t^*+k}, -\Delta y_{i,t^*+1}\right) = \text{cov}\left(\eta_{i,t^*+k} + \varepsilon_{i,t^*+k} - \varepsilon_{i,t^*+k-1}, -\eta_{i,t^*+1} - \varepsilon_{i,t^*+1} + \varepsilon_{i,t^*}\right) = 0$$

It is therefore important to normalize relative to a pre-period. We similarly have a prediction regarding pre-trends in consumption:

Testable Prediction E5 (Parallel Pre-event Consumption Trends)

If Assumption IID and CE hold, the covariance between the (negative) one-year income growth lead and lagged (normalized) log consumption is zero during periods before t^*:

$$\text{cov}\left(\tilde{c}_{i,t^*+k}, -\Delta y_{i,t^*+1}\right) = 0 \quad \text{for} \quad k < 0$$

Proof. Note that normalized consumption in period $t^* + k$, $k < 0$ can be written as:

$$\begin{aligned}
\tilde{c}_{i,t^*+k} &= c_{i,t^*+k} - c_{i,t^*+r} \\
&= (-1)^{\mathbf{1}(r>k)} \sum_{s=t^*+\min(k,r)+1}^{t^*+\max(k,r)} \Delta c_{is} \\
&= (-1)^{\mathbf{1}(r>k)} \sum_{s=t^*+\min(k,r)+1}^{t^*+\max(k,r)} (\beta_\eta \eta_{is} + \beta_\varepsilon \varepsilon_{is} + \zeta_{is})
\end{aligned}$$

We then have for $k < 0$:

$$\begin{aligned}
\text{cov}\left(\tilde{c}_{i,t^*+k}, -\Delta y_{i,t^*+1}\right) &= \text{cov}\left((-1)^{\mathbf{1}(r>k)} \sum_{s=t^*+\min(k,r)+1}^{t^*+\max(k,r)} (\beta_\eta \eta_{is} + \beta_\varepsilon \varepsilon_{is} + \zeta_{is}), -\eta_{i,t^*+1} - \varepsilon_{i,t^*+1} + \varepsilon_{i,t^*}\right) \\
&= 0
\end{aligned}$$

where the last line follows from Assumptions IID and CE and the fact that the shocks contributing to normalized log consumption do not overlap with the shocks contributing to the instrument. □

All of the above covariances can be tested within a single event-study type regression, as in equation (17). The coefficents in that regression are proportional to the covariances above, scaled by a factor of $1/\text{var}(-\Delta y_{i,t^*+1}) = 1/(\sigma_\eta^2 + 2\sigma_\varepsilon^2)$.

Appendix F Sample Restrictions

We impose sample screens to ensure that we focus on employed households whose primary bank accounts are at Chase and for whom we reliably measure labor income. For income, we restrict to households who receive a checking account direct deposit that JPMCI categorizes as a paycheck. Each paycheck includes a field describing the counterparty who initiated the direct deposit. To ensure that these direct deposits are actually labor income and that we are reliably measuring them, we exclude all paychecks coming from jobs that appear to be paid on an irregular pay schedule. This drops 8% of paychecks. After this restriction, we aggregate all remaining paychecks paid in each month by each employer to calculate employer-level variables. We then take a 10% sample of employed households for computational feasibility.

In order to run our primary empirical specifications, we require non-missing household-level labor income, coworker labor income, and nondurable spending in a given month as well as the prior month, and a non-missing checking account buffer in the current month. We also require non-missing twelfth lags of household-level labor income and coworker labor income in order to run our seasonal adjustment procedure. Together, these requirements drop 45% of households and 55% of household-months.

Finally, we impose an activity screen to ensure that we focus on households whose primary bank account is at Chase. We drop all household-months without five checking account outflows in that month and in the previous month. This sample restriction drops 9% of households and 15% of household months. Our final analysis sample includes 1.3 million households, 1.7 million household-employers, 25.2 million household-months, and 27.9 million household-employer-months.

Appendix G Sources of Variation

We use PayrollCompany data to examine the underlying sources of earnings variation observed in JPMCI bank account data. PayrollCompany data is described in more detail in Section 3 and in Ganong et al. (2025).

The unit of analysis in PayrollCompany data is a "pay item," which captures a distinct component of a worker's paycheck-specific earnings. The three most common pay items are base pay, bonuses, and overtime (Ganong et al. 2025). Typically, a single paycheck includes multiple pay items, the sum of which are deposited into the employee's bank account at the end of a pay period. Earnings data structure differs significantly between hourly and salaried workers. For hourly workers, data explicitly include hours worked and hourly wage rates, whereas salaried workers have only a monthly wage rate. Given these differences, we conduct separate analyses for hourly and salaried employees.

Ganong et al. (2025) find that earnings instability among salaried workers is not primarily driven by fluctuations in their wage rates. Similarly, for hourly workers, changes in earnings per hour account for only a small portion of total earnings variability. Instead, fluctuations in hours worked constitute a substantial share of the overall earnings variation observed for hourly workers. We show below that the same patterns are apparent when subsetting just to the volatility identified by our instrument.

Figure G-1 documents the underlying sources of earnings variation picked up by our instrument by separating earnings changes into their respective components: wages and hours. The top panel focuses on salaried workers, plotting the normalized mean change in total earnings and the monthly wage rate across deciles of our instrument, defined as the seasonally adjusted change in period-ahead coworker earnings per check. The bottom panel repeats this analysis for hourly workers, substituting the monthly wage rate with the hourly wage rate and including the normalized mean change in hours worked. These results lead us to conclude that wage rate fluctuations account for a minimal share of the earnings variability identified by our instrument, whereas changes in hours worked explain much of the earnings instability observed among hourly workers and changes in non-base salary compensation account for most of the variation for salaried workers.

Figure G-1: Sources of Variation

(a) Salaried Workers

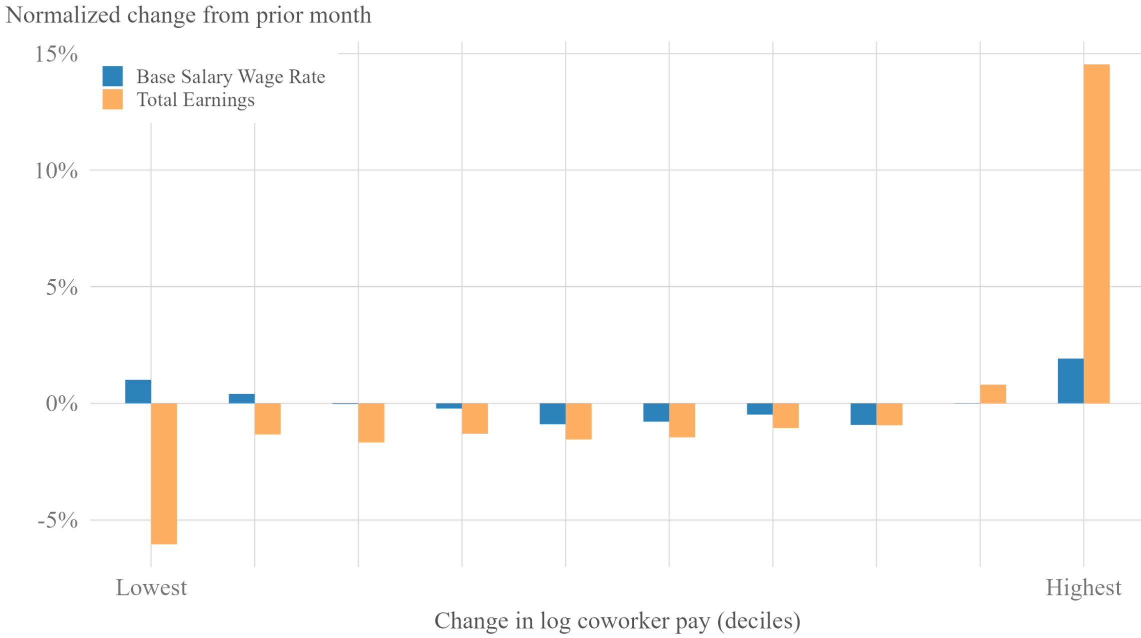

(b) Hourly Workers

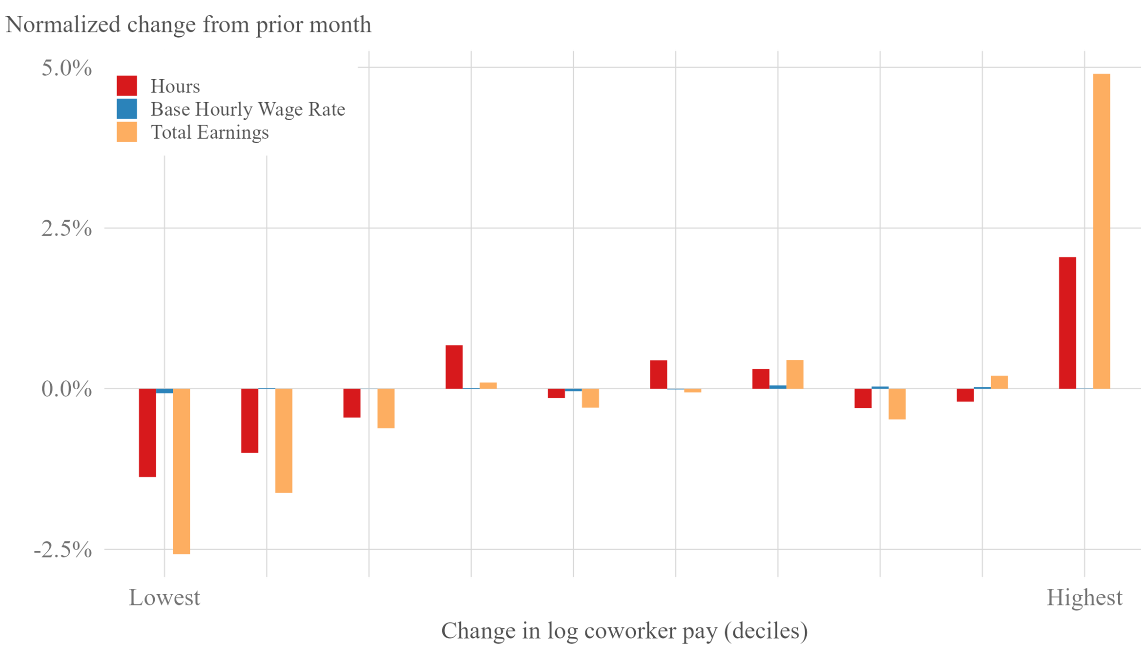

Notes: This figure separates changes in earnings into their underlying components using administrative payroll data. The top panel restricts the sample to salaried workers and plots the mean change in total earnings and base wages across 10 deciles of the instrument $(-\Delta y_{j(-i,t),t+1})$ conditional on the instrument value within each decile and net of the unconditional mean change across all deciles. The bottom panel repeats this exercise for hourly workers, plotting mean changes in total earnings, hours worked, and earnings per hour across the same deciles. See Appendix G for details.

Appendix H Simulating an Income Process

In Figure 3 we plot coefficients from regressions using actual data and simulated data. For the simulated data, we generate an income process for pay per paycheck, following Section 2, with additional components as explained in Appendix B.2:

$$\begin{aligned} y^p_{i,t} &= z_{i,t} + \varepsilon^w_{i,t} + \varepsilon^f_{j(i,t),t} \\ z_{i,t} &= z_{i,t-1} + \eta^w_{i,t} + \eta^f_{j(i,t),t} \quad \text{for} \quad t \geq 1 \end{aligned} \tag{H.1}$$

Where i indexes the N workers, j indexes the J firms, and t ranges from zero to T. We draw an initial, iid income component for each worker from the $z_{i,0} \sim \mathcal{N}(0, \sigma^2_{z_0})$. We then draw iid temporary worker shocks for each worker starting in period 0, $\varepsilon^w_{i,t} \sim \mathcal{N}(0, \sigma^2_{\varepsilon w})$, and iid permanent shocks for each worker staring in period 1, $\eta^w_{i,t} \sim \mathcal{N}(0, \sigma^2_{\eta w})$. For firm shocks, we draw iid transitory firm shocks starting in period 0, $\varepsilon^f_{j,t} \sim \mathcal{N}(0, \sigma^2_{\varepsilon f})$, and iid permanent firm shocks $\eta^f_{j,t} \sim \mathcal{N}(0, \sigma^2_{\eta f})$ starting in period 1. Each worker is randomly assigned, with equal probability, to one firm for the entire series. Finally, for each worker, in each period, we randomly draw a number of paychecks received, $M_{i,t}$, using a discrete distribution over the mass points 1 through 9. Letting $m_{i,t} \equiv \log M_{i,t}$, total pay in a period is then:

$$y_{i,t} \equiv y^p_{i,t} + m_{i,t} \tag{H.2}$$

We next simulate, for each firm, a coworker instrument that is used for each worker at that firm in our simulated data. Rather than average over the workers in our simulated data set, we generate the instrument directly as follows:

$$\Delta y^p_{j(-i,t),t+1} = \eta^f_{j(i,t),t+1} + \varepsilon^f_{j(i,t),t+1} - \varepsilon^f_{j(i,t),t} + \vartheta_{j(-i,t),t+1} \tag{H.3}$$

where $\vartheta_{j(-i,t),t+1}$ is a measurement error term with distribution $\mathcal{N}(0, \sigma^2_\vartheta)$.

In Figure 3, there are separate simulations for our regressions using binary and continuous instruments. To match the levels of the graphs, we use separate values for the variance of the permanent firm shock, $\eta^f_{j,t}$, and the coworker instrument measurement error, $\vartheta_{j(-i,t),t+1}$, for the binary and continuous versions of the graphs. The values used for the parameters underlying our simulations are provided below in Table H-1:

Table H-1: Income Simulation Parameter Values

Parameter	Value
Number of Simulated Workers: N	150,000
Number of Simulated Periods: T	120
Number of Simulated Firms: J	500
Variance of Worker Transitory Shocks: $\sigma^2_{\varepsilon w}$	0.00752
Variance of Worker Permanent Shocks: $\sigma^2_{\eta w}$	0.00557
Variance of Firm Transitory Shocks: $\sigma^2_{\varepsilon f}$	0.00255
Variance of Firm Permanent Shocks:	
Binary Instrument	0.000810
Continuous Instrument	0.000210
Variance of Instrument Measurement Error: σ^2_{ϑ}	
Binary Instrument	0.0181
Continuous Instrument	0.00885
Distribution of Number of Paychecks:	
$\Pr(M_{i,t} = 1)$	9.43%
$\Pr(M_{i,t} = 2)$	61.9%
$\Pr(M_{i,t} = 3)$	12.6%
$\Pr(M_{i,t} = 4)$	10.4%
$\Pr(M_{i,t} = 5)$	5.50%
$\Pr(M_{i,t} = 6)$	0.164%
$\Pr(M_{i,t} = 7)$	0.0146%
$\Pr(M_{i,t} = 8)$	0.00441%
$\Pr(M_{i,t} = 9)$	0.00205%

Notes: This table reports the inputs for the simulations used in Figure 3.

Once the data are simulated, we have a series $\left\{\Delta y_{i,t}, \Delta y^p_{i,t}\right\}_{t=0,\dots,T}$ for N workers, a series $\left\{\Delta y^p_{j,t}\right\}_{t=0,\dots,T}$ for J firms, and a mapping $j(i,t) : i,t \Rightarrow j$ assigning each worker to a firm in each period. Since workers do not change firms in our simulation, this assignment happens once and remains throughout. We can then define our two key instruments $D^{\text{Binary}}_{i,t} \equiv \mathbf{1}\left(-\Delta y^p_{j(-i,t),t+1} > M_{-\Delta y}\right)$ and $D^{\text{Continuous}}_{i,t} \equiv -\Delta y^p_{j(-i,t),t+1}$, where $M_{-\Delta y} \equiv median\left(-\Delta y^p_{j(-i,t),t+1}\right)$. These variables are then used to estimate our event study-like regression in equation (17), reproduced here:

$$y_{i,t^*+s} = \sum_{k=-12,k\neq-4}^{12} \rho_k \times \mathbf{1}\{k = t^* + s\} \times D_{i,t^*} + \phi_s + \nu_{i,t^*+s}$$

We stack the data, treating each period as $t = t^*$, and plot the coefficients in Figure 3. The key goal of this exercise is to demonstrate that in simulated data, an income process with the features of equations (H.1) and (H.2), used with an instrument as in equation (H.3), will yield the specific sawtooth pattern we find in the data. The fact that we can match the levels of the plot is less important, as we have enough degrees of freedom in this simulation to do so.

Appendix I MPC Scaling Factor to Match NIPA Benchmarks

In Section 4.3, we report MPC estimates by scaling our elasticity estimates by the ratio of mean nondurable consumption to mean labor income in our sample. However, it is possible that we have an incomplete lens on total nondurable consumption and total labor income. We could be missing some nondurable consumption because we do not include any spending on paper checks, because we categorize only a subset of electronic ACH payments, and because, despite our activity screens, households may still do some spending via non-Chase credit cards or out of non-Chase bank accounts. Similarly, we could be missing some labor income because we do not include labor income paid via cash or paper check, and because households may receive labor income into non-Chase bank accounts. To account for this potential unobserved spending and income, we develop a scaling factor similar to the procedure in Broda and Parker (2014), who scale up from the subset of consumer spending observed in the Nielsen Consumer Panel to a measure that aligns with National Income and Product Account (NIPA) spending per capita.

To compute an MPC scaling factor in our application, we need two components: mean NIPA household spending on nondurables, and mean NIPA net income for employed households. To compute mean nondurable spending, we first classify the PCE spending categories from the 2019 NIPA table 1.10 as either durable or nondurable.[1] Since the PCE captures total medical and drug spending, but we want only out of pocket spending for our MPC, we account for this as in Ganong and Noel (2019). In particular, we multiply PCE medical spending by 12.4 percent, and PCE spending on drugs and pharmaceutical products by 14 percent, to obtain a measure of out of pocket medical spending. After these adjustments, we sum spending on the PCE categories we categorize as nondurable, and divide the total by the number of households in 2019, which was about 123 million according to the ACS.[2] Using this procedure, we estimate that, on average, households spent $3,995 per month on nondurables in 2019.

We perform a similar exercise for income, the denominator of our scaling factor. We start from total gross wages paid to persons in NIPA table 1.10, which gives us roughly $9.3 trillion. To align this more closely with the type of data we observe from the Chase bank accounts, we net out federal and state tax withholding. Using the 2019 IRS data book[3], we estimate that roughly $1.3 trillion was withheld in federal taxes. Using the FRED series on National Totals of State Tax Revenue[4], we then estimate total state tax revenue to be about $0.4 trillion in 2019. We further net out employee contributions to pension, healthcare and social security. From the 2019 Private Pension Plan bulletin of the DOL[5], we find that employees contributed a total of $0.3 trillion to private pension plans. Using a 2019 Policy Research Perspectives report on National Health Expenditures by the AMA[6], we find that employees contributed about $0.3 trillion to their health insurance plans. Finally, using the SSA's Contributions to the Social Security and Medicare Trust Funds tables[7], we find that employee social security contributions totaled roughly $0.6 trillion. That leaves total net income to employed

[1] We start from the classification in Ganong and Noel (2019) and update it for the 2019 list of NIPA categories. The crosswalk can be found here: https://docs.google.com/spreadsheets/d/1RjYtJZVy8DPrfMJfxv3hT4lPzfKvdEeGYNPS20LpySo

[2] https://data.census.gov/table/ACSST1Y2019.S1101?q=ACSST1Y2019.S1101

[3] https://www.irs.gov/pub/irs-prior/p55b–2020.pdf

[4] https://fred.stlouisfed.org/series/QTAXTOTALQTAXCAT1USYES

[5] https://www.dol.gov/sites/dolgov/files/ebsa/researchers/statistics/retirement-bulletins/private-pension-plan-bulletins-abstract-2019.pdf

[6] https://www.ama-assn.org/system/files/2021-05/prp-annual-spending-2019.pdf

[7] https://www.ssa.gov/oact/STATS/table3c3.html

households at about \$6 trillion. To get at an average number, we must estimate the total number of employed households in the U.S. in 2019. Using the 2019 ACS table on household size by number of workers in the household[8], we find that about 91 million households had at least one employed member in 2019. Putting it all together, we estimate that employed households earned \$5,781 per month in net labor income, on average, in 2019.

Combining these two inputs, the ratio of average nondurable consumption to average net labor income from 2019 NIPA data is 0.69, whereas the ratio from the JPMCI data is 0.45, as shown in Table A-10 column (5). Therefore, in order to rescale our MPC estimates to the NIPA benchmarks, as in Broda and Parker (2014), readers should multiply them by 1.5, the ratio of the PCE scaling factor to the JPMCI scaling factor.

Appendix J Estimating a Dynamic Consumption Response

We now extend the consumption model from Section 2.1 to include a dynamic response to transitory shocks. We allow transitory income shocks in a given month to impact consumption for three months, i.e. one quarter. We maintain the income process from equation (3), but amend the consumption function in equation (4) to include potential responses to lagged transitory shocks:

$$\begin{aligned} \Delta c_{it} &= \beta_\eta \eta_{it} + \beta_{\varepsilon,0}\varepsilon_{it} + \beta_{\varepsilon,1}\varepsilon_{i,t-1} + \beta_{\varepsilon,2}\varepsilon_{i,t-2} + \beta_{\varepsilon,3}\varepsilon_{i,t-3} + \zeta_{it} \\ \beta_{\varepsilon,3} &= -\left(\beta_{\varepsilon,0} + \beta_{\varepsilon,1} + \beta_{\varepsilon,2}\right) \end{aligned} \tag{J.1}$$

The constraint on $\beta_{\varepsilon,3}$ limits the impact of a transitory shock to three periods. If equation (J.1) holds, we can identity the parameters that govern the dynamic response relying on the same set of assumptions underlying Result HM-BPP:[9,10]

Result Dynamic (Identification of $\beta_{\varepsilon,0}$, $\beta_{\varepsilon,1}$, and $\beta_{\varepsilon,2}$)

If the consumption function follows equation (J.1) and Assumptions IID and CE hold, the dynamic response parameters can be identified as follows:

[8] https://data.census.gov/table/ACSDT1Y2019.B08202?t=Employment+and+Labor+Force+Status&y=2019&d=ACS+1-Year+Estimates+Detailed+Tables

[9] In appendix C "more general model," Blundell, Pistaferri, and Preston (2008) consider a more general consumption function with dynamic responses to both transitory and permanent shocks. They show that without further assumptions, only the contemporaneous response to a transitory shock is identified, while all responses to lagged shocks are not. We add the assumption that consumption fully adjusts to permanent shocks in the month that a permanent shock is realized, which enables us to identify the dynamic response to transitory shocks. In practice this assumption has little quantitative importance in our setting because we find that there are few permanent shocks at a monthly frequency.

[10] This dynamic version of the model is over-identified, given our constraint on $\beta_{\varepsilon,3}$. The moments we use are just one of many possible combinations of moments. We have confirmed that using alternative combinations of moments yield similar qualitative results. Namely, the moment used to estimate $\beta_{\varepsilon,0}$ holds in all cases; MPCs decrease over the three months; and the cumulative MPC over three months remains strictly positive.

$$
\begin{aligned}
\hat{\beta}_{\varepsilon,0} &\equiv \frac{\operatorname{cov}\left(\Delta c_{it}, -\Delta y_{i,t+1}\right)}{\operatorname{cov}\left(\Delta y_{it}, -\Delta y_{i,t+1}\right)} \\
&= \beta_{\varepsilon,0} \\
\hat{\beta}_{\varepsilon,1} &\equiv \frac{\operatorname{cov}\left(\Delta c_{it}, \Delta y_{i,t-1} + \Delta y_{i,t-2} + \Delta y_{i,t-3}\right)}{\operatorname{cov}\left(\Delta y_{it}, -\Delta y_{i,t+1}\right)} \\
&= \beta_{\varepsilon,1} \\
\hat{\beta}_{\varepsilon,2} &\equiv \frac{\operatorname{cov}\left(\Delta c_{it}, \Delta y_{i,t-2} + \Delta y_{i,t-3}\right)}{\operatorname{cov}\left(\Delta y_{it}, -\Delta y_{i,t+1}\right)} \\
&= \beta_{\varepsilon,2}
\end{aligned}
$$

Proof. First note that in all three cases, the denominator simplifies as follows:

$$
\begin{aligned}
\operatorname{cov}\left(\Delta y_{it}, -\Delta y_{i,t+1}\right) &= \operatorname{cov}\left(\eta_{it} + \varepsilon_{it} - \varepsilon_{i,t-1}, -\eta_{i,t+1} - \varepsilon_{i,t+1} + \varepsilon_{it}\right) \\
&= \sigma_{\varepsilon}^2
\end{aligned}
$$

Next, we can simplify each of the three numerators:

$$
\begin{aligned}
\operatorname{cov}\left(\Delta c_{it}, -\Delta y_{i,t+1}\right) &= \operatorname{cov}\left(\beta_\eta \eta_{it} + \beta_{\varepsilon,0}\varepsilon_{it} + \beta_{\varepsilon,1}\varepsilon_{i,t-1} + \beta_{\varepsilon,2}\varepsilon_{i,t-2} + \beta_{\varepsilon,3}\varepsilon_{i,t-3} + \zeta_{it},\right. \\
&\qquad\qquad \left. -\eta_{i,t+1} - \varepsilon_{i,t+1} + \varepsilon_{it}\right) \\
&= \beta_{\varepsilon,0}\sigma_{\varepsilon}^2 \\
\operatorname{cov}\left(\Delta c_{it}, \Delta y_{i,t-1} + \Delta y_{i,t-2} + \Delta y_{i,t-3}\right) &= \operatorname{cov}\left(\Delta c_{it}, y_{i,t-1} - y_{i,t-2} + y_{i,t-2} - y_{i,t-3} + y_{i,t-3} - y_{i,t-4}\right) \\
&= \operatorname{cov}\left(\Delta c_{it}, y_{i,t-1} - y_{i,t-4}\right) \\
&= \operatorname{cov}\left(\Delta c_{it}, z_{i,t-1} + \varepsilon_{i,t-1} - z_{i,t-4} - \varepsilon_{i,t-4}\right) \\
&= \operatorname{cov}\left(\beta_\eta \eta_{it} + \beta_{\varepsilon,0}\varepsilon_{it} + \beta_{\varepsilon,1}\varepsilon_{i,t-1} + \beta_{\varepsilon,2}\varepsilon_{i,t-2} + \beta_{\varepsilon,3}\varepsilon_{i,t-3} + \zeta_{it},\right. \\
&\qquad\qquad \left. \eta_{i,t-1} + \eta_{i,t-2} + \eta_{i,t-3} + \varepsilon_{i,t-1} - \varepsilon_{i,t-4}\right) \\
&= \beta_{\varepsilon,1}\sigma_{\varepsilon}^2 \\
\operatorname{cov}\left(\Delta c_{it}, \Delta y_{i,t-2} + \Delta y_{i,t-3}\right) &= \operatorname{cov}\left(\Delta c_{it}, y_{i,t-2} - y_{i,t-3} + y_{i,t-3} - y_{i,t-4}\right) \\
&= \operatorname{cov}\left(\Delta c_{it}, y_{i,t-2} - y_{i,t-4}\right) \\
&= \operatorname{cov}\left(\Delta c_{it}, z_{i,t-2} + \varepsilon_{i,t-2} - z_{i,t-4} - \varepsilon_{i,t-4}\right) \\
&= \operatorname{cov}\left(\beta_\eta \eta_{it} + \beta_{\varepsilon,0}\varepsilon_{it} + \beta_{\varepsilon,1}\varepsilon_{i,t-1} + \beta_{\varepsilon,2}\varepsilon_{i,t-2} + \beta_{\varepsilon,3}\varepsilon_{i,t-3} + \zeta_{it},\right. \\
&\qquad\qquad \left. \eta_{i,t-2} + \eta_{i,t-3} + \varepsilon_{i,t-2} - \varepsilon_{i,t-4}\right) \\
&= \beta_{\varepsilon,2}\sigma_{\varepsilon}^2
\end{aligned}
$$

The results follow by taking the ratio of these numerators and the denominator above.

□

We next show how the parameters of the consumption function can be used to trace out the dynamic response to a transitory shock. That is, we would like to characterize the change in log consumption, in periods t, $t+1$, and $t+2$, as a result of a shock ε_t in period t.

Note that we can write log consumption in period t as:

$$c_t = c_{t-1} + \Delta c_t$$

Thus, the effect on consumption in period t is:

$$\begin{aligned}\frac{\partial c_t}{\partial \varepsilon_t} &= \frac{\partial (c_{t-1} + \Delta c_t)}{\partial \varepsilon_t} \\ &= \frac{\partial (\Delta c_t)}{\partial \varepsilon_t} \\ &= \beta_{\varepsilon,0}\end{aligned}$$

where the second line follows from Assumption IID. Similarly, we can write log consumption in period $t+1$ as:

$$c_{t+1} = c_{t-1} + \Delta c_t + \Delta c_{t+1}$$

It follows that the effect on consumption in period $t+1$ is:

$$\begin{aligned}\frac{\partial c_{t+1}}{\partial \varepsilon_t} &= \frac{\partial (c_{t-1} + \Delta c_t + \Delta c_{t+1})}{\partial \varepsilon_t} \\ &= \frac{\partial (\Delta c_t + \Delta c_{t+1})}{\partial \varepsilon_t} \\ &= \beta_{\varepsilon,0} + \beta_{\varepsilon,1}\end{aligned}$$

Finally, we have, for c_{t+2}:

$$c_{t+2} = c_{t-1} + \Delta c_t + \Delta c_{t+1} + \Delta c_{t+2}$$

Likewise, the effect on consumption in period $t+2$ would be:

$$\begin{aligned}\frac{\partial c_{t+2}}{\partial \varepsilon_t} &= \frac{\partial (c_{t-1} + \Delta c_t + \Delta c_{t+1} + \Delta c_{t+2})}{\partial \varepsilon_t} \\ &= \frac{\partial (\Delta c_t + \Delta c_{t+1} + \Delta c_{t+2})}{\partial \varepsilon_t} \\ &= \beta_{\varepsilon,0} + \beta_{\varepsilon,1} + \beta_{\varepsilon,2}\end{aligned}$$

As in our main analysis, the results above can be extended to account for seasonality, variation due to paycheck frequency, and the use of a coworker instrument to account for labor supply endogeneity by using seasonally adjusted, pay per paycheck, coworker income growth for the second terms in all of the covariances above. Our estimate for $\beta_{\varepsilon,0}$ is the same as before: 0.221. Our estimates for the remaining parameters are $\beta_{\varepsilon,1} = -0.100$ and $\beta_{\varepsilon,2} = -0.0368$. Note that negative estimates of the latter two parameters do not mean that consumption responds negatively to lagged transitory shocks. The effect on the log-level of consumption is the cumulative sum of the coefficients. Rather, these parameter values imply that the consumption responses are increasingly smaller as time passes, but remain positive. Our constraint that $\beta_{\varepsilon,3} = -(\beta_{\varepsilon,0} + \beta_{\varepsilon,1} + \beta_{\varepsilon,2})$ then ensures that the response ceases after three months, i.e. $\partial c_{t+k}/\partial \varepsilon_t = \beta_{\varepsilon,0} + \beta_{\varepsilon,1} + \beta_{\varepsilon,2} + \beta_{\varepsilon,3} = 0$ for $k \geq 3$.

After using the parameters $\beta_{\varepsilon,k}$ to calculate the change in log consumption in each of these three periods, we can rescale each by the ratio of consumption to income to convert these changes into marginal propensities to consume (MPCs). These estimates are shown period-by-period in the left panel of Figure 5 and cumulatively in the right panel of the same figure. In practice, we obtain estimates of the elasticity parameters and MPCs, and their standard errors, using GMM, with two-way clustering of our variance covariance matrices in the household dimension and separately in the firm dimension.

Appendix K Notes on Welfare Calculations

K.1 Derivation of Equation (25)

We begin with equation (24) from Section 5.1:

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

Passing the expectations operator through the summation operator, and using the definitions of C_{it} and $\tilde{C_{it}}$ from Section 5.1, we have:

$$\sum_{t=1}^{T}\delta^t \frac{(1+\lambda)^{1-\gamma}\,\mathbb{E}\left(\tilde{C_t}^{1-\gamma}\right)\mathbb{E}\left(e^{(1-\gamma)\beta_\varepsilon \varepsilon_t}\right)}{1-\gamma} = \sum_{t=1}^{T}\delta^t \frac{\mathbb{E}\left(\tilde{C_t}^{1-\gamma}\right)}{1-\gamma}$$

Subtracting the left-hand side from both sides, we have:

$$\begin{aligned} 0 &= \sum_{t=1}^{T}\left\{\delta^t \frac{\mathbb{E}\left(\tilde{C_t}^{1-\gamma}\right)}{1-\gamma}\left(1-(1+\lambda)^{1-\gamma}\,\mathbb{E}\left(e^{(1-\gamma)\beta_\varepsilon \varepsilon_t}\right)\right)\right\} \\ &= \left(1-(1+\lambda)^{1-\gamma}\,\mathbb{E}\left(e^{(1-\gamma)\beta_\varepsilon \varepsilon_t}\right)\right)\sum_{t=1}^{T}\left\{\delta^t \frac{\mathbb{E}\left(\tilde{C_t}^{1-\gamma}\right)}{1-\gamma}\right\}, \end{aligned}$$

which, assuming $\tilde{C}_t^{1-\gamma}$ is on average nonzero, leaves us with:

$$\begin{aligned} 1-(1+\lambda)^{1-\gamma}\,\mathbb{E}\left(e^{(1-\gamma)\beta_\varepsilon \varepsilon_t}\right) &= 0 \\ \Rightarrow (1+\lambda)^{\gamma-1} &= \mathbb{E}\left(e^{(1-\gamma)\beta_\varepsilon \varepsilon_t}\right) \end{aligned}$$

For the term on the right, recall that we have assumed that ε is a normally distributed variable and $\mathbb{E}\left(e^{\varepsilon}\right) = 1$, which means that $\varepsilon \sim \mathcal{N}(-\sigma_\varepsilon^2/2, \sigma_\varepsilon^2)$. It then follows that:

$$(1-\gamma)\beta_\varepsilon \varepsilon \sim \mathcal{N}((\gamma-1)\beta_\varepsilon \sigma_\varepsilon^2/2, (\gamma-1)^2\beta_\varepsilon^2\sigma_\varepsilon^2)$$

Thus, we have:

$$
\begin{aligned}
(1+\lambda)^{\gamma-1} &= \mathbb{E}\left(e^{(1-\gamma)\beta_\varepsilon \varepsilon_t}\right) \\
&= e^{(\gamma-1)\beta_\varepsilon \frac{\sigma_\varepsilon^2}{2}+(\gamma-1)^2\beta_\varepsilon^2\frac{\sigma_\varepsilon^2}{2}} \\
&= e^{(\gamma-1)\left(\beta_\varepsilon+(\gamma-1)\beta_\varepsilon^2\right)\frac{\sigma_\varepsilon^2}{2}}
\end{aligned}
$$

Taking logs of both sides, we have:

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

Finally, if we use the approximation $\log(1+\lambda)\simeq\lambda$, we have equation (25):

$$
\lambda \simeq \left(\beta_\varepsilon+(\gamma-1)\beta_\varepsilon^2\right)\frac{\sigma_\varepsilon^2}{2}
$$

Note, that since $\beta_\varepsilon=\frac{\partial c_t}{\partial \varepsilon_t}$, the expression can also be written as:

$$
\lambda \simeq \left(\frac{\partial c_t}{\partial \varepsilon_t}+(\gamma-1)\left(\frac{\partial c_t}{\partial c_t}\right)^2\right)\frac{\sigma_\varepsilon^2}{2}
$$

Note, when $\beta_\varepsilon=1$, the expression reduces to the Lucas (1987) result:

$$
\lambda \simeq \gamma\frac{\sigma_\varepsilon^2}{2}
$$

As we can see, λ is increasing in σ_ε^2 and increasing in γ. Finally, we have:

$$
\frac{\partial\lambda}{\partial\beta_\varepsilon} = \left(1+2(\gamma-1)\beta_\varepsilon\right)\frac{\sigma_\varepsilon^2}{2}
$$

When $\gamma\geq1$, λ is unambiguously increasing in β_ε.

Note, when $\gamma=1$, utility takes a log functional form, but equation (25) still holds. Equation (24) is now:

$$
\mathbb{E}\left(\sum_{t=1}^{T}\delta^t\log\left((1+\lambda)C_t\right)\right) \equiv \mathbb{E}\left(\sum_{t=1}^{T}\delta^t\log\left(\tilde{C}_t\right)\right)
$$

Simplifying, we have:

$$
\begin{aligned}
\mathbb{E}\left(\sum_{t=1}^{T}\delta^{t}\left(\log\left(1+\lambda\right)+\log\left(\tilde{C}_{t}\right)+\log\left(e^{\beta_{\varepsilon}\varepsilon_{t}}\right)\right)\right) &= \mathbb{E}\left(\sum_{t=1}^{T}\delta^{t}\log\left(\tilde{C}_{t}\right)\right)\\
\Rightarrow \sum_{t=1}^{T}\delta^{t}\log\left(1+\lambda\right) &= -\sum_{t=1}^{T}\delta^{t}\mathbb{E}\left(\log\left(e^{\beta_{\varepsilon}\varepsilon_{t}}\right)\right)\\
&= -\sum_{t=1}^{T}\delta^{t}\mathbb{E}\left(\beta_{\varepsilon}\varepsilon_{t}\right)\\
&= \sum_{t=1}^{T}\delta^{t}\beta_{\varepsilon}\frac{\sigma_{\varepsilon}^{2}}{2}\\
\Rightarrow \log\left(1+\lambda\right) &= \beta_{\varepsilon}\frac{\sigma_{\varepsilon}^{2}}{2}
\end{aligned}
$$

Once again, using the approximation $\log\left(1+\lambda\right)\simeq\lambda$, we have:

$$
\lambda\simeq\beta_{\varepsilon}\frac{\sigma_{\varepsilon}^{2}}{2},
$$

which is the same as equation (25), when $\gamma=1$.

K.2 Welfare Analysis with a Dynamic Consumption Response

We can extend the calculations in Appendix K.1 to account for a consumption function that includes more general responses to lagged transitory shocks, as in Appendix J. We maintain the same income process as in equation (20), and consider a generalization of the dynamic consumption function in equation (J.1):

$$
\begin{aligned}
\Delta c_{t} &= \Delta\tilde{c}_{t}\left(\Gamma_{t},z_{t},\zeta_{t}\right)+\beta_{\varepsilon,0}\varepsilon_{t}+\beta_{\varepsilon,1}\varepsilon_{t-1}+\beta_{\varepsilon,2}\varepsilon_{t-2}+\beta_{\varepsilon,3}\varepsilon_{t-3}\\
\beta_{\varepsilon,3} &= -\left(\beta_{\varepsilon,0}+\beta_{\varepsilon,1}+\beta_{\varepsilon,2}\right)
\end{aligned}
\tag{K.1}
$$

which in turn implies that the *log* of consumption is:

$$
c_{t}=\tilde{c}_{t}(\Gamma_{t},z_{t},\zeta_{t})+\frac{\partial c_{t}}{\partial\varepsilon_{t}}\varepsilon_{t}+\frac{\partial c_{t}}{\partial\varepsilon_{t-1}}\varepsilon_{t-1}+\frac{\partial c_{t}}{\partial\varepsilon_{t-2}}\varepsilon_{t-2},
$$

where we recall that, as detailed in Appendix J:

$$
\frac{\partial c_{t}}{\partial\varepsilon_{t-k}}=\begin{cases}\beta_{\varepsilon,0} & \text{if}\quad k=0\\ \beta_{\varepsilon,0}+\beta_{\varepsilon,1} & \text{if}\quad k=1\\ \beta_{\varepsilon,0}+\beta_{\varepsilon,1}+\beta_{\varepsilon,2} & \text{if}\quad k=2\\ 0 & \text{otherwise}\end{cases}
$$

For convenience, we assume that lagged shocks, ε_0 and ε_{-1} are drawn, even though consumption and income are not realized in those periods.[11] Once again, the constraint that $\beta_{\varepsilon,3}=-\left(\beta_{\varepsilon,0}+\beta_{\varepsilon,1}+\beta_{\varepsilon,2}\right)$

[11] Alternatively, we can assume that lagged shocks are missing from the expressions for c_1 and c_2, which results in a slight approximation error in the derivations below.

confines the consumption response to transitory shocks to three periods.

We can now write the *level* of consumption as:

$$\begin{aligned} C_t &= e^{\tilde{c}_t} e^{\frac{\partial c_t}{\partial \varepsilon_t}\varepsilon_t + \frac{\partial c_t}{\partial \varepsilon_{t-1}}\varepsilon_{t-1} + \frac{\partial c_t}{\partial \varepsilon_{t-2}}\varepsilon_{t-2}} \\ &= \tilde{C}_t e^{\frac{\partial c_t}{\partial \varepsilon_t}\varepsilon_t + \frac{\partial c_t}{\partial \varepsilon_{t-1}}\varepsilon_{t-1} + \frac{\partial c_t}{\partial \varepsilon_{t-2}}\varepsilon_{t-2}} \end{aligned}$$

As in equation (24), we have an implicit definition of λ_{dynamic}:

$$\mathbb{E}\left(\sum_{t=1}^{T} \delta^t \frac{((1+\lambda_{\text{dynamic}})C_t)^{1-\gamma}}{1-\gamma}\right) \equiv \mathbb{E}\left(\sum_{t=1}^{T} \delta^t \frac{\tilde{C}_t^{1-\gamma}}{1-\gamma}\right)$$

Using a set of steps very similar to those in Appendix K.1, we arrive at the following expression:[12]

$$\begin{aligned} \lambda_{\text{dynamic}} &\simeq \Big(\beta_{\varepsilon,0} + (\beta_{\varepsilon,0} + \beta_{\varepsilon,1}) + (\beta_{\varepsilon,0} + \beta_{\varepsilon,1} + \beta_{\varepsilon,2}) \\ &\qquad + (\gamma-1)\left(\beta_{\varepsilon,0}^2 + (\beta_{\varepsilon,0} + \beta_{\varepsilon,1})^2 + (\beta_{\varepsilon,0} + \beta_{\varepsilon,1} + \beta_{\varepsilon,2})^2\right)\Big)\frac{\sigma_\varepsilon^2}{2} \\ &= \left(\frac{\partial c_t}{\partial \varepsilon_t} + \frac{\partial c_t}{\partial \varepsilon_{t-1}} + \frac{\partial c_t}{\partial \varepsilon_{t-2}} + (\gamma-1)\left(\left(\frac{\partial c_t}{\partial \varepsilon_t}\right)^2 + \left(\frac{\partial c_t}{\partial \varepsilon_{t-1}}\right)^2 + \left(\frac{\partial c_t}{\partial \varepsilon_{t-2}}\right)^2\right)\right)\frac{\sigma_\varepsilon^2}{2} \qquad \text{(K.2)} \\ &= \sum_{j=0}^{2}\left(\frac{\partial c_t}{\partial \varepsilon_{t-j}} + (\gamma-1)\left(\frac{\partial c_t}{\partial \varepsilon_{t-j}}\right)^2\right)\frac{\sigma_\varepsilon^2}{2} \end{aligned}$$

K.3 Incorporating Predictable Fluctuations in Income

Above we derived an expression for welfare gain when eliminating variation driven by unpredictable and transitory shocks to income. It may be of interest to expand the types of variation included in this calculation, for example, if liquidity constraints prevent households from completely smoothing even predictable shocks. The two main sources of predictable income that we account for in our method are changes due to seasonal variation in income and changes due to the number of paychecks in a month. We show in this appendix how one might incorporate these two sources of variation, beginning with number of paychecks. In that case, we also quantity the welfare gains of additionally eliminating this source of variation. We then show how one might also separately account for seasonal variation, but leave quantification of that additional component for future work.

To begin, we highlight these two sources of arguably predictable variation by decomposing the term Γ_t:

$$\Gamma_t = m_t + \varepsilon_t^s \qquad \text{(K.3)}$$

where m_t, as in Appendix B.2, is the log of the number of paychecks in a month and ε_t^s is a seasonal shock to income. We then extend the log consumption function to be linear in the log of number of paychecks:

[12] As in Appendix K.1 , this expression holds both for $\gamma > 1$ and the case where $\gamma = 1$, that is, log utility.

$$\begin{aligned} c_t &= \tilde{c}_t\left(\Gamma_t, z_t, \zeta_t\right) + \beta_\varepsilon \varepsilon_t \\ &= \check{c}_t\left(z_t, \varepsilon_t^s, \zeta_t\right) + \beta_m m_t + \beta_\varepsilon \varepsilon_t \end{aligned} \tag{K.4}$$

This results in an modified expression for the *level* of consumption:

$$\begin{aligned} C_t &= e^{\check{c}_t\left(z_t, \varepsilon_t^s, \zeta_t\right) + \beta_m m_t + \beta_\varepsilon \varepsilon_t} \\ &= \check{C}_t M_t^{\beta_m} e^{\beta_\varepsilon \varepsilon_t} \end{aligned} \tag{K.5}$$

where $\check{C}_t = e^{\check{c}_t}$ and M_t is the number of paychecks. Similar to our exercise above, we eliminate the variation driven by unpredictable transitory shocks by setting $\varepsilon_t = 0$, which preserves the mean of income. Likewise, we substitute M_t with the average number of paychecks per month, $\mathbb{E}\left(M_t\right)$, which also preserve the mean of income. For now, we still allow seasonal variation due to ε_t^s to remain. Finally, we ask what increase in consumption is equivalent to the elimination of both the unpredictable transitory shocks and variation due to number of paychecks. In the case where $\gamma > 1$, we have:

$$\mathbb{E}\left(\sum_{t=1}^{T} \delta^t \frac{\left((1+\lambda_m) C_t\right)^{1-\gamma}}{1-\gamma}\right) \equiv \mathbb{E}\left(\sum_{t=1}^{T} \delta^t \frac{\left(\check{C}_t \mathbb{E}\left(M_t\right)^{\beta_m}\right)^{1-\gamma}}{1-\gamma}\right)$$

We can solve for λ_m using steps similar to those in Appendix K.1. First, we must make an additional orthogonality assumption:

$$m_t \perp \left(z_t, \varepsilon_t, \varepsilon_t^s, \zeta_t\right)$$

Passing the expectations operator through the summation operator, and, using the definitions of C_{it} from equation (K.5), we have:

$$\begin{aligned} \sum_{t=1}^{T} \delta^t \frac{\left(1+\lambda_m\right)^{1-\gamma} \mathbb{E}\left(\check{C}_t^{1-\gamma}\right) \mathbb{E}\left(M_t^{(1-\gamma)\beta_m}\right) \mathbb{E}\left(e^{(1-\gamma)\beta_\varepsilon \varepsilon_t}\right)}{1-\gamma} &= \sum_{t=1}^{T} \delta^t \frac{\mathbb{E}\left(\check{C}_t^{1-\gamma}\right) \mathbb{E}\left(M_t\right)^{(1-\gamma)\beta_m}}{1-\gamma} \\ \Rightarrow \left(1+\lambda_m\right)^{1-\gamma} \mathbb{E}\left(M_t^{(1-\gamma)\beta_m}\right) \mathbb{E}\left(e^{(1-\gamma)\beta_\varepsilon \varepsilon_t}\right) \sum_{t=1}^{T} \delta^t \frac{\mathbb{E}\left(\check{C}_t^{1-\gamma}\right)}{1-\gamma} &= \mathbb{E}\left(M_t\right)^{(1-\gamma)\beta_m} \sum_{t=1}^{T} \delta^t \frac{\mathbb{E}\left(\check{C}_t^{1-\gamma}\right)}{1-\gamma} \\ \Rightarrow \left(1+\lambda_m\right)^{1-\gamma} \mathbb{E}\left(M_t^{(1-\gamma)\beta_m}\right) \mathbb{E}\left(e^{(1-\gamma)\beta_\varepsilon \varepsilon_t}\right) &= \mathbb{E}\left(M_t\right)^{(1-\gamma)\beta_m} \\ \Rightarrow \left(1+\lambda_m\right)^{1-\gamma} \mathbb{E}\left(M_t^{(1-\gamma)\beta_m}\right) e^{(\gamma-1)\left(\beta_\varepsilon + (\gamma-1)\beta_\varepsilon^2\right)\frac{\sigma_\varepsilon^2}{2}} &= \mathbb{E}\left(M_t\right)^{(1-\gamma)\beta_m}, \end{aligned}$$

where in the last line, we made use of the fact that ε_t is distributed $\mathcal{N}\left(-\sigma_\varepsilon^2/2, \sigma_\varepsilon^2\right)$. Taking logs of both sides, we have:

$$
\begin{aligned}
(1-\gamma)\log(1+\lambda_m)+\log\mathbb{E}\left(M_t^{(1-\gamma)\beta_m}\right)+(\gamma-1)\left(\beta_\varepsilon+(\gamma-1)\beta_\varepsilon^2\right)\frac{\sigma_\varepsilon^2}{2}&=\log\mathbb{E}\left(M_t\right)^{(1-\gamma)\beta_m}\\
\Rightarrow \log(1+\lambda_m)+\frac{1}{1-\gamma}\log\mathbb{E}\left(M_t^{(1-\gamma)\beta_m}\right)-\left(\beta_\varepsilon+(\gamma-1)\beta_\varepsilon^2\right)\frac{\sigma_\varepsilon^2}{2}&=\frac{1}{1-\gamma}\log\mathbb{E}\left(M_t\right)^{(1-\gamma)\beta_m}\\
\Rightarrow \log(1+\lambda_m)+\frac{\beta_m}{(1-\gamma)\beta_m}\log\mathbb{E}\left(M_t^{(1-\gamma)\beta_m}\right)-\left(\beta_\varepsilon+(\gamma-1)\beta_\varepsilon^2\right)\frac{\sigma_\varepsilon^2}{2}&=\frac{1}{1-\gamma}\log\mathbb{E}\left(M_t\right)^{(1-\gamma)\beta_m}\\
\Rightarrow \log(1+\lambda_m)+\beta_m\log\left(\mathbb{E}\left(M_t^{(1-\gamma)\beta_m}\right)^{\frac{1}{(1-\gamma)\beta_m}}\right)-\left(\beta_\varepsilon+(\gamma-1)\beta_\varepsilon^2\right)\frac{\sigma_\varepsilon^2}{2}&=\beta_m\log\mathbb{E}\left(M_t\right)
\end{aligned}
$$

After gathering terms, we can show:

$$
\log(1+\lambda_m)=\left(\beta_\varepsilon+(\gamma-1)\beta_\varepsilon^2\right)\frac{\sigma_\varepsilon^2}{2}+\beta_m\left(\log\mathbb{E}\left(M_t\right)-\log\left(\mathbb{E}\left(M_t^{(1-\gamma)\beta_m}\right)^{\frac{1}{(1-\gamma)\beta_m}}\right)\right)
$$

And finally, using an approximation for $\log(1+\lambda_m)$, we have:

$$
\lambda_m\simeq\left(\beta_\varepsilon+(\gamma-1)\beta_\varepsilon^2\right)\frac{\sigma_\varepsilon^2}{2}+\beta_m\left(\log\mathbb{E}\left(M_t\right)-\log\left(\mathbb{E}\left(M_t^{(1-\gamma)\beta_m}\right)^{\frac{1}{(1-\gamma)\beta_m}}\right)\right)
$$

In the case where $\gamma=1$, i.e. log utility, we start again with:

$$
\sum_{t=1}^{T}\delta^t\mathbb{E}[\log((1+\lambda_m)C_t)]=\sum_{t=1}^{T}\delta^t\mathbb{E}\left(\log\left(\check{C}_t\mathbb{E}\left(M_t\right)^{\beta_m}\right)\right)
$$

Substituting for C_t using equation (K.5), expanding the log of products, and passing through the expectation function, we have:

$$
\begin{aligned}
\sum_{t=1}^{T}\delta^t\left(\log(1+\lambda_m)+\mathbb{E}\left(\log\check{C}_t\right)+\mathbb{E}\left(\beta_m\log M_t\right)+\mathbb{E}\left(\beta_\varepsilon\varepsilon_t\right)\right)&=\sum_{t=1}^{T}\delta^t\left(\mathbb{E}\left(\log\check{C}_t\right)+\beta_m\log\left(\mathbb{E}\left(M_t\right)\right)\right)\\
\Rightarrow \log(1+\lambda_m)+\beta_m\mathbb{E}\left(\log M_t\right)-\beta_\varepsilon\frac{\sigma_\varepsilon^2}{2}&=\beta_m\log\left(\mathbb{E}\left(M_t\right)\right)\\
\Rightarrow \log(1+\lambda_m)&=\beta_\varepsilon\frac{\sigma_\varepsilon^2}{2}+\beta_m\left(\log\left(\mathbb{E}\left(M_t\right)\right)-\mathbb{E}\left(\log M_t\right)\right)
\end{aligned}
$$

Again approximating $\log(1+\lambda_m)$, and combining with the case of $\gamma>1$ above, we have:

$$
\lambda_m\simeq\begin{cases}\beta_\varepsilon\frac{\sigma_\varepsilon^2}{2}+\beta_m\left(\log\left(\mathbb{E}\left(M_t\right)\right)-\mathbb{E}\left(\log M_t\right)\right) & \text{if}\quad \gamma=1\\[2ex] \left(\beta_\varepsilon+(\gamma-1)\beta_\varepsilon^2\right)\frac{\sigma_\varepsilon^2}{2}+\beta_m\left(\log\mathbb{E}\left(M_t\right)-\log\left(\mathbb{E}\left(M_t^{(1-\gamma)\beta_m}\right)^{\frac{1}{(1-\gamma)\beta_m}}\right)\right) & \text{if}\quad \gamma>1\end{cases}\tag{K.6}
$$

Relative to the case in Section 5, there is an additional term that captures the consumption equivalent of replacing a variable number of paychecks with the mean number of paychecks per month.

To calculate expectations over the distribution of M_t, we use the discrete, empirical distribution of M_t from Table H-1. The expression also requires an estimate of β_m: the elasticity of consumption with respect to predictable changes in number of paychecks per month. We detail below in Appendix K.4 how to estimate this term. The welfare expression is increasing in this elasticity.

We can, in addition, amend our extended welfare calculation to accommodate a dynamic response to shocks, as in Appendix K.2. We extend the consumption function in (K.1) to include responses to predictable variation in income, including, potentially, lags of seasonal variation:

$$\begin{aligned}\Delta c_t &= \Delta \check{c}_t\left(z_t, \varepsilon_t^s, \varepsilon_{t-1}^s, \varepsilon_{t-2}^s, \varepsilon_{t-3}^s, \zeta_t\right) + \beta_m\left(m_t - m_{t-1}\right)\\ &\quad + \beta_{\varepsilon,0}\varepsilon_t + \beta_{\varepsilon,1}\varepsilon_{t-1} + \beta_{\varepsilon,2}\varepsilon_{t-2} + \beta_{\varepsilon,3}\varepsilon_{t-3}\\ \beta_{\varepsilon,3} &= -\left(\beta_{\varepsilon,0} + \beta_{\varepsilon,1} + \beta_{\varepsilon,2}\right)\end{aligned}$$

This implies the following expression for the *log* of consumption:

$$\begin{aligned}c_t &= \check{c}_t\left(z_t, \varepsilon_t^s, \varepsilon_{t-1}^s, \varepsilon_{t-2}^s, \varepsilon_{t-3}^s, \zeta_t\right) + \beta_m m_t\\ &\quad + \beta_{\varepsilon,0}\varepsilon_t + \left(\beta_{\varepsilon,0} + \beta_{\varepsilon,1}\right)\varepsilon_{t-1} + \left(\beta_{\varepsilon,0} + \beta_{\varepsilon,1} + \beta_{\varepsilon,2}\right)\varepsilon_{t-2},\\ &= \check{c}_t\left(z_t, \varepsilon_t^s, \varepsilon_{t-1}^s, \varepsilon_{t-2}^s, \varepsilon_{t-3}^s, \zeta_t\right) + \beta_m m_t\\ &\quad + \frac{\partial c}{\partial \varepsilon_t}\varepsilon_t + \frac{\partial c}{\partial \varepsilon_{t-1}}\varepsilon_{t-1} + \frac{\partial c}{\partial \varepsilon_{t-2}}\varepsilon_{t-2},\end{aligned} \tag{K.7}$$

and a related expression for the *level* of consumption:

$$C_t = \check{C}_t M_t^{\beta_m} e^{\frac{\partial c}{\partial \varepsilon_t}\varepsilon_t + \frac{\partial c}{\partial \varepsilon_{t-1}}\varepsilon_{t-1} + \frac{\partial c}{\partial \varepsilon_{t-2}}\varepsilon_{t-2}}$$

We can now solve for $\lambda_{m,\text{dynamic}}$, i.e. the gain in welfare from eliminating unpredictable transitory shocks and variability in the number of paychecks per month in our more dynamic setting:

$$\mathbb{E}\left(\sum_{t=1}^{T}\delta^t \frac{\left((1+\lambda_{m,\text{dynamic}})C_t\right)^{1-\gamma}}{1-\gamma}\right) \equiv \mathbb{E}\left(\sum_{t=1}^{T}\delta^t \frac{\left(\check{C}_t \mathbb{E}\left(M_t\right)^{\beta_m}\right)^{1-\gamma}}{1-\gamma}\right)$$

Following a series of steps similar to those used above to derive equations (K.2) and (K.6), we arrive at the following expression for $\lambda_{m,\text{dynamic}}$:

$$\lambda_{m,\text{dynamic}} \simeq \begin{cases} \sum_{j=0}^{2}\left(\frac{\partial c_t}{\partial \varepsilon_{t-j}}\frac{\sigma_\varepsilon^2}{2}\right) + \beta_m\left(\log\left(\mathbb{E}\left(M_t\right)\right) - \mathbb{E}\left(\log M_t\right)\right) & \text{if} \quad \gamma = 1\\ \sum_{j=0}^{2}\left(\frac{\partial c_t}{\partial \varepsilon_{t-j}} + (\gamma-1)\left(\frac{\partial c_t}{\partial \varepsilon_{t-j}}\right)^2\right)\frac{\sigma_\varepsilon^2}{2} & \\ \quad + \beta_m\left(\log \mathbb{E}\left(M_t\right) - \log\left(\mathbb{E}\left(M_t^{(1-\gamma)\beta_m}\right)^{\frac{1}{(1-\gamma)\beta_m}}\right)\right) & \text{if} \quad \gamma > 1 \end{cases} \tag{K.8}$$

We can, in principle, extend this analysis even further to also calculate the welfare gain of eliminating predictable, transitory variability due to seasonality. However, we leave for future work an analysis of a sufficient set of assumptions that allow for a welfare calculation that also considers seasonal variation in income.

K.4 Estimation of Welfare Calculation Inputs

A number of parameters are needed as inputs for our welfare calculations in Tables 2 and K-1. First, equation (25) requires an estimate of the variance of unpredictable shocks, including not only the firm component, but also the worker component of ε_t. We use the following:

$$\hat{\sigma}^2_{\varepsilon,\text{worker}} \equiv \text{cov}(\Delta y^p_t, -\Delta y^p_{t+1})$$

In this case, to isolate unpredictable shocks, we use seasonally adjusted pay per paycheck. The key here is that instead of a coworker instrument, we use the worker's own lead of income growth.

Equations (K.6) and (K.8) require the additional input β_m, the response of consumption to changes in the number of paychecks. Starting with our generalized income process in equation (20) and noting our decomposition in equation (K.3), we have:

$$\begin{aligned} y_t &= \Gamma_t + z_t + \varepsilon_t \\ &= z_t + \varepsilon_t + \varepsilon^s_t + m_t \end{aligned}$$

We can recover our parameter of interest using the following:

$$\begin{aligned} \hat{\beta}_m &\equiv \frac{\text{cov}\left(\Delta c_t, -\Delta m_{t+1}\right)}{\text{cov}\left(\Delta y_t, -\Delta m_{t+1}\right)} \\ &= \frac{\text{cov}\left(\Delta \check{c}_t + \beta_\varepsilon\left(\varepsilon_t - \varepsilon_{t-1}\right) + \beta_m\left(m_t - m_{t-1}\right), -m_{t+1} + m_t\right)}{\text{cov}\left(\eta_{t+1} + \varepsilon_t - \varepsilon_{t-1} + \varepsilon^s_t - \varepsilon^s_{t-1} + m_t - m_{t-1}, -m_{t+1} + m_t\right)} \\ &= \frac{\beta_m \sigma^2_m}{\sigma^2_m} \\ &= \beta_m \end{aligned}$$

where σ^2_m is the variance of the log of the number of paychecks in a month, and c_t is defined as in equation (K.4). This identification result continues to hold in the case of our dynamic model of consumption in equation (K.7).

Table K-1 reports our estimated welfare gain from eliminating part of the variation in income. In the first two rows, we repeat our calculations from Table 2, which only feature an elimination of variation due to unpredictable transitory shocks. The next two rows calculate the additional gain of also removing the variability of the number of paychecks per month.

Table K-1: Welfare Gain of Eliminating Transitory Shocks Including Paycheck Variation

	Elasticity of consumption to transitory shocks:			Transitory	Elasticity of consumption to pay frequency	Consumption equivalent of paycheck frequency variation				*Welfare gain:* λ			
						Coefficient of relative risk aversion:				Coefficient of relative risk aversion:			
	$\partial c/\partial \varepsilon_t$	$\partial c/\partial \varepsilon_{t-1}$	$\partial c/\partial \varepsilon_{t-2}$	variance: σ_ε^2	shocks: β_m	$\gamma = 1$	$\gamma = 2$	$\gamma = 3$	$\gamma = 4$	$\gamma = 1$	$\gamma = 2$	$\gamma = 3$	$\gamma = 4$
Unpredictable Shocks													
Monthly	0.22	-	-	0.05	-	-	-	-	-	0.55%	0.67%	0.79%	0.91%
Quarterly (Dynamic)	0.22	0.12	0.08	0.05	-	-	-	-	-	1.05%	1.23%	1.40%	1.57%
Unpredictable/Predictable Mix													
Monthly	0.22	-	-	0.05	0.12	0.96%	1.08%	1.20%	1.32%	1.51%	1.75%	1.99%	2.23%
Quarterly (Dynamic)	0.22	0.12	0.08	0.05	0.12	0.96%	1.08%	1.20%	1.32%	2.01%	2.31%	2.60%	2.89%

Notes: This table shows the welfare gains from eliminating unpredictable transitory shocks (first two rows) as well as from a mix of unpredictable and predictable transitory shocks, including predictable variation in the number of paychecks (second two rows). The welfare gains are computed both in our preferred specification, as well as in its dynamic extension. The table displays all the necessary inputs to compute the structural welfare gains.

K.5 Structural Model

The advantage of our statistical model in Section 5 is that it only requires estimates of the consumption elasticity and transitory income volatility to summarize welfare costs. We now turn to estimating a structural, life-cycle model of consumption and saving, which, among other things, allows households to fully reoptimize under different volatility scenarios and incorporates wealth into the consumption decision. This has several advantages, including capturing a broader measure of welfare and tying more directly into our discussion of liquid assets and consumption smoothing. To demonstrate the welfare implications of volatility we build a model that can match the elasticities we observe in our data. Within this model, we can then shut off transitory income volatility and calculate the expected lifetime welfare improvement for each agent. We find that the type of volatility we observe in our data is economically significant.

We consider a two-asset, discrete time, life-cycle model—a simplified version of Kaplan and Violante (2014). Each period represents a month, matching the frequency of our empirical analysis. Households work until period T_{ret}. The probability of surviving until period t is Λ_t. No households survive beyond period T. Households choose consumption, C_t, to maximize time-separable expected utility:

$$\mathbb{E}_0 \left[\sum_{t=1}^{T} \Lambda_t \delta^{t-1} u(C_t) \right] \tag{K.9}$$

where δ is the per-period discount factor and utility takes a CRRA form: $u(C_t) = \frac{C^{1-\gamma}}{1-\gamma}$. The parameter γ captures risk aversion and intertemporal substitution preferences. The expectation is taken over permanent and transitory shocks to income.

While working, log income in period t, y_t, is a generalization of the income process in equation (20):

$$\begin{aligned} y_t &= \Gamma_t + z_t + \varepsilon_t \\ z_t &= \rho z_{t-1} + \theta_t \eta_t \end{aligned} \tag{K.10}$$

where Γ_t is a predictable, life-cycle growth parameter. The transitory shock, ε_t, is a normal random variable with variance σ_ε and mean $-\sigma_\varepsilon^2/2$, so that $\mathbb{E}[e^\varepsilon] = 1$. The permanent component, z_t, follows an auto-regressive process, with an AR(1) coefficient of ρ. The permanent shock, η_t, is likewise a normal random variable with variance σ_η^2 and mean $-\sigma_\eta^2/2$. The parameter θ_t, a Bernoulli random variable equaling one with probability π_η, governs the arrival rate of the permanent shock. Upon reaching retirement age, T_{ret}, households receive a stream of Social Security payments, which are a fixed percentage of permanent income just prior to retirement. There is no income uncertainty during retirement.

The household may save liquid assets in each period, A_t, which earn a gross return R at the beginning of the next period. Alternatively, they can make a deposit or withdrawal, D_t, using an illiquid asset, which requires an adjustment cost of K dollars. Each period, a gross return of R_N is earned on the balance of the illiquid asset, B_t. As in Kaplan and Violante (2022), we assume that households cannot borrow. The state variables at the beginning of each period are cash on hand, M_t, the pre-decision balance of the illiquid account, N_t, and the level of permanent income, P_t. The transition equations and constraints on decision and state variables are summarized as follows:

$$\begin{aligned}
M_t &= RA_{t-1} + Y_t \\
A_t &= M_t - C_t - D_t - K \cdot \mathbf{1}\{D_t \neq 0\} \\
N_t &= R_N B_{t-1} \\
B_t &= N_t + D_t \\
C_t &> 0 \\
A_t &\geq 0 \\
B_t &\geq 0
\end{aligned} \tag{K.11}$$

We solve the model by adapting the Nested Endogenous Grid Method (NEGM) of Druedahl (2021), which accounts for non-convexities that arise in the presence of the illiquid account.[13] In practice, we normalize all variables by permanent income, $x_t \equiv X_t/P_t$. With a unit root income process, this would reduce the dimensionality of the problem by one. However, because we have a more general AR(1) process, we must retain P_t as a state variable. Nevertheless, this normalization allows the grid points used in solving the model to flexibly adjust with changes in the state variable P_t.[14]

We calibrate several input parameters based on representative survey data and prior literature. Our life-cycle growth parameters, Γ_t, are chosen to match age gradients in median income from the CPS, and are also set so that the steady-state mean of our income process matches mean monthly income at age 35, \$4,242.[15] We set our survival probabilities using mortality tables from the Centers for Disease Control and Prevention (CDC)[16]. The model begins at age 25, and people live at most until age 85. We set the Social Security replacement rate in retirement to 45 percent, which matches average replacement rates (Mitchell and Phillips 2006). The retirement age is set to 62, matching the median age of retirement (Biggs and Springstead 2008). Following Kaplan and Violante (2022), we

[13] In practice, we solve this model using the `EconModel` and `ConSav` Python libraries developed by Jeppe Druedhal and available on GitHub at `https://github.com/NumEconCopenhagen`.

[14] Additional details on solving the model with normalized variables are available upon request.

[15] We use the CPS ASEC and Unbanked/Underbanked supplement from 2017 to 2023 (US Census Bureau 2017–2023; US Census Bureau and FDIC 2017–2023) We focus on respondents aged 25 to 75. Our life-cycle growth parameters are inflation-adjusted to 2023 dollars using the CPI (Bureau of Labor Statistics 2016–2023).

[16] We use the CDC Underlying Cause of Death table, with data from 1999 to 2019 (CDC, National Center for Health Statistics 1999–2019), and focus on individuals aged 25 to 85.

set gross returns on assets such that they match *annualized* values of $R = 0.98$ and $R_N = 1.062$, and we set our transaction cost to $K = \$1,400$.[17] We set our CRRA parameter to $\gamma = 1$, which results in a log utility functional form.

We set the variance of transitory income shocks to match our estimate of 0.05 in Table 2. We set $\pi_\eta = 1/12$, so that permanent shocks arrive on average once per year, as in Kaplan and Violante (2022). Finally, we calibrate the remaining parameters of the AR(1) process for permanent income—$\rho = 0.982$ and $\sigma_\eta^2 = 0.220$—to match two higher order moments of *annual* income, as in Kaplan and Violante (2022): $\text{var}(y_t) = 0.504$ and $\text{var}(\Delta y_t) = 0.142$. These moments are matched conditional on our parameters for the transitory shock and arrival rate π_η. Table K-2 reports the input parameters for the structural model.

Once our model is solved, we simulate 10,000 households at a monthly frequency, from age 25 to 85. We then draw from the steady-state age distribution and calculate various simulated moments. We choose our monthly discount factor, δ, so that an elasticity estimated using our simulated data matches our main estimate of the elasticity of consumption with respect to transitory shocks: $\beta_\varepsilon = 0.221$. We achieve an elasticity of 0.222. Although the elasticity is generally declining in the discount factor over long ranges, there are multiple elasticities in a narrow band that can match this moment. We therefore choose a relatively high discount factor among this set, which helps in matching the wealth moments discussed below. This results in a monthly value of $\delta \approx 0.990$, or at an annualized frequency, 0.888.

Table K-3 reports additional simulated moments, again averaged over the steady-state age distribution. As is the case in the data, the model generates a significant gradient in the elasticity when moving from the lowest to highest asset buffer quartiles, albeit at an even steeper slope than in the data. The monthly MPC implied by the model is comparable to that in the data, while the quarterly MPC is somewhat larger than its empirical counterpart. The model comes close to matching the overall share of households that are hand-to-mouth, although the makeup of this group is slightly skewed toward the "poor" hand-to-mouth, relative to the data. Following Kaplan and Violante (2022), we define a hand-to-mouth household as one that has less than two months of income in liquid assets. We have a reasonable good fit on the ratio of median total wealth to mean income, 2.06 in the model versus 2.03 in the SCF.[18] On the other hand, the ratio of mean total wealth to mean income is significantly higher in our model. The remainder of the moments show that our model cannot fully account for the right tail of the liquid and illiquid asset distributions, although it is able to fit the bottom end of the liquid asset distribution.

At the bottom of Table K-3, we report the welfare gain from eliminating transitory income shocks. We solve the model twice, with our baseline inputs and again with the variance of transitory shocks set to zero. In each case, we calculate the certainty equivalent of the value function at period zero. The certainty equivalent increases by 0.413 percent when transitory shocks are shut off.

[17] Our transaction cost is rounded down from the $1,456.80 in Kaplan and Violante (2022)

[18] We use the 2019 SCF (Board of Governors of the Federal Reserve System 2019), focusing on banked households. We further trim the data of the top five percent of net worth households, following Kaplan and Violante (2022).

Table K-2: Structural Model Parameters

Parameter	Value	Source
Coefficient of relative risk aversion: γ	1.0	N/A
Monthly discount factor: δ	0.990	Calibrated to match Elasticity and Wealth to Income Ratio
Monthly persistence of permanent shocks: ρ	0.982	Calibrated to Kaplan and Violante (2022)
Monthly variance of permanent shocks: σ_η^2	0.220	Calibrated to Kaplan and Violante (2022)
Arrival rate of permanent shocks: π_η	1/12	Kaplan and Violante (2022)
Monthly variance of transitory shocks: σ_ε^2	0.0493	Estimated from Data
Annualized gross liquid return: R	0.980	Kaplan and Violante (2022)
Annualized gross illiquid return: R_N	1.062	Kaplan and Violante (2022)
Illiquid account adjustment cost: K	1,400	Kaplan and Violante (2022)
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
Mean income at 35	4,242	US Census Bureau (2017–2023); US Census Bureau and FDIC (2017–2023)
Lifecycle income growth: Γ_t		US Census Bureau (2017–2023); US Census Bureau and FDIC (2017–2023) Mitchell and Phillips (2006); Biggs and Springstead (2008)
25 to 35	7.09%	
35 to 45	2.25%	
45 to 55	-2.07%	
55 to 61	-13.1%	
61 to 62	-55.9%	
63 to 85	0.00%	
Survival Probabilities: Λ_t		CDC, National Center for Health Statistics (1999–2019)
Pr(Live to 35)	98.9%	
Pr(Live to 45)	97.0%	
Pr(Live to 55)	93.1%	
Pr(Live to 65)	84.9%	
Pr(Live to 75)	69.2%	

Notes: This table reports the inputs for the the structural model in Appendix K.5 and their respective sources. The illiquid adjustment cost from Kaplan and Violante (2022) is rounded down.

Table K-3: Structural Model Results

	Data/Literature	Model
A. Target moments		
Average elasticity: β_ε	0.221	0.22
B. Non-Target Moments		
Pct. Change in elasticity b/w Top and Bottom asset Quartile	-76.5%	-95.8%
Per-Period MPC	0.1	0.103
Quarterly MPC	0.2	0.258
Share HtM	0.41	0.358
Share PHtM	0.14	0.21
Total Asset to Income Ratio (mean)	4.98	7.82
Total Asset to Income Ratio (median)	2.03	2.06
Liquid Asset to Income Ratio (mean)	0.689	1.28
Liquid Asset to Income Ratio (median)	0.0687	0.844
Share with Liquid Assets <$1K	31.5	27.6
Share with Liquid Assets <$5K	51.3	68.5
Share with Liquid Assets <$10K	60.0	87.6
Share with Total Assets <$1K	14.1	21.6
Share with Total Assets <$5K	17.7	44.4
Share with Total Assets <$10K	21.6	52.1
Share with Total Assets <$50K	33.9	79.2
Share with Total Assets <$100K	44.0	92.2
C. Welfare gain from eliminating transitory income shocks		
Average gain (consumption equivalent):	0.413%	

Notes: This table reports simulated moments from the model in Appendix K.5. Panel A shows the target moments from the data, which we use to calibrate the monthly discount factor. Panel B shows a selection of non-targeted moments. Elasticity/MPC estimates are from our data, values for hand-to-mouth households are from Kaplan and Violante (2022), and information on asset distributions and ratios are from the 2019 SCF (Board of Governors of the Federal Reserve System 2019). The definition of liquid assets follows Kaplan and Violante (2014), who subtract credit card debt from assets held in checking and savings accounts. Panel C shows the results from eliminating temporary income volatility and calculating the welfare gain in lifetime consumption equivalents. All simulated moments are averaged over the steady-state age distribution. All numbers are rounded to three significant digits.