Transmission of Income Variations to Consumption Variations: The Role of the Firm

Miao Jin, Yu-Jane Liu, Juanjuan Meng and Yu Zhang[1]

Abstract: We use matched employer-employee data to study the role of the firm in the transmission of income growth into consumption growth. We find that growth in income relative to the firm average (the within-firm component) translates significantly less into consumption than growth in firm average income (the between-firm component). These findings are explained by lower persistence of the within-firm component of income, better self-insurance for workers more exposed to variations in income growth from the within-firm component, and peer effects in the workplace. Quantitatively, income persistence provides 43% of the explanatory power, self-insurance provides 35%, and peer effects provide 22%.

JEL: D31, E21, G51 **Keywords**: income inequality, consumption transmission, firm, permanent income, self-insurance, peer effect

---

[1] Miao Jin is from Renmin University of China; Yu-Jane Liu, Juanjuan Meng and Yu Zhang are from Guanghua School of Management, Peking University. Corresponding author: Yu Zhang, E-mail: yuzhang@gsm.pku.edu.cn. All authors contributed equally to this work. We thank Shachar Kariv (the editor), an anonymous referee, Richard Blundell, Harrison Hong, Yi Huang, Dun Jia, Richard Rogerson, Wei Xiong, Shenghao Zhu, and participants at seminars and conferences for their valuable suggestions. We gratefully acknowledge the financial support of the National Natural Science Foundation of China (grant numbers 71471004, 71673007, 71822301, 71872006, 72073004); China Postdoctoral Science Foundation (grant number 2019M650364); the Solareast Corporation Guanghua Research Fellowship, the Guanghua Thought Leadership Research Grant, and the Key Laboratory of Mathematical Economics and Quantitative Finance (Peking University). All errors are our own.

1 Introduction

Income inequality has risen to unprecedented levels in the United States, China, Russia, and India, and, to a lesser extent, in Europe (Alvaredo, Chanel, Piketty, Saez, and Zucman, 2018). Much has been written about the role of firms in shaping income inequality and the disparity in income growth (Card, Heining, and Kline, 2013; Barth, Bryson, Davis, and Freeman, 2016; Song, Price, Guvenen, Bloom, and von Wachter, 2018; Alvarez, Benguria, Engbom, and Moser, 2018). These studies find that firm average income, the "between-firm" component, and income relative to the firm average, the "within-firm" component, both contribute substantially to the variation in workers' income growth, but these two components have different importance in different economies. Because income influences welfare through consumption, it is crucial to understand the consumption consequences of differences in income growth within and between firms, especially as variations in income growth do not necessarily correspond directly with consumption (Blundell, Pistaferri, and Preston, 2008; Krueger and Perri, 2006; Heathcote, Storesletten, and Violante, 2014).

Thus, in this study we perform the first investigation of whether firms play a role in translating income growth to consumption growth, in addition to firms' traditionally investigated role in income inequality. We use a 2-year, transaction-level dataset of the incomes, consumption patterns, and financials of employees at firms that use the same large commercial bank in Asia for direct deposits. We find that while differences in income growth translate into substantial differences in consumption growth, differences in income growth arising from within-firm sources translate much less into consumption growth than differences in income growth arising from between-firm sources; specifically, the transmission effect of the former is approximately half that of the latter. This shows that the within-firm and between-firm origins of income growth over time have materially different effects on consumption growth.

Our empirical analysis proceeds in several steps. First, to document the role of the firm in

the transmission of income growth to consumption growth over time, we decompose income growth into a between-firm component (growth in firm average earnings) and a within-firm component (growth in log earnings relative to the firm average). We then estimate what proportion of the income growth from the between-firm component and the within-firm component, respectively, translate into consumption growth, and test whether these degrees of consumption transmission differ statistically.

At an annual horizon, we find that the elasticity of consumption growth to income growth arising from the between-firm component is 0.315, and the elasticity of consumption growth to income growth arising from the within-firm component is 0.129, a 59.0% lower degree of consumption transmission than the former. That is, the consumption consequences of income growth from the between-firm component are substantially greater than those of income growth from the within-firm component of income. This result is robust to controlling for demographic factors and holds at different time horizons.

Next, we explore three possible channels to explain the difference in consumption transmission for the two income components. The first channel is income persistence. The permanent income hypothesis states that the sensitivity of consumption growth to income growth depends on persistence. If income growth from the within-firm component is less persistent than income growth from the between-firm component, the former will have a lower effect on consumption growth than the latter. This will occur if the productivity and market power of a firm are more stable than the pay, bonuses, and commissions for its employees. We find empirically that income variations due to the between-firm component are more persistent at the annual frequency than income variations due to the within-firm component, which shows that the income persistence channel is relevant.

The second channel is self-insurance. Wealthier employees who are more able to self-insure (Jappelli and Pistaferri, 2014; Kaplan and Violante, 2014) may be disproportionately

exposed to fluctuations in within-firm incomes than employees who are less able to self-insure. For instance, wealthy employees may work in industries that rely more on performance pay (e.g., bonuses and commissions). Additionally, precautionary savings motives will mean that these employees are likely to accumulate more wealth (Aiyagari, 1994; Kazarosian, 1997; Carroll and Samwick, 1998) than those in other industries. Therefore, observed variations in income growth from the within-firm component are disproportionately generated by wealthier employees. Thus, via a composition effect, we estimate a lower effect on consumption. We find that, in comparison to the between-firm component, variances in income growth from the within-firm component indeed disproportionally arise from wealthy employees, which shows the self-insurance channel is relevant.

The third channel is a behavioral one: the peer effect on consumption. The firm is an ideal setting for the peer effect (Cornelissen, Dustmann, and Schönberg, 2017; De Giorgi, Frederiksen, and Pistaferri, 2019), as employees spend a substantial amount of their time with their peers. Thus, employees may influence each other's consumption decisions, which can lower the sensitivity of consumption growth to income growth from the within-firm component. We find that the transmission of income growth to consumption growth is even weaker in in-firm groups defined by income clusters, which is consistent with the peer effect. Moreover, we control for individuals' income, lagged consumption, peer income, and multiple fixed effects, and thus obtain direct evidence for the relevance of the peer effect channel: lagged peer consumption significantly and positively predicts individuals' consumption and debt, but negatively predicts savings.

With this preliminary support for these three channels, we investigate whether they jointly or separately explain the role of the firm in affecting the transmission of income growth to consumption growth. We examine the three channels' explanatory power with a joint test that uses region × industry variations in the strength of our core findings and the strength of the

prerequisite conditions for the three channels. For each region × industry subsample, we estimate our benchmark specification and use the estimated difference in consumption transmission as the dependent variable in the joint test. Next, we compute measures for the three channels for each subsample: a persistence difference between between-firm and within-firm income, which represents the strength of the persistence channel; a variance ratio for the degree of disproportionate exposure of wealthy employees to income growth from the within-firm component, which represents the strength of the self-insurance channel; and an average degree of consumption peer effects, which represents the strength of the peer effect channel. These measures are the dependent variables.

The joint test indicates that the three channels combine to explain our core findings. One standard deviation change in the strength of the income persistence channel explains 43.9% of the difference in consumption transmission. The corresponding numbers are 35.7% for the self-insurance channel, and 22.6% for the peer effect channel. Summing the three percentages shows that if each of the three channels is weakened by one standard deviation, the difference between the transmissions of income growth from the within-firm and the between-firm components to consumption growth may disappear, and thus that the three channels may collectively account for the difference.

Our study is the first to reveal distinct differences between the consumption consequences of differential income growth arising within a firm and between firms, thereby contributing to the literature on the role of the firm in the determination of income inequality (Card, Heining, and Kline, 2013; Barth, Bryson, Davis, and Freeman, 2016; Song, Price, Guvenen, Bloom, and von Wachter, 2018; Alvarez, Benguria, Engbom, and Moser, 2018; Sorkin, 2018). Previous studies have found that the within-firm and the between-firm components of income growth have different effects on income inequality. Our analysis of our dataset of matched income–consumption financials of firms and their employees shows that this distinction has substantial

implications for employees' consumption.

Our study also contributes to the literature on consumption determination and consumption inequality (Krueger and Perri, 2006; Blundell, Pistaferri, and Preston, 2008; Heathcote, Storesletten, and Violante, 2014; Aguiar and Bils, 2015). Our analysis joins two channels from the macro consumption literature, income persistence (Karahan and Ozcan, 2013; Arellano, Blundell, and Bonhomme, 2017) and self-insurance (Campbell and Deaton, 1989; Kaplan and Violante, 2010), with a channel that has been more widely discussed in the micro consumption literature, the peer effect (Gaviria and Raphael, 2001; Clark and Lohéac, 2007; Moretti, 2011; Smith, Windmeijer, and Wright, 2015; Bertrand and Morse, 2016; De Giorgi, Frederiksen, and Pistaferri, 2019; Han. Hirshleifer, and Walden, 2019; Coibion, Gorodnichenko, Kudlyak, and Mondragon, 2020). The ability of these channels to explain the consumption consequences of income growth differences from within-firm and between-firm sources demonstrates how bridging the two bodies of literature leads to a better understanding of consumption.

The remainder of this paper is structured as follows. Section 2 presents our data. Section 3 lays out the statistical framework. Section 4 presents our findings on the differences in the degree to which the within-firm and the between-firm components of income are transmitted to consumption. In Section 5, we outline three potential channels supporting these differences, examine them individually, and jointly test their ability to explain these differences. Section 6 concludes the study.

2 Data

To study the role of the firm in the transmission of income growth to consumption growth, we use a dataset on the income, consumption, and financials of employees at firms that use the same leading commercial bank in Asia (henceforth "Bank") for direct deposits. The Bank is the largest commercial bank in a high-income Asian economy. It provides direct deposit services for many firms and a wide variety of financial services for these firms' employees.

Our data cover a period of 24 months, from July 2013 to June 2015, which allows us to trace patterns of income, consumption, and financial outcomes for a broad set of employees.

The data are categorized through anonymized identifiers corresponding to the paying firms and the industries to which the firms belong. We also examine information on paycheck and credit-card consumption of the firms' employees. Employees collect paychecks, with income tax deducted, via a checking account at the Bank, allowing us to observe all incoming paychecks for each employee. Credit cards and cash are the two dominant methods of payment in this economy. While we do not observe cash transactions, we have extensive records on credit card transactions for employees in our sample firms. Employees in our dataset use credit cards to pay for a wide range of goods and services, ranging from daily necessities to luxury items, so we can observe their credit card consumption at the transactional level. We supplement the dataset with information on financial behaviors, such as savings account balances, and outstanding debt in credit card and installment loan accounts. We observe the demographic characteristics of the employees, such as age, gender, educational attainment, job position, and marital status. All monetary variables used in the analysis are inflation-adjusted.[1]

We perform the following dataset-cleaning procedure. First, we identify a sample of employees who continuously receive paychecks via the Bank during the 24-month sample period, have a credit card account at the Bank and actively use it every month, and use the Bank as their main credit-card bank (the "main credit card" sample). A main credit-card bank is defined as the bank with which an employee has the highest average credit-card balance. Using this main credit-card sample is advantageous because it covers consumption expenditures well. This sample contains 25,996 employees. Second, we further restrict our

[1] The data on the consumer price index for this economy are obtained from the World Bank. All variables are converted into real values for June 2013.

sample to employees who stay at the same firm throughout the sample period (95.04% of the 25,996 employees), because the small proportion of employees who switch firms (due to low turnover and a relatively short sample period) does not provide sufficient statistical power for tests on these employees. Third, we restrict our sample to full-time employees from 18 to 65 years of age (94.58% of the 25,996 employees), which is a common restriction used in the literature on firms and income distributions (e.g., Song et al., 2018; Alvarez et al., 2018) to minimize the influence of variations in working hours on income measures. Furthermore, following this literature, we only consider employees in firms with ≥20 employees (88.57% of the 25,996 employees), to ensure the within-firm measure is meaningful. This selection procedure yields a main credit-card sample comprising 23,025 employees in 784 firms.

To address the possibility that the selection of employees with a main credit-card relationship with the Bank may not be random, we compare the main credit-card sample to an all direct depositor sample. This sample is a superset of the main credit-card sample that contains all employees who receive paychecks via direct deposit at the Bank and who have a credit card account at the Bank and actively spend on it every month, but whose highest credit-card balance may or may not be on their Bank credit card. After applying the same restrictions on employees and firms, the all direct depositor sample contains 101,493 employees in 861 firms. We show in the analysis that our results are not affected by the above selection criterion because the results for both samples are similar.

DESCRIPTIVE STATISTICS. – Table 1 provides summary statistics for the main credit-card sample. We also provide summary statistics for the all direct depositor sample in Online Appendix Table A1. The main credit card sample is similar to the all direct depositor sample, in terms of the distribution of income, age, marital status, education attainment, employee positions, and firm size.

[Table 1 here]

While the main credit card and all direct depositor samples are similar, they are not representative of the labor force in this economy, as they represent a more skilled and educated component. For example, 65.2% of employees in our dataset are university educated or have a post-graduate degree. In comparison, in official surveys implemented by the economy's government, the percentage of the labor force with a university or post-graduate education is 35.3%. While this difference may be partially due to the inability of official surveys to accurately sample highly educated employees, this difference is also consistent with the possibility that these employees may have more access to banking services. The average annual income in our dataset is 1,110,695 in local currency (~35,000 USD), compared to the official survey average of 602,080 in local currency (~20,000 USD). Therefore, our results only apply to a group of highly educated and wealthy employees who have access to financial services.

The average annual credit-card consumption in our dataset is 157,104 in local currency (~5,300 USD). This consumption is approximately a third (34.9%) of the average personal consumption (450,000 in local currency) in the official survey. This is consistent with credit cards being one of the two major methods of payment in this economy and with the fact that employees in the main credit-card sample also use other credit cards, while maintaining their highest balance with the Bank. Other recent research has used credit card spending to measure consumption (e.g., Gan, 2010; Agarwal and Qian, 2014, 2017). The ratio of credit card consumption to income in our dataset is 14%, similar to Agarwal and Qian (2014) and Agarwal and Qian (2017), who report ratios of 12% and 11%, respectively. These studies use credit-card consumption data from a major bank in Singapore, a comparable economy, to examine the consumption response to income and to housing equity.

The distribution of firm sizes in our dataset is similar to that in the official business census. The average firm size (number of employees) in our dataset is 244, while the median firm size is 74, so the distribution is skewed to the right. The official business census provides a

coarsened distribution of firm size, as it decomposes the number of firms in this economy to those with <5 employees, those with 5–29, 30–49, 50–199, 200–499 employees, and those with >500 employees. We accordingly coarsen the firm size distribution in our data from the Bank and compare the resulting distribution to that of the business census. Figure 1 shows the result: except for a lower density of the smallest firms (≤5 employees) in our dataset, the distribution of firm size in our dataset closely corresponds to that of the official business census.

[Figure 1 here]

3 Statistical Framework

In this section, we describe our statistical framework for understanding the role of a firm in determining how income growth translates into consumption outcomes.

Suppose that the log earnings of employee i in firm f in period t are given by:

$$y_{it} = \alpha_i^w + \alpha_f^F + z_{it}^w + \epsilon_{it}^w + z_{ft}^F + \epsilon_{ft}^F, \tag{1}$$

where the variables indexed by F correspond to firm-level characteristics that apply to all employees in a firm; the variables indexed by w correspond to within-firm characteristics specific to each employee in a firm; the variables α^w (α^F), z^w (z^F), and ϵ^w (ϵ^F) are employee (firm)-fixed effects and the permanent and transitory income components, respectively. In turn, we define η^w and η^F as the innovations in the permanent employee and firm components, respectively: $z_{it}^w = z_{i,t-1}^w + \eta^w$ and $z_{ft}^F = z_{f,t-1}^F + \eta^F$. Following Abowd, Kramarz, and Margolis (1999) and Meghir and Pisteferri (2004), ϵ^F, η^F, ϵ^w and η^w are variations in income growth arising between firms and within the firm. In this study, we examine how much of the growth in average earnings in a firm, the between-firm component of income growth ($\eta^F + \epsilon^F$), translates into consumption growth compared to within-firm individual earning growth relative to the firm average, the within-firm component of income growth ($\eta^w + \epsilon^w$). We also investigate the possible economic channels that may drive this

difference.

Following Blundell, Pistaferri, and Preston (2008) and Heathcote, Storesletten, and Violante (2014), we develop a framework to study the differences between the degree of transmission of income growth to consumption growth due to the between-firm component of employee's income and that due to the within-firm component of employee's income. We formulate the log credit-card consumption of employee i in firm f in period t as:

$$c_{it} = \gamma_i^w + \gamma_f^F + \phi^w z_{it}^w + \psi^w \epsilon_{it}^w + \phi^F z_{ft}^F + \psi^F \epsilon_{ft}^F + \zeta_{it}, \quad (2)$$

where γ_i^w (γ_f^F) are employee (firm)-fixed effects in log credit-card consumption and ϕ^w (ϕ^F); ψ^w (ψ^F) are the degrees of consumption transmission of the permanent and transitory innovations in the within-firm (between-firm) component of the employee's income, η^w (η^F) and ϵ^w (ϵ^F), respectively; and ζ_{it} is an error term that captures the unobserved heterogeneity in consumption.

We observe each employee for 2 years. Therefore, at an annual time horizon, we are restricted to estimating how much of the total innovations in log average earnings in a firm (the between-firm component of income growth, $\eta^F + \epsilon^F$) and how much of the total innovations to log earnings relative to the firm average (the within-firm component of income growth, $\eta^w + \epsilon^w$), translate into consumption.

The main equation in our empirical analysis is developed based on the above framework. Specifically, let $\bar{y}_{ft}^F = \alpha_f^F + z_{ft}^F + \epsilon_{ft}^F$, be firm f's average earnings at time t, defined as the between-firm component of income. This corresponds to Equation (1), where $y_{it} - \bar{y}_{ft}^F = \alpha_i^w + z_{it}^w + \epsilon_{it}^w$ is the amount of an employee's earnings in excess or shortfall of the firm average (the within-firm component of income). Because we focus on the employees who remain at a firm throughout the 2-year sample period, the employee-fixed effects absorb the

firm-fixed effects and we let $\lambda_i = \gamma_i^w + \gamma_f^F$, which results in the following equation:

$$c_{it} = \lambda_i (\equiv \gamma_i^w + \gamma_f^F) + b^F \underbrace{\bar{y}_{ft}^F}_{\text{Between-firm}} + b^w \underbrace{(y_{it} - \bar{y}_{ft}^F)}_{\text{Within-firm}} + \zeta_{it}, \tag{3}$$

where b^F is the statistical degree of consumption transmission of income growth from the between-firm component, and b^w is the statistical degree of consumption transmission of income growth from the within-firm component. We then take the first difference, which eliminates the employee-fixed effects, and obtain the main equation in our empirical analysis, which relates consumption growth to income growth in the two components, as follows:

$$\Delta c_{it} = b^F \underbrace{\Delta \bar{y}_{ft}^F}_{\text{Between-firm}} + b^w \underbrace{\Delta (y_{it} - \bar{y}_{ft}^F)}_{\text{Within-firm}} + \Delta \zeta_{it} \cdot \tag{4}$$

where b^F is a weighted average of ϕ^F and ψ^F, and b^w is a weighted average of ϕ^w and ψ^w.

We focus on estimating b^F and b^w, and the difference between them, i.e., the difference between the statistical degree of consumption transmission of income growth from the between-firm component (growth in firm average earnings) and that of income growth from the within-firm component (growth of employee earnings relative to the firm average). The baseline regression specification of our study, Equation (5), is a transformation of Equation (4). Specifically:

$$\begin{aligned} \Delta c_{it} &= b^F \underbrace{\Delta \bar{y}_{ft}^F}_{\text{Between-firm}} + b^w \underbrace{\Delta (y_{it} - \bar{y}_{ft}^F)}_{\text{Within-firm}} + \Delta \zeta_{it} \\ &= b^F \Delta y_{it} + (b^w - b^F) \Delta (y_{it} - \bar{y}_{ft}^F) + \Delta \zeta_{it} \\ &\equiv b^F \Delta y_{it} + d \cdot \Delta (y_{it} - \bar{y}_{ft}^F) + \Delta \zeta_{it} \end{aligned} \tag{5}$$

where $d \equiv b^w - b^F$ denotes the difference between the within-firm (b^w) and between-firm (b^F) coefficients of consumption transmission. The null hypothesis $H_0 : d = 0$ ($b^F = b^w$), i.e., that the difference between the two consumption sensitivities is zero, allows for a direct test of

the role of firms in the transmission of income growth to consumption growth.

We carry out several robustness tests to ensure that the results obtained from Equation (5) are not driven by empirical specifications or by a spurious correlation of the between-firm and within-firm income components caused by omitted variables.

First, we estimate Equation (5) for different lengths of time to ensure that our results are not driven by time horizons. As it is standard in the consumption literature, we use an annual period length for our baseline. We also report results for semi-annual and monthly period lengths, to examine whether income growth measured at different time horizons accounted for by the within-firm component and the between-firm component also lead to the same difference in the degree of consumption transmission.

Second, we follow a two-step procedure that is a standard procedure in the consumption literature (Browning and Lusardi, 1996) to remove the effect of observable demographic characteristics on income and consumption. In the first step, a set of residual measures of log consumption and log income are constructed as the regression residuals of log consumption and log income on a vector of demographic characteristics. Specifically,

$$c_{it} = \alpha + X_{it}\beta + u^{c}_{it}$$

$$\tilde{c}_{it} \equiv u^{c}_{it} \tag{6}$$

where X_{it} represents consumption taste shifters, such as age, age squared, gender, marital status, and educational attainment. The resulting regression residual $\tilde{c}_{it}$ is the residual measure of consumption. The residual measure of income is defined in the same way. In the second step, we obtain estimates from Equation (5) using residual measures of consumption and income. This alternative estimation ensures that the estimates obtained from Equation (5) are not driven by omitted consumption taste shifters. The estimates based on residual measures of consumption growth and income growth supplement the baseline estimates, for which we

use transparent measures of the raw values of consumption growth and income growth. The first difference in Equation (5) addresses individual-level, time-invariant heterogeneity in consumption tastes.

4 The Difference between Consumption Transmission of Income Growth from the Between-Firm Component and of Income Growth from the Within-firm component

In this section, we report our main results from the estimation of Equation (5). We demonstrate that the consumption transmission of income growth from the within-firm component is much weaker than the consumption transmission of income growth from the between-firm component.

4.1 Main Result

In Panel A of Table 2, we regress the log growth in consumption on the log growth in income as an initial check of whether income growth translates into consumption growth in a direct manner. In this preliminary check, we use the overall income growth as the dependent variable and do not distinguish between income growth from the between-firm component and income growth from the within-firm component. We show that with these specifications, employees with a 10% higher year-to-year income growth have a 1% to 2% higher consumption growth (Panel A, columns 1 and 2). The overall effect of income growth on consumption growth is within the range of estimates in the consumption literature (Hall and Mishkin, 1982; Cochrane, 1991; Attanasio and Weber, 1993; Blundell, Pistaferri, and Preston, 2008; Heathcote, Storesletten, and Violante, 2014; Arellano, Blundell, and Bonhomme, 2017).

Panel B of Table 2 shows the decomposition of income growth into the between-firm and the within-firm component described in Equation (5), which is the main result of our study. It reveals that the transmission of income growth arising from the within-firm component to consumption growth is much weaker than that of income growth arising from the between-firm component. Column (1) is our preferred specification because it uses transparent measures of

consumption growth and income growth that clearly represent our main finding. Column (1) shows that while income growth from the between-firm component has a consumption transmission coefficient of 0.315, income growth from the within-firm component has a consumption transmission coefficient of only 0.129. The difference between the two consumption transmission coefficients, -0.186, is highly statistically significant ($p < 0.01$).

[Table 2 here]

This result is observed for all specifications in Table 2. First, we observe that our results are not purely driven by the observable characteristics of employees. Krueger and Perri (2006) show that income variations attributed to gender, race, education, and experience closely correspond with consumption variations. The same, however, is not true for residual variations in income, meaning that income variations cannot be attributed to these characteristics. To control for these factors, we replace raw measures of income growth and consumption growth in the preferred specification with residual measures of income growth and consumption growth, as defined in Equation (6). We report the estimate of Equation (5) based on residual measures in column (2). We find that the consumption transmission coefficient is 0.219 (0.124) for growth in the between-firm (within-firm) component of income. The difference between the two consumption transmission coefficients, -0.095, remains highly significant ($p < 0.01$). This shows that our results hold even after controlling for the factors outlined in Krueger and Perri (2006).

Second, we show that our result is robust to the sample inclusion criterion. As discussed in Section 2, we use the main credit-card sample as the preferred estimation sample. We address the concern that this sample may be specific to the characteristics of employees represented in this sample by estimating Equation (5) using the all direct depositor sample, which is larger. The results for this sample, as reported in columns (3) and (4) in Table 2, are quantitatively similar to those for the main credit-card sample.

There is a large difference between the degrees of consumption transmission of the within-firm and between-firm components of income growth. Figure 2 visualizes this difference. After using observed (residual) measures, the degree to which income growth is transmitted to consumption growth is 59% (43%) less for income growth from the within-firm component than for the between-firm component.

[Figure 2 here]

This difference in consumption transmission of income growth has implications for the traditionally defined role of the firm in driving differences in income growth between employees. If the differences in income growth in the economy are due to differences between the growth in earnings of employees within each firm (the within-firm component), then our estimate indicates such difference in income growth will be associated with a smaller difference than if the difference in income growth is due to differences between the growth in firm average earnings (the between-firm component).

To present the magnitude of the two components of income growth, Panel B of Table 1 shows the variance of the log levels and the log growth of total income and the between-firm and within-firm components. These variances align with the literature on firms and income distribution, such as Comin, Groshen, and Rabin (2009), Barth, Bryson, Davis, and Freeman (2016), Song, Price, Guvenen, Bloom, and von Wachter (2018) and Alvarez, Benguria, Engbom, and Moser, (2018). The variance of income growth from the within-firm component is larger than that of income growth from the between-firm component. However, even though the income growth from the within-firm component appears to have a larger variance, the effect of this variance on consumption growth is dampened by its smaller consumption transmission coefficient in comparison to that of the between-firm component. The variance of income growth from the within-firm component is 5.4 times the variance of income growth from the between-firm component (0.2048 vs 0.0379). Conversely, the consumption transmission of

income growth from the between-firm component is 2.4 times the consumption transmission of income growth from the within-firm component (0.315 vs 0.129). Therefore, the total consumption effect of income growth from the within-firm component is only 0.91 times that of income growth from the between-firm component (0.2048×0.129^2 vs 0.0379×0.315^2).

4.2 Effects on Debt and Savings

In addition to data on credit card consumption, the Bank provides data on debt and bank deposits.[2] We exploit the data on debt and bank deposits to analyze whether employees accumulate more or less credit card debt and bank deposits after firm-level or within-firm shocks. This is a natural test because we expect that our main result will also show on the other end of the employee's budget constraint. If that is true, this test will provide additional evidence for our main result. We therefore examine the effect of income growth from the between- and within-firm components on changes in the flows into the employee's debt levels and the flows into the employee's net asset positions.

We report the results in Table 3. As in Equation (5), the variables on the right are income growth from the within-firm and the between-firm components for the employee. The dependent variables are the second differences in debt and net assets that correspond to income growth according to the budget constraint. Overall, we find that a lower income growth accounted for by either the between- or within-firm component leads to an acceleration in the accumulation of credit card debt, but with the within-firm component having a significantly stronger influence on such debt (column 1). Thus, individuals increase debt to smooth over negative growth from the within-firm component more than to smooth over negative growth from the between-firm component. This is consistent with our main result. Our result also holds

[2] In the online appendix, we use data on bank deposits to show that our main result is unaffected when we include ATM cash withdrawals in the consumption measure.

if we include installment debt in addition to credit card debt in the transactional dataset, although it is used less frequently (column 2). Columns (3) and (4) show the results for net asset accumulation, for which we add information on the deposit side to the debt side in our banking dataset. Once again, we find that income growth from the within-firm component transmits more to changes in net asset accumulation than income growth from the between-firm component. Thus, employees increase debt and draw down assets in the face of a negative change in the within-firm income component, rather than reducing consumption.

[Table 3 here]

5 Channels

In this section, we outline the channels that may contribute to the results discussed in Section 4: income persistence, self-insurance, and the peer effect. We first empirically describe, examine, and measure the effect of each of the three channels. We then test whether these measures have explanatory power to account for the main results.

5.1 Income Persistence

The permanent income hypothesis (Friedman, 1957) states that transitory fluctuations in income do not influence consumption. As tests of this hypothesis (Johnson, Parker, and Souleles, 2006; Blundell, Pistaferri, and Preston, 2008) have yielded mixed results, a looser interpretation of the hypothesis has been incorporated into macroeconomics. In this interpretation, transitory income shocks are more insurable than permanent income shocks (Heathcote, Storesletten, and Violante 2014).

The income persistence channel accounts for the possibility that the between-firm component of income is more persistent than the within-firm component of income, i.e., it holds that consumption sensitivity to income growth from the within-firm component is smaller than that to income growth from the between-firm component. This is plausible because income growth from the within-firm and between-firm components of income stems from different

sources. The former may arise from variations over time in firm-level differences in productivity and market power, while the latter may derive from variations over time in employees' performances, hierarchies, and incentive pay. It is not empirically obvious which component is more persistent.

We begin testing the income persistence channel by estimating a measure of persistence for each of the two components. Using annual panel data, we estimate a set of statistical AR(1) specifications in Equations (7) and (8), where κ^F, κ^w are the unconditional means of each of the two components, and ρ^F, ρ^w are the persistence parameters of interest.

$$\underbrace{\bar{y}_{ft}^{\,F}}_{\text{Between-firm income (t)}} = \rho^F \underbrace{\bar{y}_{f,t-1}^{\,F}}_{\text{Between-firm income (t-1)}} + (1-\rho^F)\kappa^F + e_{ft}^{\,F} \tag{7}$$

$$\underbrace{(y_{it} - \bar{y}_{ft}^{\,F})}_{\text{Within-firm income (t)}} = \rho^w \underbrace{(y_{i,t-1} - \bar{y}_{f,t-1}^{\,F})}_{\text{Within-firm income (t-1)}} + (1-\rho^w)\kappa^w + e_{it}^{\,w} \tag{8}$$

Table 4 presents the estimation results of Equations (7) and (8). To the best of our knowledge, this is the first study to evaluate the relative persistence of the within-firm and between-firm components of income. The statistical AR(1) coefficient of the within-firm income component ρ^w (0.802) is significantly lower than that of the between-firm income component ρ^F (0.930). The difference between the AR(1) coefficients is not driven by an attenuation bias because we observe the true direct deposit of labor income, without a reporting error. This difference indicates that the within-firm component of income is more transitory than the between-firm component of income. To explain why the within-firm component is less persistent, we show that bonuses, a key part of income, are more transitory at the employee level than at the firm level (Online Appendix Table A2). In sum, the findings regarding differential persistence show that the income persistence channel may account for our main result.

[Table 4 here]

5.2 Self-insurance

The self-insurance channel may also account for our results, as this channel represents the possibility that variations in income growth from the within-firm component disproportionately affect employees who are more able to self-insure. For example, employees at firms whose pay structure involves disproportionately large fluctuations in the within-firm component of their income (e.g., bonuses and commissions) may have a strong precautionary savings motive to accumulate larger amounts of wealth to safeguard against these fluctuations. As a result, relative to other employees, they may have both more exposure to variations in income growth from the within-firm component of income and a greater ability to self-insure. The possible presence of such employees due to a composition effect may significantly lower the estimated sensitivity of consumption growth to income growth from the within-firm income component. This is not likely to hold for the between-firm component because it is more persistent and difficult for workers to self-insure against. Nonetheless, the existence of a self-insurance channel, in the form of disproportionate within-firm variances in the exposure of wealthier employees to income variances remains to be explored.

To examine the self-insurance channel, we first estimate the sensitivity of consumption growth to the income growth of wealthier employees. Specifically, in Panel A of Table 5, we re-estimate a regression of consumption growth on income growth and interact income growth with a dummy for high financial wealth. This dummy equals one if the employee has financial wealth (column 1) or a financial wealth-to-income ratio (column 2) above the median sample value, and equals zero otherwise. The results in Panel A of Table 5 consistently show that the coefficient for the interaction term between high financial wealth and income growth is negative. The consumption transmission coefficient for employees with a high financial wealth-to-income ratio (a coefficient of 0.111 when using the preferred specification in column 2) is approximately half that of other employees (a consumption transmission coefficient of

0.228), which is a significant difference.

We then examine whether wealthier employees, relative to other employees, have disproportionately more exposure to variances in income growth from the within-firm component. We compute the variances in the log changes for the between-firm and within-firm components of income separately for the groups of employees with financial wealth-to-income ratios higher and lower than the median. The results (Panel B of Table 5) show that the group with a higher financial wealth-to-income ratio are exposed to a variance of income growth from the within-firm income component of 0.2703, approximately twice that to which the lower wealth-to-income ratio group are exposed (0.1393). This is not driven by larger income variance across the board because the variance of income growth from the between-firm component is approximately equal between the groups (0.0381 vs 0.0374). Panel B of Table 5 shows that the group with a higher financial wealth-to-income ratio are 32.0% more exposed to within-firm income variance than the lower ratio group, but the former group are not more exposed to between-firm income variances. The same findings hold if we split workers by financial wealth instead of financial wealth-to-income ratios. Overall, wealthy employees disproportionately account for the exposure to within-firm income variations over time, so the self-insurance channel helps to explain our main result.

[Table 5 here]

5.3 Peer Effect

The peer effect channel accounts for the possibility that an individual's consumption behavior is influenced by the consumption behaviors of others in an individual's peer group. We consider a firm as a natural location for the peer effect to exist, and define a peer group as a group of employees in a firm who have similar incomes (as employees with similar incomes are likely hold similar positions and therefore interact with one another) and categorize workers by their average income over the sample period. We obtain quantitatively similar results using groups

that evolve over time due to changes in worker income. The income categories for workers in each firm are determined by an optimal clustering algorithm (Makles, 2012) that balances reducing intra-group differences in income with reducing the number of groups. The median number of income groups in a firm is three, and Online Appendix Figure A1 gives the distribution of the number of clusters.

After identifying the peer groups within firms, we find that the transmission of income growth to consumption is weaker for variations within peer groups compared to the transmission of income growth variations between peer groups. We let $\bar{y}_{gt}^{\,G}$ denote the log average income for income cluster g in firm j in year t, and thus $y_{it} - \bar{y}_{gt}^{\,G}$ and $\bar{y}_{gt}^{\,G} - \bar{y}_{ft}^{\,F}$ refer to the within-cluster and the between-cluster (within-firm) income components, respectively. We estimate the following regression, where $d^{GF} \equiv b^G - b^F$ and $d^{wF} \equiv b^w - b^F$ denote the difference in the transmission of between-cluster (within-firm) and between-firm components of income growth variations to consumption, and the difference in the transmission of within-cluster and between-firm components of income growth variations to consumption, respectively:

$$\begin{aligned}\Delta c_{it} &= b^F \underbrace{\Delta \bar{y}_{ft}^{\,F}}_{\text{Between-firm}} + b^G \underbrace{\Delta(\bar{y}_{gt}^{\,G} - \bar{y}_{ft}^{\,F})}_{\text{Between-cluster (Within-firm)}} + b^w \underbrace{\Delta(y_{it} - \bar{y}_{gt}^{\,G})}_{\text{Within-cluster}} + \Delta\zeta_{it} \\ &\equiv b^F \Delta y_{it} + d^{GF}\Delta(\bar{y}_{gt}^{\,G} - \bar{y}_{ft}^{\,F}) + d^{wF}\Delta(y_{it} - \bar{y}_{gt}^{\,G}) + \Delta\zeta_{it}\end{aligned} \tag{9}$$

Table 6 presents the estimation result of Equation (9). We estimate that b^F, d^{GF} and d^{wF} are 0.381, -0.069, and -0.270, respectively, and that d^{wF} is significantly lower than d^{GF} at the 1% confidence level. These estimates indicate that income growth from the within-firm component transmits less to consumption than income growth from the between-firm component, and that income growth variations within income cluster groups within the firm transmit even less to consumption than income growth variations between income cluster

groups within the firm. This preliminary evidence shows that a peer group can affect the transmission of income to consumption.

[Table 6 here]

We also use higher-frequency consumption data at the monthly level to perform a predictive test of the peer effect: i.e., we estimate how the average consumption of an income cluster group affects an employee's consumption decisions. Following the peer effect literature (Clark and Lohéac, 2007; Smith, Windmeijer, and Wright, 2015; Dimmock, Gerken, and Graham, 2018), we estimate the following predictive consumption peer-effect regression:

$$c_{it} = \beta_0 + \beta_1 c_{-i,t-1} + \beta_2 c_{i,t-1} + \beta_3 y_{it} + \beta_4 y_{-i,t} + v_i^{\ w} + \omega_{gt}^{\ G} + \delta_{ft}^{\ F} + \theta_{gj}^{\ GF} + e_{it} \qquad (10)$$

The coefficient of interest is β_1, which captures the effect of peers' lagged consumption ($c_{-i,t-1}$) on an employee's current consumption (c_{it}) after controlling for the employee's lagged consumption ($c_{i,t-1}$), the current income of the employee (y_{it}), and the current income of peers ($y_{-i,t}$). There are well-known empirical challenges to identifying the peer effect, such as the reflection problem, the sorting effect, and the common shock problem (Manski, 1993).[3] Our estimation process addresses these considerations. First, following Clark and Lohéac (2007) and Smith, Windmeijer, and Wright (2015), we focus on the predictive effect of peers' lagged consumption, which addresses the reflection problem. Second, following Card, Heining, and Kline (2013) and Cornelissen, Dustmann, and Schönberg (2017), we address the sorting effect and the common shock problem by controlling for an extensive set of fixed effects: an

[3] The reflection problem is due to employees and peers affecting each other's consumption; the sorting effect is due to similar employees sorting into the same firm, and to peers with similar background characteristics having similar consumption behavior; and the common shock problem is that employees in a firm may be affected by the same shock.

individual fixed effect, time-varying cluster-fixed effects, time-varying firm-fixed effects, and firm-specific cluster effects. Additionally, we address the omitted variable problem created by a possible correlation between peers' lagged consumption and permanent income by controlling for an employee's lagged consumption, which contains information on permanent income expectations.

[Table 7 here]

Columns (1) to (3) in Table 7 present the estimation results, with more controls added in each column. In column (1), we control for the extensive set of fixed effects mentioned above. In column (2), we add an employee's lagged consumption. In column (3), we control for the current income of both an employee and her peers. The results are similar for different specifications and show a significant and positive peer effect. Columns (4) and (5) in Table 7 show that to match peers' consumption, employees will increase their debt and decrease their savings. Overall, the results in Table 7 show that the peer effect on consumption is significantly positive within employee groups with similar incomes.

5.4 Joint Test of Channels

In this section, we investigate whether the three channels, income persistence, self-insurance, and the peer effect, jointly or separately explain the effect of the baseline role of the firm on the consumption transmission of the within-firm and between-firm components of income growth variations. We test this by using region × industry variations in the strength of our main finding as well as the strength of the prerequisite conditions for the three channels.

Our data span 18 regions and 7 industries: agriculture, manufacturing, utilities and public service, commerce, finance, science and education, and healthcare. Every industry is not present in all regions, and we have 101 region × industry groups. For each group, we estimate Equation (5) and use the estimated average differences in consumption pass-throughs ($d_{region \times industry}$) between within-firm and between-firm income changes as the dependent

variables in the joint test of channels.

The regression specification for the joint test of channels is as follows:

$$d_{region\times industry} = \beta_0 + \beta_1 \cdot Persistence_{region\times industry} + \beta_2 \cdot SelfInsuarnce_{region\times industry} + \beta_3 \cdot PeerEffect_{region\times industry} + \varepsilon_{region\times industry}, \quad (11)$$

where $Persistence_{region\times industry}$, $SelfInsurance_{region\times industry}$, and $PeerEffect_{region\times industry}$ are independent variables at the region × industry level that measure the strength of the prerequisite conditions for the income persistence, self-insurance, and peer effect channels, respectively; the regression residual $\varepsilon_{region\times industry}$ represents the unexplained variations in $d_{region\times industry}$; and the coefficients β_1, β_2, and β_3 represent the ability of each channel to explain our core findings.

To construct $Persistence_{region\times industry}$, we estimate the set of statistical AR(1) specifications in Section 5.1 by using Equations (7) and (8) for each region × industry subsample. We develop two AR(1) coefficients, $\rho^F{}_{region\times industry}$ and $\rho^w{}_{region\times industry}$, that measure the persistence of the between-firm and within-firm components of income in each region × industry subsample. We define $Persistence_{region\times industry} = \rho^F{}_{region\times industry} - \rho^w{}_{region\times industry}$, which measures the extent of the persistence difference between the between-firm component of income and the within-firm component at the region × industry level.

To construct $SelfInsurance_{region\times industry}$, we estimate the variances in income growth arising from the between-firm and within-firm components for employees with a higher financial wealth-to-income ratio than the median (wealthier employees) and for employees with a ratio lower than the median (poorer employees) by following panel B of Table 5 in Section 5.2 for each region × industry subsample. $SelfInsurance_{region\times industry}$ is defined as the following variance ratio in the region × industry subsample: $(\sigma^2_{w,h}/\sigma^2_{w,l})/(\sigma^2_{b,h}/\sigma^2_{b,l})$, where $\sigma^2_{w,h}$ ($\sigma^2_{w,l}$), $\sigma^2_{b,h}$ ($\sigma^2_{b,l}$) are the variances in the log changes in the within-firm component of income for

wealthier (or poorer) employees, and the corresponding variances in the log changes in the between-firm component of income for wealthier (or poorer) employees, respectively. This variance ratio measures the degree to which wealthier employees are more exposed to variances of income growth from the within-firm component than poorer employees, that is, the strength of the prerequisite condition of the self-insurance channel at the region × industry level.

We construct $PeerEffect_{region\times industry}$ for each region × industry with two steps. First, we run a time-series consumption peer-effect regression for each employee in our sample, (Equation 10). We follow Section 5.3 and regress monthly observations of individual consumption on the lagged average peers' consumption in the same income-cluster group. We observe each employee for 24 months, so this procedure amounts to running a time-series predictive regression using these monthly consumption observations. We denote the coefficient of the lagged average peers' consumption as the degree of peer effect for the employee if the degree is significantly positive at the 10% confidence level, and zero otherwise. Second, for each region × industry, we define $PeerEffect_{region\times industry}$ as the average degree of peer effect for all employees in the region × industry subsample. $PeerEffect_{region\times industry}$ therefore measures the average strength of the peer effect on consumption at the region × industry level.

After the dependent variable and independent variables are constructed, we estimate Equation (11) at the region × industry level, weighted by the number of employees in each region × industry group. The independent variables are standardized to have a mean of zero and a unit variance. Table 8 presents the estimation results for the joint test of channels represented by Equation (11), which shows that the three channels have substantial explanatory power. Columns (1)–(3) of Table 8 contain the univariate tests for each channel. Column (1) shows that when the persistence difference between the between-firm and the within-firm components of income is larger, the difference in consumption pass-throughs ($d_{region\times industry}$) between within-firm and between-firm income changes becomes more negative, and thus our

core finding becomes more salient. Similarly, column (2) of Table 8 shows that when wealthier employees are disproportionately more exposed to variances of income growth from the within-firm component, $d_{region\times industry}$ also becomes more negative. Column (3) of Table 8 shows that the average degree of the peer effect for region × industry also corresponds to our core finding $d_{region\times industry}$ in the expected direction.

[Table 8 here]

Next, we evaluate how the three channels jointly explain our core finding. The results in column (4) of Table 8 show that all three channels significantly explain our core finding. If we decompose the combined explanatory power of the three channels, the persistence channel contributes 43.0% of the combined explanatory power, while the self-insurance channel contributes 34.9% and the peer effect channel contributes 22.1%.

We also present an alternative way to interpret the economic magnitude of the three competing explanations. Our reduced-form method of examining the statistical importance of the three channels can estimate how much the different levels of transmission of income growth from within-firm and between-firm components to consumption may be weakened if the strength of each channel is reduced by one standard deviation, while holding other factors constant. The results of this estimation are 43.9% (computed by dividing -0.101 by -0.230) for the income persistence channel, 35.7% for the self-insurance channel, and 22.6% for the peer effect channel. Summing the three percentages shows that if each of the three channels is weakened by one standard deviation, the difference between the degree of consumption transmission of income growth from the within-firm and the between-firm components may disappear, and thus that the three channels may collectively account for the difference.

6 Conclusion

Recent research has emphasized the distinct roles of firm-level factors and within-firm factors in driving income inequality and differential income growth. However, little is known about

whether income growth arising from the components of employees' income within and between firms have similar or different effects on outcomes other than income. Using a matched employer-employee dataset on income and consumption, we provide evidence that income growth from the within-firm component and income growth from the between-firm component generate materially different levels of consumption response. Specifically, a same-sized income growth from the within-firm component drives a level of consumption growth that is on average 43% to 59% smaller than that driven by income growth from the between-firm component (i.e., a firm's average earnings). We provide preliminary evidence that this difference in the transmissions of income growth from the within-firm and the between-firm components to consumption is explained by lower persistence of the within-firm component of income, better self-insurance for workers more exposed to variations in income growth from the within-firm component, and peer effects in the workplace.

Our research is of interest to academics and regulators alike. Our results provide the first evidence that fluctuations in income arising from within-firm and between-firm sources have different effects on employees' consumption. Our results also indicate that if a social planner wishes to limit the degree of dispersion in the growth of consumption, she must pay attention to the extent to which the variations over time in the income of employees in the economy are driven by changes in firm-level factors (e.g., market power, see Van Reenen, 2018; Furman and Orszag, 2018) or by within-firm factors (e.g., managerial compensation and pay-structures, see Murphy and Zabojnik, 2004; Bandiera, Barankay, and Rasul, 2007; Pupato, 2017; Mueller, Ouimet, and Simintzi, 2017).

Working individuals spend a third of their time at work (Ramey and Francis, 2009). Due to the increasing availability of matched firm-worker data covering many aspects of employee behavior, we have unprecedented opportunities to grow our understanding of the role of firms in shaping individual economic and financial decisions. Further research should examine

additional aspects of how firms may affect the economic and financial decisions of their employees.

REFERENCES

Abowd, John M., Francis Kramarz, and David N. Margolis. 1999. "High wage workers and high wage firms." *Econometrica*, 67(2): 251-333.

Agarwal, Sumit, and Wenlan Qian. 2014. "Consumption and Debt Response to Unanticipated Income Shocks: Evidence from A Natural Experiment in Singapore." *American Economic Review,* 104(12): 4205-4230.

Agarwal, Sumit, and Wenlan Qian. 2017. "Access to Home Equity and Consumption: Evidence from A Policy Experiment." *Review of Economics and Statistics*, 99(1): 40-52.

Aguiar, Mark, and Mark Bils. 2015. "Has Consumption Inequality Mirrored Income Inequality?" *American Economic Review,* 105(9): 2725-2756.

Aiyagari, S. Rao. 1994. "Uninsured Idiosyncratic Risk and Aggregate Saving." *The Quarterly Journal of Economics*, 109(3): 659–684.

Alvaredo, Facundo, Lucas Chancel, Thomas Piketty, Emmanuel Saez, Gabriel Zucman. 2018. "World Inequality Report 2018." *Belknap Press*.

Alvarez, Jorge, Felipe Benguria, Niklas Engbom, and Christian Moser. 2018. "Firms and the Decline in Earnings Inequality in Brazil." *American Economic Journal*: *Macroeconomics,* 10(1): 149-189.

Arellano, Manuel, Richard Blundell, and Stéphane Bonhomme. 2017. "Earnings and Consumption Dynamics: A Nonlinear Panel Data Framework." *Econometrica*, 85(3): 693-734.

Attanasio, Orazio P., and Guglielmo Weber. 1993. "Consumption Growth, the Interest Rate and Aggregation." *The Review of Economic Studies*, 60(3): 631-649.

Bandiera, Oriana, Iwan Barankay, and Imran Rasul. 2007. "Incentives for Managers and Inequality Among Workers: Evidence from A Firm-level Experiment." *The Quarterly Journal of Economics*, 122(2): 729-773.

Barth, Erling, Alex Bryson, James C. Davis, and Richard Freeman. 2016. "It's Where You Work: Increases in the Dispersion of Earnings across Establishments and Individuals in the United States." *Journal of Labor Economics*, 34(S2): S67-S97.

Bertrand, Marianne, and Adair Morse. 2016. "Trickle-Down Consumption." *Review of Economics and Statistics*, 98(5): 863-879.

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

Browning, Martin, and Annamaria Lusardi. 1996. "Household Saving: Micro Theories and Micro Facts." *Journal of Economic Literature,* 34(4): 1797-1855.

Campbell, John, and Angus Deaton. 1989. "Why is Consumption So Smooth?" *The Review of Economic Studies*, 56(3): 357-373.

Card, David, Joerg Heining, and Patrick Kline. 2013. "Workplace Heterogeneity and the Rise of West German Wage Inequality." *The Quarterly Journal of Economics*, 128(3): 967-1015.

Carroll, C.D. and Samwick, A.A. 1998. "How important is precautionary saving?" *Review of Economics and Statistics*, 80(3): 410-419.

Clark, Andrew E., and Youenn Lohéac. 2007. "It Wasn't Me, It Was Them! Social Influence in Risky Behavior by Adolescents." *Journal of Health Economics*, 26(4): 763-784.

Cochrane, John H. 1991. "A Simple Test of Consumption Insurance." *Journal of Political Economy*, 99(5): 957-976.

Coibion, Olivier, Yuriy Gorodnichenko, Marianna Kudlyak, and John Mondragon. 2020. "Greater Inequality and Household Borrowing: New Evidence from Household Data." *Journal of the European Economic Association*, forthcoming.

Comin, Diego, Erica L. Groshen, and Bess Rabin. 2009. "Turbulent Firms, Turbulent Wages?" *Journal of Monetary Economics*, 56(1): 109-133.

Cornelissen, Thomas, Christian Dustmann, and Uta Schönberg. 2017. "Peer Effects in the Workplace." *American Economic Review,* 107(2): 425-456.

De Giorgi, Giacomo, Anders Frederiksen, and Luigi Pistaferri. 2019. "Consumption Network Effects." *The Review of Economic Studies*, 87(1): 130-163.

Dimmock, Stephen G., William Christopher Gerken, and Nathaniel P. Graham. 2018. "Is Fraud Contagious? Coworker Influence on Misconduct by Financial Advisors." *Journal of Finance*, 73(3): 1417-1450.

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

Furman, Jason, and Peter Orszag. 2018. "A Firm-Level Perspective on the Role of Rents in the Rise in Inequality." *Toward a Just Society. Columbia University Press*, 19-47.

Gan, Jie, 2010. "Housing Wealth and Consumption Growth: Evidence from A Large Panel of

Households. "*Review of Financial Studies*, 23(6): 2229–2267.

Gaviria, Alejandro, and Steven Raphael. 2001. "School-based Peer Effects and Juvenile Behavior." *Review of Economics and Statistics*, 83(2): 257-268.

Hall, Robert E., and Frederic S. Mishkin. 1982. "The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data on Households." *Econometrica*, 50(2): 461-481.

Han, Bing, David A. Hirshleifer, and Johan Walden. 2019. "Visibility Bias in the Transmission of Consumption Beliefs and Undersaving." *Working Paper*.

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante. 2014. "Consumption and Labor Supply with Partial Insurance: An Analytical Framework." *American Economic Review*, 104(7): 2075-2126.

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

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

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

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

Karahan, Fatih, and Serdar Ozkan. 2013. "On the Persistence of Income Shocks Over the

Life Cycle: Evidence, Theory, and Implications." *Review of Economic Dynamics*, 16(3): 452-476.

Kazarosian, M., 1997. "Precautionary savings—a panel study." *Review of Economics and Statistics*, 79(2): 241-247.

Krueger, Dirk, and Fabrizio Perri. 2006. "Does Income Inequality Lead to Consumption Inequality? Evidence and Theory." *The Review of Economic Studies*, 73(1): 163-193.

Makles, Anna. 2012. "Stata Tip 110: How to Get the Optimal K-Means Cluster Solution." *The Stata Journal*, 12(2): 347-351.

Manski, Charles F. 1993. "Identification of Endogenous Social Effects: the Reflection Problem." *The Review of Economic Studies*, 60(3): 531-542.

Meghir, Costas, and Luigi Pistaferri. 2004. "Income variance dynamics and heterogeneity." *Econometrica*, 72(1): 1-32.

Moretti, Enrico. 2011. "Social Learning and Peer Effects in Consumption: Evidence from Movie Sales." *The Review of Economic Studies,* 78(1): 356-393.

Mueller, Holger M., Paige P. Ouimet, and Elena Simintzi. 2017. "Within-Firm Pay Inequality." *The Review of Financial Studies*, 30(10): 3605-3635.

Murphy, Kevin J., and Jan Zabojnik. 2004. "CEO Pay and Appointments: A Market-Based Explanation for Recent Trends." *American Economic Review,* 94(2): 192-196.

Pupato, Germán. 2017. "Performance Pay, Trade and Inequality." *Journal of Economic Theory*, 172: 478-504.

Ramey, Valerie A., and Neville Francis. 2009. "A Century of Work and Leisure." *American*

Economic Journal: Macroeconomics, 1(2): 189-224.

Smith, Sarah, Frank Windmeijer, and Edmund Wright. 2015. "Peer Effects in Charitable Giving: Evidence from the (Running) Field." *The Economic Journal*, 125(585): 1053-1071.

Song, Jae, David J. Price, Fatih Guvenen, Nicholas Bloom, and Till von Wachter. 2018. "Firming Up Inequality." *The Quarterly Journal of Economics*, 134(1): 1-50.

Sorkin, Isaac. 2018. "Ranking Firms Using Revealed Preference." *The Quarterly Journal of Economics*, 133(3): 1331-1393.

Van Reenen, John. 2018. "Increasing Differences Between Firms: Market Power and the Macro-Economy." *Working Paper*.

Figures

Figure 1: Distribution of firm size

This figure presents the distribution of firm size in our sample and in the business census of this economy. The sample period is from July 2013 to June 2015. Firm size is measured as the average number of employees collecting direct deposits in any given month. The business census is an official survey that covers all firms in the economy and is conducted every 5 years. We use the 2016 wave, and the firm size is the number of employees at the time of the survey. The business census reports a coarsened firm-size distribution (firms with <5, 5–29, 30–49, 50–199, 200–499, and >500 employees) and we likewise coarsen our data. The empirical analysis in our study focuses on firms with 20 employees or more to ensure that the within-firm measure is meaningful.

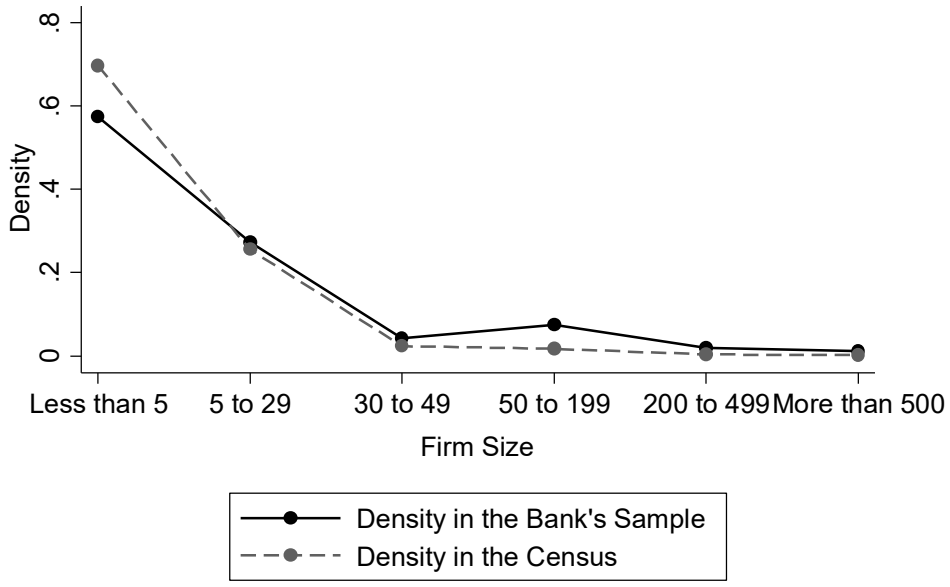

Figure 2: The within-firm versus between-firm difference in consumption transmission

This figure presents the regression estimate of the difference in the extent of consumption transmission of income growth from the between-firm and within-firm components. Detailed results are shown in Table 2. Here we present the results from column (1), in which we use the main credit-card sample. The estimated coefficients b^F and d are 0.315 and -0.186, respectively, and that means the between-firm consumption transmission coefficient and within-firm consumption transmission coefficient are 0.315 and 0.129, respectively.

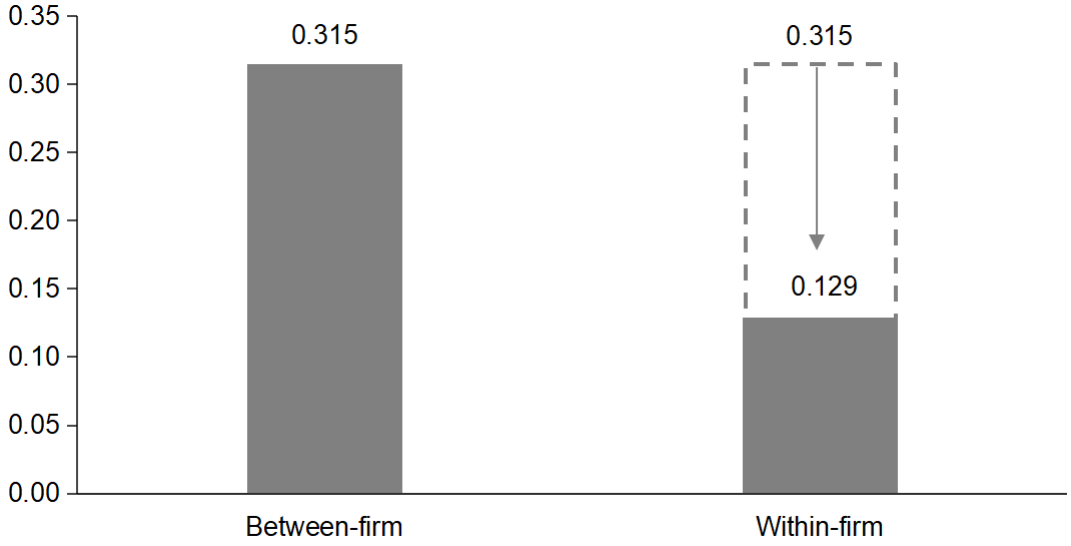

Tables

Table 1: Summary statistics for the baseline sample

Panel A presents the summary statistics of the main credit-card sample at the annual level. Panel B gives the variance and skewness of the log levels and log growth in income and the within-firm and the between-firm components of income. The between-firm component of income is defined as the logarithm of firm average earnings. The within-firm component of income is defined as the difference between log individual earnings and log firm average earnings.

Panel A: Summary statistics

	Count	Mean	S.D.	p25	p50	p75
Income	46050	1,110,695	783,998	729,548	1,064,180	1,663,36
Consumption	46050	157,104	374,292	90,794	148,090	248,568
Savings	46050	632,122	2,102,40	25,927	119,077	505,215
Debt	46050	89,353	464,883	6,965	15,700	35,576
Age	46050	37.7	7.4	32	37	42
Firm size	784	243.9	913.8	38	74	162

Gender	Female	Male
	0.322	0.678
Marital status	Married	Not married
	0.372	0.628
Education	Graduate degree	Bachelor’s degree
	0.246	0.406
	Associate’s degree	High school
	0.188	0.151
	Middle school and below	

	0.009	
Job position	Blue-collar	White-collar
	0.244	0.605
	Managerial	
	0.151	

Panel B: Variances of components of income and components of income growth

	Variance	Skewness
Log of annual value		
Log income	0.7171	1.1180
Between-firm component	0.2746	0.2374
Within-firm component	0.5425	0.5452
Log change in annual value		
Income growth	0.2426	-0.0671
Income growth from between-firm component	0.0378	-0.3313
Income growth from within-firm component	0.2048	-0.0264

Table 2: The within-firm versus between-firm difference in consumption transmission

This table presents results from regressions that estimate the difference in the extent of consumption transmission of growth in the between-firm and within-firm components of income. The data is at the individual level. Columns (1) and (3) use the observed income and consumption values. Columns (2) and (4) use residual income and consumption, as defined in Equation (6). d ($\equiv b^w - b^F$) is the estimated difference in the elasticity of consumption growth to income growth from the within-firm component, b^w, versus that to income growth from the between-firm component, b^F. Heteroscedasticity-consistent standard errors are shown in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% level, respectively.

	(1)	(2)	(3)	(4)
Dependent variable:	*Consumption growth*			
Sample：	Main credit-card		All direct depositor	
	Raw	Residual	Raw	Residual
Panel A: overall effect of income on consumption				
Income growth	0.147***	0.132***	0.127***	0.127***
	(0.010)	(0.010)	(0.008)	(0.008)
R^2	0.012	0.010	0.004	0.004
Panel B: comparison of between-firm and within-firm effects				
b^F (between firm)	0.315***	0.219***	0.336***	0.285***
	(0.021)	(0.023)	(0.022)	(0.025)
d (within firm – between firm)	-0.186***	-0.095***	-0.225***	-0.165***
	(0.021)	(0.023)	(0.022)	(0.025)
Observations	23,025	22,648	101,493	99,896
R^2	0.015	0.010	0.005	0.004

Table 3: The within-firm versus between-firm difference in debt and asset transmission

This table presents the results from regressions that estimate the responses of changes in debt accumulation as well as the responses of changes in net asset accumulation to income growth from the between-firm component and income growth from the within-firm component. The estimation uses the main credit-card sample at the individual level. All debt and net asset variables are normalized by annual labor earnings. In column (1), the dependent variable is the change in credit card debt accumulation. In column (2), the dependent variable is the change in total debt accumulation, which includes credit card debt and installment loans. In column (3), the dependent variable is the change in net asset accumulation, which accounts for deposits and credit card debt. In column (4), the dependent variable is the change in net asset accumulation, which includes deposits, credit card debt, and installment loans. d ($\equiv b^W - b^F$) is the estimated difference in the response to income growth from the within-firm component on the dependent variable, b^W, compared to the response to income growth from the between-firm component, b^F. Heteroscedasticity-consistent standard errors are shown in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% level, respectively.

	(1)	(2)	(3)	(4)
Dependent variable:	Change in debt accumulation	Change in debt accum. (incl. installment loans)	Change in net asset accumulation	Change in net asset accum. (incl. installment loans)
b^F (between firm)	-0.104***	-0.090***	0.205***	0.191***
	(0.006)	(0.007)	(0.014)	(0.013)
d (within firm – between firm)	-0.021***	-0.015**	0.039***	0.032**
	(0.006)	(0.007)	(0.014)	(0.014)
Observations	23,025	23,025	23,025	23,025
R^2	0.019	0.011	0.005	0.005

Table 4: Persistence of within-firm and between-firm income components

This table presents the results from regressions that estimate persistence via a simple AR(1) specification for the logarithms of the between-firm and within-firm components of income by using Equations (7) and (8). The estimations use the main credit-card sample at the annual level. Columns (1) and (2) show the results of the within-firm component of income and the between-firm component of income, respectively. In column (1), the dependent variable is the within-firm component of income, defined as the difference between the log of income and the log of the firm average income, and the independent variable is its lagged value. In column (2), the dependent variable is the between-firm component of income, defined as the log of firm average income, and the independent variable is its lagged value. Heteroscedasticity-consistent standard errors are shown in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% level, respectively.

	(1)	(2)
Dependent variable Y:	*Within-firm income*	*Between-firm income*
Lagged Y	0.802***	0.930***
	(0.005)	(0.003)
Constant	-0.053***	0.907***
	(0.004)	(0.037)
Observations	23,025	23,025
R^2	0.656	0.874
Prob > chi^2 [H0: coef. in (1) = (2)]	0.000	

Table 5: Prerequisites for the self-insurance channel

This table presents the prerequisite tests for the self-insurance channel. The tests use the main credit-card sample at the individual level. In Panel A, *High wealth* is an indicator that equals one if the individual has financial wealth or a financial wealth-to-income ratio above the sample median value. Heteroscedasticity-consistent standard errors are shown in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% level, respectively. Panel B presents the variance comparisons of the income growth components for wealthier employees and the rest of the employees.

Panel A: Consumption transmission of income growth for wealthier and poorer employees

Dep. Var.: Consumption growth	(1)	(2)
	High/Low Wealth	High/Low Wealth-to-Income
Income growth	0.228***	0.228***
	(0.017)	(0.017)
Income growth	-0.115***	-0.117***
× High wealth	(0.021)	(0.020)
Observations	23,025	23,025
R^2	0.014	0.014

Panel B: Variances of income growth components for wealthier and poorer employees

	(1)	(2)	(3)	(4)	(5)
	All	Wealth		Wealth-to-Income	
		Low	High	Low	High
Variance (Between income growth)	0.0378	0.0371	0.0384	0.0374	0.0381
Variance (Within income growth)	0.2048	0.1360	0.2735	0.1393	0.2703
Group value/Population value:					
Variance (Between income growth):		0.9838	1.0162	0.9902	1.0097
Variance (Within income growth):		0.6641	1.3355	0.6802	1.3195

Table 6: Consumption transmission of income growth variations from the within-cluster component of worker income

This table presents regression estimates of differences in the degree of consumption transmission of income growth from the within-cluster component relative to that of income growth from the between-cluster component, and income growth from the between-firm component, using Equation (9). The estimation in this table uses the main credit-card sample at the individual level. b^F is the estimated elasticity of consumption growth to growth in the log firm average income. d^{GF} ($\equiv b^G - b^F$) is the estimated difference in the elasticity of consumption growth to income growth from the between-cluster component, b^G, compared to that for income growth from the between-firm component, b^F. d^{wF} ($\equiv b^w - b^F$) is the corresponding difference in the elasticity of consumption growth to income growth from the within-cluster component, b^w, versus that for income growth from the between-firm component, b^F. Heteroscedasticity-consistent standard errors are shown in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% level, respectively.

Dependent variable:	*Consumption growth*
b^F (between firm)	0.381***
	(0.025)
d^{GF} (between cls (within firm) – between firm)	-0.069***
	(0.027)
d^{wF} (within cls – between firm)	-0.270***
	(0.026)
Observations	23,025
R^2	0.017
Prob > F (F-test: $d^{GF} = d^{wF}$)	0.000

Table 7: Peer effects in consumption

This table presents the results of Equation (10), which estimates the effect of the lagged average peers' consumption on an employee's own current consumption and financial outcomes. The estimation uses the main credit-card sample at the monthly level. The peer group is defined as employees within the same income cluster in the same firm. In columns (1), (2), and (3), the main independent variable is the lagged log average peers' consumption. In columns (4) and (5), the dependent variable is the IHS (inverse hyperbolic sine) of financial outcomes. The main independent variable is the IHS of the lagged average peers' consumption. Standard errors are in parentheses and are clustered at the individual level. *, **, and *** denote significance at the 10%, 5%, and 1% level, respectively.

	(1)	(2)	(3)	(4)	(5)
Dependent variable:	*Log consumption*			*ΔDebt*	*ΔSavings*
Lagged average peers' consumption	0.024***	0.012**	0.012**	0.053***	-0.153***
	(0.005)	(0.005)	(0.005)	(0.010)	(0.042)
Lagged self-consumption		0.007***	0.005**	0.326***	-0.470***
		(0.002)	(0.002)	(0.004)	(0.014)
Self-income			0.100***	0.052***	5.055***
			(0.004)	(0.010)	(0.047)
Average income of colleagues			-0.142***	-0.022	5.369***
			(0.007)	(0.014)	(0.072)
Observations	417,742	417,742	417,742	417,742	417,742
R^2	0.373	0.377	0.379	0.452	0.238
Individual Fixed Effects (FE)	Yes	Yes	Yes	Yes	Yes
Income Cluster × Time FE	Yes	Yes	Yes	Yes	Yes
Firm × Time FE	Yes	Yes	Yes	Yes	Yes
Income Cluster × Firm FE	Yes	Yes	Yes	Yes	Yes

Table 8: Joint tests of channels: income persistence, self-insurance, and peer effect

This table presents the results from regressions that estimate the explanatory power of three channels: income persistence, self-insurance, and the peer effect, by using data at the region × industry level weighted by the number of employees in each region × industry group. The dependent variable *Consumption transmission difference* (*d*) is the estimated average difference in consumption pass-through between "within-firm" and "between-firm" income changes when we run the baseline regression Equation (5) for each region × industry subsample. The channel measures, *Persistence, Self-insurance,* and *Peer effect* are standardized to have a mean of zero with unit variance. Heteroscedasticity-consistent standard errors are shown in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% level, respectively.

	(1)	(2)	(3)	(4)
Dependent variable:	*Consumption transmission difference (d) at region × industry level*			
Persistence	-0.098**			-0.101***
	(0.040)			(0.020)
Self-insurance		-0.045**		-0.082***
		(0.023)		(0.015)
Peer effect			-0.111***	-0.052***
			(0.027)	(0.016)
Constant	-0.230***	-0.230***	-0.230***	-0.230***
	(0.047)	(0.054)	(0.046)	(0.042)
Observations	101	101	101	101
R^2	0.286	0.060	0.365	0.578