

The Matching Multiplier and the Amplification of Recessions*

Christina Patterson[†]

February 2022

First Draft: November 2018

Abstract

This paper shows that the unequal incidence of recessions in the labor market amplifies aggregate shocks. Using administrative data from the United States, I document a positive covariance between worker marginal propensities to consume (MPCs) and their elasticities of earnings to GDP, which is a key moment for a new class of heterogeneous-agent models. I define the *Matching Multiplier* as the increase in the multiplier stemming from this matching of high MPC workers to more cyclical jobs. I show that this covariance is large enough to increase the aggregate MPC by 20 percent over an equal exposure benchmark.

Keywords: Marginal Propensity to Consume, Amplification, Labor Market Inequality
JEL Classification: E21, J11, J23

* I am very grateful to Daron Acemoglu, David Autor, Heidi Williams, and Jonathan Parker for their invaluable guidance and support. A special thanks to the incredibly helpful comments and suggestions from Otis Reid and a thank you to Adrien Auclert for sharing useful code. This paper also benefited greatly from discussions with Martin Beraja, Corina Boar, Jane Choi, Joel Flynn, Colin Gray, Jonathon Hazell, Erik Hurst, Ryan Hill, Raymond Kluender, Benjamin Marx, Pooya Molavi, Toshihiko Mukoyama, Tamar Ostrom, Christopher Palmer, Matt Rognlie, Aysegul Sahin, Carolyn Stein, Michael Stepner, Ludwig Straub, John Sturm, John Van Reenen, Daniel Waldinger, Ivan Werning and Nathan Zorzi and seminar participants at the Minneapolis Fed, Michigan, Fed Board, NY Fed, BU, Chicago Booth, Columbia Business School, Yale, NYU Stern, Maryland, Wharton, Berkeley, Davis, IIES, NBER Summer Institute, Women in Macro conference and CopenhagenMacro days. I gratefully acknowledge financial support from the Alfred P. Sloan Foundation Pre-doctoral Fellowship on the Economics of an Aging Workforce and the Washington Center for Equitable Growth Doctoral Grants and Junior Fellowship programs. Any views expressed are those of the authors and not those of the U.S. Census Bureau. The Census Bureau's Disclosure Review Board and Disclosure Avoidance Officers have reviewed this information product for unauthorized disclosure of confidential information and have approved the disclosure avoidance practices applied to this release. This research was performed at a Federal Statistical Research Data Center under FSRDC Project Number 1616. (CBDRB-FY22-P1616-R9424). This research uses data from the Census Bureau's Longitudinal Employer Household Dynamics Program, which was partially supported by the following National Science Foundation Grants SES-9978093, SES-0339191 and ITR-0427889; National Institute on Aging Grant AG018854; and grants from the Alfred P. Sloan Foundation.

[†]University of Chicago, Booth School of Business. Email: christina.patterson@chicagobooth.edu. All errors are my own.

1 Introduction

The postwar U.S. economy is characterized by periodic large recessions. In the 11 recessions between 1945 and 2011, gross domestic product fell by an average of 2 percent, and the unemployment rate spiked by an average of 2.3 percentage points. Recessions are also unequally distributed across firms and workers. In the labor market, the employment of small and young firms is particularly volatile, as are the earnings of both very low and very high earners (Fort et al. (2013), Guvenen et al. (2017)). This paper shows a link between the heterogeneous impact of aggregate fluctuations and their size. I find that the unequal incidence of aggregate fluctuations in the labor market increases the aggregate marginal propensity to consume, providing a measurement for a key moment in a new class of heterogeneous agent models. I summarize the amplification which comes from matching workers with different MPCs to jobs with different exposures to business cycles by defining what I call the *matching multiplier*, and further show that local economies with greater inequality in the incidence of recessions experience more volatile business cycles.

The response of aggregate consumption to business cycle shocks has been a central focus of macroeconomic research (Mian et al. (2013), Kaplan and Violante (2014)). In both Keynesian and New Keynesian models, the aggregate MPC plays a critical role. However, the aggregate MPC is not itself a fundamental parameter of the economy. Unlike time or risk preferences, which are inputs to a model, the aggregate MPC is a feature of the model that depends on the model's other assumptions. For example, in a representative agent New Keynesian model, the aggregate MPC in response to a transitory income shock is very small because the agent saves to smooth consumption, while in models featuring substantial heterogeneity (e.g., with constrained hand-to-mouth consumers), the aggregate MPC can be an order of magnitude larger (Bilbiie (2020), Galí et al. (2007), Kaplan and Violante (2014)). To elucidate the core mechanism through which the incidence of recessions in the labor market increases the aggregate MPC, consider the case where there are no profits or capital and all income comes from labor earnings. The aggregate MPC is simply given by:

$$MPC = \underbrace{\sum_i \frac{dC_i}{dE_i} \frac{dE_i}{dY}}_{\text{Aggregate MPC}} = \underbrace{\sum_i \frac{E_i}{Y} \frac{dC_i}{dE_i}}_{\text{Avg. MPC}} + \underbrace{Cov\left(\frac{dC_i}{dE_i}, \gamma_i\right)}_{\text{Covariance}} \quad (1)$$

where E_i is the income of individual i , C_i is the consumption of individual i , Y is aggregate output, and $\gamma_i = \frac{dE_i}{dY} \frac{Y}{E_i}$ is the elasticity of individual i 's earnings to aggregate output. The aggregate MPC is made up of two terms. First, it depends on the earnings-weighted average *level* of individual MPCs, $\sum_i \left(\frac{E_i}{Y} \frac{dC_i}{dE_i} \right)$. When individual consumption responds more on average to changes in incomes, the MPC is higher. Second, and importantly for this paper, it depends on the *covariance* between individual MPCs and the elasticities of individual incomes to aggregate movements, $Cov\left(\frac{dC_i}{dE_i}, \gamma_i\right)$. This covariance term captures how the matching of workers with different MPCs to jobs with different sensitivities to aggregate fluctuations affects the magnitude of the aggregate MPC – when shocks disproportionately hit the incomes of individuals whose consumption is more sensitive, the aggregate MPC is larger.

This particular covariance is a core amplification mechanism in a set of new HANK models, highlighted in a series of papers by Bilbiie (2008, 2018a, and 2018b) and by Auclert (2019). For example, in a two-agent New Keynesian model, Bilbiie (2008) derives the result that monetary policy shocks are amplified with agent heterogeneity only when the elasticity of income of the constrained high-MPC agent is above 1. I capture this mechanism empirically with a simple summary measure of how this type of earnings inequality affects amplification – a mechanism I term the *Matching Multiplier*.

The core contribution of this paper is to estimate this key object in the heterogeneous agent macro literature and quantify the degree to which the earnings of high-MPC workers are more exposed to aggregate fluctuations. The main challenge in estimating this empirical moment is that we do not directly observe workers' MPCs.¹ I overcome this challenge by combining information from two data sets. I first use the Panel Study of Income Dynamics (PSID), which is a longitudinal survey that includes measures of both consumption and income, to estimate the MPC for individuals based on their detailed demographic characteristics. I then use several assumptions that I define and test to impute MPCs using those demographics in the matched worker-firm earnings data recorded in the Longitudinal Employer-Household Dynamics data set (LEHD) from the U.S. Census Bureau. While it is possible to impute MPCs and measure earnings elasticities entirely within the PSID, combining the PSID with the LEHD solves three issues: (1) it substantially improves earning measurement and sample size (2) econometrically, it breaks the correlated error structure present when estimating MPCs and earnings elasticities in the same sample (3) it allows me to

¹One exception to this is survey datasets that ask workers to self-report their MPC. I explore these estimates using data from the Italian Survey of Income and Wealth in Section 5.1.

explore firm and geographic heterogeneity.² Within the LEHD, I then explore the degree to which workers with different MPCs are more exposed to movements in aggregate GDP.

Similar to prior literature, I find that individual MPCs, when estimated as the consumption drop per dollar lost upon unemployment, are sizable and heterogeneous. This identifying shock is large and persistent and captures the income variation most similar to what is experienced by workers in recessions. Consistent with other papers in the literature, I estimate an average MPC out of lost labor income of 0.5, with young, black, and low-income workers having higher MPCs. These patterns hold both when examining alternative consumption measures within the PSID and when using other types of income shocks.

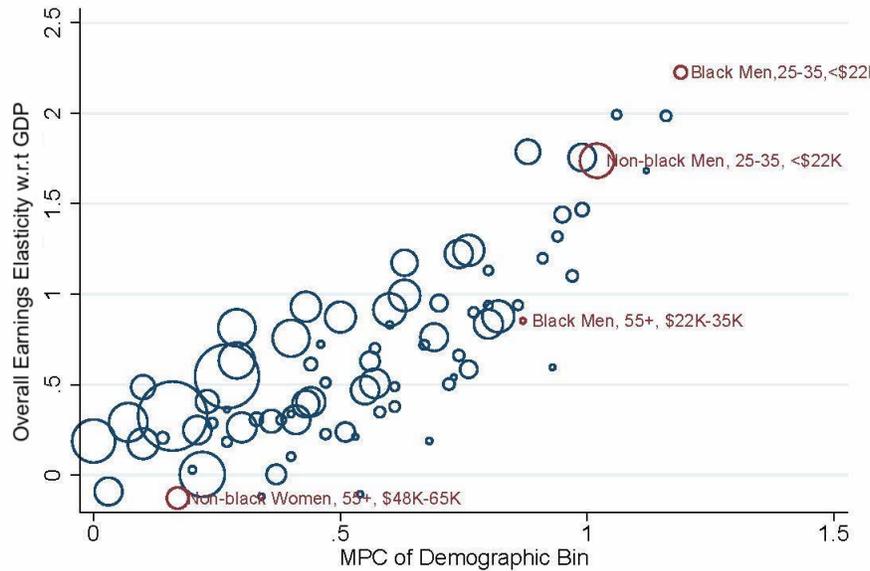
Examining the relationship between MPCs and earnings cyclicalities, I uncover a large, positive covariance between these estimates of a group's MPC and the sensitivity of the group's earnings to aggregate GDP. Figure 1 displays this key positive relationship. Each circle represents a detailed demographic group, the y-axis plots the average earnings elasticity to GDP within that group, and the x-axis plots the average MPC of that demographic group. Across groups, there is a strong positive relationship, with the earnings of high-MPC workers being more exposed to aggregate fluctuations than those of low-MPC workers. I also examine this moment in other contexts and find that the results are similar, holding in the Italian Survey of Income and Wealth (SHIW), holding with MPCs estimated out of purely transitory tax rebates in the United States, and holding within the PSID using total income instead of only labor income. The fact that these patterns are consistent indicates that this mechanism likely operates in a broad set of contexts.

The magnitude of this covariance between worker MPCs and earnings elasticities is large enough to have a meaningful effect on the aggregate MPC. I show this in two distinct ways. First, using the estimated covariance and a partial equilibrium model, I find that the heterogeneous incidence of shocks increases the aggregate marginal propensity to consume out of labor income by 7 percentage points. Due to the non-linearity of the multiplier, this difference in MPCs leads to an increase in the multiplier of around 17 percentage points. Since this multiplier captures the general equilibrium response to a broad class of demand shocks, this amplification mechanism has implications not only for the magnitude of business cycles in general but also for the response of the economy to both fiscal and monetary policy.

Second, in order to empirically validate the importance of this moment, I turn to the geographic variation in the administrative earnings data and show descriptively that areas with a larger covariance be-

²Nonetheless, in Section 5.1, I show that the results are similar when using the PSID alone.

Figure 1: Recession Exposure and MPC by Demographic Group



Notes: Sample includes the set of all workers employed in a sample state in year $t - 1$ from 1995 to 2011. The dependent variable in the regression producing the y-axis estimates is the total change in log earnings for the demographic group. The subgroup is defined in year $t - 1$ and earnings are not conditional on the subgroup into which individuals move in period t . The size of each bubble represents the earnings share of that demographic group. The coefficient on the fitted line for this plot is 1.33. The 80 bins in the figure are defined by the combination of the following categories: gender (men, women), race (black, non-black), lagged income (< 22k, 22-35k, 35-48k, 48-65k, > 65k) and age (25-35, 36-45, 46-55, 56-62). Appendix Figure A8 shows the corresponding figure separately for the intensive and extensive margin of earnings.

tween MPCs and earnings elasticities experience deeper recessions and larger booms. This is exactly what this mechanism would predict under the assumption that a significant share of demand within a commuting zone (CZ) is derived locally. I also find that this difference across local labor markets is entirely concentrated in nontradable industries, where the local consumption response should be much more important. Together, these local estimates provide additional evidence for the empirical relevance of this key covariance.

Overall, estimating this moment is important for three main reasons. First, as I emphasized above, this is a crucial moment for new heterogeneous agent New Keynesian models and the relatively sizable magnitude I find here suggests that these models incorporating this moment better capture dynamics of the economy. Second, my results indicate that structural changes in the economy, such as decreasing unionization or rising inequality, may have cyclical consequences, as any forces that contribute to the matching of high MPC workers to more cyclical jobs will accentuate the multiplier. Finally, my findings have direct implications for macroeconomic stabilization policies. Several countries, such as Germany, offer more gen-

erous unemployment insurance to older workers. My results suggest that in fact the opposite approach might be more effective at reducing the severity of recessions, since it is younger, higher MPC workers who tend to be more exposed to unemployment risk and contribute more to the aggregate MPC.

Related Literature: The analysis in this paper adds to a large literature emphasizing that micro heterogeneity in the consumption responses to income changes is critical in determining aggregate dynamics. Important work in this area includes empirical studies documenting substantial heterogeneity in MPCs at the individual level (Johnson et al. (2006), Fagereng et al. (2021), Jappelli and Pistaferri (2014)) and quantitative models demonstrating the importance of agent heterogeneity in the determining the effectiveness of fiscal (Galí et al. (2007), Kaplan and Violante (2014)) and monetary policy (Auclert (2019), Kaplan et al. (2018), McKay et al. (2016)). Additionally, several other papers have highlighted the key role that the general equilibrium redistribution of income plays in the effectiveness of automatic stabilizers (McKay and Reis (2016)) or fiscal redistribution over the business cycle (Oh and Reis (2012)). This analysis builds on these papers but differs from them in two critical respects. First, I focus on a particular amplification mechanism coming from the covariance of worker MPCs and the sensitivity of their incomes to the business cycle. Second, I undertake a detailed empirical analysis of this channel to quantify this covariance.

One important emphasis in the heterogeneous-agent New-Keynesian literature is that the introduction of constrained, high-MPC workers does little to increase the aggregate MPC, as when these models are calibrated to match the empirical distribution of wealth, high-MPC workers comprise a small share of the economy (Kaplan et al. (2014)). Recent work addresses this concern by adding realistic features of the economy that work to increase the prevalence of high-MPC households (two types of assets in the case of Kaplan and Violante (2014) and Kaplan et al. (2018) or preference heterogeneity in the case of Krusell and Smith (1998) and Carroll et al. (2017)). The mechanism in this paper is a complementary mechanism that increases the aggregate MPC – while high MPC workers may constitute a small share of the economy, if their income is most affected by the aggregate shock, they will become disproportionately important in determining the response.³

As I discussed above, the importance of this particular mechanism is highlighted analytically in several

³ It is important to note that the mechanism emphasized in this paper also differs from the amplification that comes from countercyclical income risk, featured in various ways in Werning (2015), McKay (2017), Heathcote and Perri (2018), and Ravn and Sterk (2017), among others. The matching multiplier mechanism explored here focuses instead on the distribution of realized income, rather than cyclical changes in income risk. Indeed, recent work by Bilbiie (2018) clearly disentangles these two channels and shows that these two forces – countercyclical income risk and heterogeneous incidence of shocks – reinforce each other.

heterogeneous-agent New Keynesian models. Bilbiie (2008, 2018a, and 2018b) terms this channel “cyclical inequality” and demonstrates that the covariance between household MPCs and earnings inequality is a sufficient statistic for whether household heterogeneity amplifies or dampens inequality.⁴ Auclert (2019) similarly highlights the theoretical role that the covariance between worker MPCs and the sensitivity of their incomes to aggregate output plays in amplifying monetary policy shocks. In newer work, Alves et al. (2020) and Slacalek et al. (2020) also model this mechanism and find it is important for amplification in their quantitative model. This paper differs from those papers in its focus on the direct measurement of this key covariance.

Lastly, this project also connects to a growing empirical literature examining the incidence of recessions. Recessions can distribute shocks unequally through several channels, including the housing market (Mian et al. (2013)) or the asset market (Glover et al. (2020)). This project specifically focuses on heterogeneity coming through the labor market. In that context, Hoynes et al. (2012) show that the earnings of young, low-education men are more sensitive to business cycles, and, using high-quality administrative tax records, Guvenen et al. (2017) document that the earnings of both the very low and very high income workers are particularly exposed.⁵ More specifically, in the context of monetary policy shock, Coibion et al. (2017) find evidence that the labor market earnings of low net worth households (likely higher MPC) fall slightly more following a monetary contraction than those for high-net worth households (likely lower MPC). Similarly, Cloyne et al. (2019) show that the elasticity of income with respect a monetary shock is slightly higher and the MPC substantially larger for mortgagors and renters than outright home owners, suggesting a positive covariance between MPCs and income. However, both of these findings about income heterogeneity are in the context of monetary policy and the direct implications for the overall covariance between earnings responses and MPCs is not obvious.

Outline: The rest of the of the paper proceeds as follows. Section 2 defines the matching multiplier in a simplified framework. Section 3 describes the two main data sets that I combine in my empirical analysis. Section 4 presents empirical estimates of the covariance between MPCs and earnings sensitivities to the

⁴While I focus on the importance of this moment for amplification, Bilbiie (2018) shows this covariance may also have important implications for the determinacy of heterogeneous-agent New Keynesian models with interest rate rules and potentially exacerbates a wide class of New-Keynesian puzzles.

⁵ This evidence on the relationship between income sensitivities and lagged incomes does not have any immediate implications for the relationship between income sensitivities and MPCs. First, MPCs are not a direct function of income and vary with other characteristics such as time preferences (Parker (2017)) or liquid wealth (Kaplan and Violante (2014)). Second, since MPCs are generally linearly falling in wealth, there is not clear mapping to the nonlinear relationship between the elasticity of income to changes in aggregate output across the income distribution.

aggregate. Section 5 provides estimates of the matching multiplier and Section 6 uses geographic variation to empirically validate the importance of this amplification mechanism. Section 7 concludes.

2 Defining the Matching Multiplier

The equilibrium assignment of workers to jobs affects the economy's response to aggregate shocks in a wide class of models (Auclert (2019), Carroll et al. (2017), Bilbiie (2020)). While the effect of the covariance on the aggregate MPC is relatively straightforward, the effect on the aggregate multiplier from this moment will depend on the detailed structure of the model being studied. In the online appendix, I more formally model the persistence of unemployment, which is the income shock that identifies the MPCs, and the effect this has on the relationship between MPCs and the multiplier. For now, I consider a simplified 2-period model to clarify the intuition and define the matching multiplier.

To begin, consider the simple case in which worker i has a consumption function $c_i(E_i(Y), \theta_i)$, where E_i are the earnings of individual i , which are given as a function of aggregate output Y , and θ_i are other parameters affecting the consumption of the individual, such as preferences, borrowing constraints, etc. I assume that there are no capital or profits and that all output is consumed by workers, meaning that the market clearing condition dictates $Y = C$, where C represents aggregate consumption. Since I am interested in understanding the importance of demand-side heterogeneity in propagating shocks, I assume that prices are fixed and that output is demand-determined. Therefore, the total derivative of the market-clearing condition yields

$$dY = \sum_i \frac{\partial c_i}{\partial E_i} \frac{\partial E_i}{\partial Y} dY + \underbrace{\sum_i \frac{\partial c_i}{\partial \theta_i} d\theta_i}_{d\varepsilon}$$

where $d\varepsilon$ is the change in total demand in response to an exogenous shock before output adjusts. Define $\gamma_i = \frac{dE_i}{dY} \frac{Y}{E_i}$ as the elasticity of individual i 's earnings to aggregate output and $MPC_i = \frac{\partial c_i}{\partial E_i}$ as the marginal propensity to consume of individual i . Assuming that the aggregate MPC is less than 1 (i.e. $\sum_i \frac{\partial c_i}{\partial E_i} \frac{\partial E_i}{\partial Y} < 1$), this can be expressed as

$$\frac{dY}{d\varepsilon} = \frac{1}{1 - \sum_i \frac{E_i}{Y} MPC_i \gamma_i} \quad (2)$$

where $\sum_i \frac{E_i}{Y} MPC_i \gamma_i = \sum_i \frac{\partial c_i}{\partial E_i} \frac{\partial E_i}{\partial Y}$ is the *actual* aggregate marginal propensity to consume (MPC^a) in the economy – it captures how much of an additional unit of output is translated into an additional unit of consumption demand, taking into account the distribution of the aggregate shock. The multiplier, which is given in this setting by $\frac{1}{1-MPC^a}$, determines the economy’s response to any demand shock $d\varepsilon$. Equation 2 can be rewritten to highlight the role of earnings heterogeneity as

$$\frac{dY}{d\varepsilon} = \frac{1}{1 - (\bar{\gamma} \overline{MPC} + Cov(MPC_i, \gamma_i))} \quad (3)$$

where \overline{MPC} is the earnings-weighted average MPC in the economy, $\bar{\gamma}$ is the elasticity of the average dollar of earnings in the economy to aggregate output, and $Cov(MPC_i, \gamma_i)$ is the earnings-weighted covariance between MPCs and earnings elasticities.⁶ In the benchmark case in which every worker has an earnings elasticity equal to the average, $Cov(MPC_i, \gamma_i) = 0$ and the aggregate MPC is $MPC^b = \bar{\gamma} \overline{MPC}$. However, when the labor earnings of high-MPC workers are more exposed to aggregate movements in output, $Cov(MPC_i, \gamma_i) > 0$, and the aggregate MPC is larger. To explicitly capture the contribution of the covariance term to the multiplier, I define the *matching multiplier, or MM*, as difference between the multiplier when workers face their actual earnings elasticity and a benchmark multiplier where the covariance between worker MPCs and earnings elasticities is 0.

$$MM = \frac{1}{1 - MPC^a} - \frac{1}{1 - MPC^b} \quad (4)$$

3 Data description

3.1 Longitudinal Employer-Household Dynamics (LEHD)

The main data set for this analysis is the U.S. Census Bureau’s Longitudinal Employed-Household Dynamics (LEHD) data set, a longitudinal data set that provides quarterly earnings for all workers covered by the state-level Unemployment Insurance Program. The data for this paper includes a subset of 23 states in an unbalanced panel from 1995 to 2011, a period that covers two recessions. By the late 1990s, this sub-

⁶ In this simple model, the earnings-weighted average elasticity $\bar{\gamma}$ in the economy is one, but in the data, this number will be different, as labor earnings may not move one for one with output because of movements in investment, movements in net exports, or various labor market frictions.

set contained almost 50 percent of total U.S. private employment.⁷ In addition to the quarterly earnings of these workers, this dataset also includes information on their demographics (age, race, gender, and education) and their employers (location, industry, firm size, and age).⁸ Reported quarterly earnings in the LEHD include gross wages and salaries, bonuses, stock options, tips, and other gratuities, as well as the value of meals and lodging (Spletzer (2014)).

In implementing the analysis, I make several sample restrictions. First, I exclude the first two years an individual appears in any state within my sample, allowing me to construct earnings histories and thus impute MPCs for all workers. Second, to abstract from schooling and retirement decisions, I exclude workers who are younger than age 25 or older than age 62. Third, for computational reasons, I restrict my attention to a fourth-quarter snapshot of the data.⁹ Appendix Table A1 shows simple summary statistics for this sample. Over the full sample, I observe an average of 38 million workers each year, about half of whom are male and who earn an average of \$44,000 dollars per year.

3.2 Panel Study of Income Dynamics (PSID)

I supplement the detailed administrative data on earnings with marginal propensities to consume estimated using the PSID. Each household in the PSID is interviewed every year from 1968 to 1997 and every other year after that. Among other things, households are asked detailed questions on their demographics and the labor market experiences of the household head and spouse. In order to include secondary earners in my analysis and estimate a MPC at the individual level for both household heads and spouses, I transform the data to have up to two observations per household, one for the head and one for the spouse. This transformation means that each individual within the household has the same consumption, but the head and spouse differ in their demographics and labor market variables. Appendix Table A3 shows sum-

⁷ In a given state, this data set covers about 95 percent of private sector employment. See Appendix Table A2 for the list of included states, as well as the years for which each state is in the sample. I include all states to which I was given access. Appendix Table A3 shows that the demographic and labor market characteristics of workers in the LEHD states are very similar to those nationally over the same sample period.

⁸ Data on worker demographics in the LEHD come from the internal Census Personal Characteristic file, which covers 95 percent of the sample and is used to identify individuals across all census transactions and the long- and short-form censuses, which cover 61 and 12 percent of the sample, respectively (Vilhuber and McKinney (2014)). Specifically, worker age and gender are recorded in the person characteristic file, race is sourced from the short-form census, and education is sourced from the long-form census. In all cases, missing demographic data are imputed by the Census Bureau. Because education is imputed for almost 90 percent of the sample, I use this variable only in robustness exercises but not in the baseline results reported throughout this analysis. See Abowd et al. (2009) for further details on the data construction.

⁹ It is common in the literature to restrict analysis of the LEHD to an annual snapshot. See Sorokin (2018) and Abowd et al. (2003). I explore the robustness of the results to using annual earnings data rather than fourth-quarter earnings and find that results are similar.

mary statistics for the PSID sample that I include in the estimation.¹⁰ Because I include all samples within the PSID to estimate MPCs (i.e., the nationally representative sample and the oversample of low-income groups), the sample has a higher share of black and low-income workers than the total population.

For most of the PSID sample, the main expenditure variable is food consumption. There are several reasons to suspect that the response of food expenditure to income changes is not representative of overall expenditure responses. First, as food is a necessity, its share of total consumption varies across the income distribution (see Appendix Figure A1). Second, the provision of food stamps potentially distorts food consumption decisions on the margin and likely dampens fluctuations in food expenditure relative to overall expenditures (Hastings and Shapiro (2018)). In order to address these issues, the main measure of expenditure I use throughout the analysis is total expenditure, which I impute using overlapping information in the Panel Study of Income Dynamics and the Consumer Expenditure Survey (CEX), following the methodology introduced in Blundell et al. (2008) and expanded in Guvenen and Smith (2014). Using the CEX, I estimate a demand for food expenditure as a function of durable consumption, nondurable consumption, demographic variables, and relative prices. Under the assumption that food demands are monotonic, this demand function can be inverted to get predicted total consumption based on the food expenditures and demographics of the household in the PSID. Appendix Figure A2 shows that this imputation captures important cross-sectional patterns in the CEX such as the food share of consumption. Because of the limitations inherent in this measure of consumption, I explore the robustness of the results to alternate consumption measures throughout the analysis.

4 Measuring the Covariance between MPCs and Earnings Elasticities

I estimate the covariance between an individual's MPC and the sensitivity of individual earnings to aggregate shocks in two steps. In Section 4.1, I present estimates of MPCs by worker characteristic using the PSID, where there is both consumption and earnings data. In Section 4.2, I impute these MPCs in the LEHD using overlapping information on worker demographics and uncover an estimate of this key

¹⁰ I start with the source and SEO samples of the PSID from 1968 to 2015. I drop any observations that do not have two lags, used to define previous income, and one lead, used to define current income. As in the LEHD analysis, I also drop those individuals younger than age 25 or older than age 62. I also drop observations with missing food consumption or missing income. Lastly, I drop observations where the two-year change in either food consumption or income is more than fourfold. These restrictions on outliers are similar to those in Hendren (2017), who excludes individuals with more than a threefold change in food consumption, and Gruber (1997), who excludes observations with a greater than 1.1 log change in food consumption.

covariance.

4.1 Estimating marginal propensities to consume

I begin by estimating the MPCs for workers with different characteristics using the panel structure of the PSID. My estimation of MPCs borrows from a large literature that explores the response of individual consumption to income changes. A consistent finding within this large and heterogeneous literature is that households exhibit high MPCs even out of transitory income shocks and that the magnitude of these responses differs across the population.¹¹

I build most directly on a line of research beginning with Gruber (1997), who examines the consumption drop upon unemployment. Using the panel structure of the PSID, I estimate

$$\Delta C_{t,i} = \sum_x (\beta_x \Delta E_{t,i} \times x_{t-1,i} + \alpha_x x_{t-1,i}) + \delta_{t,s} + \epsilon_{t,i} \quad (5)$$

where $C_{t,i}$ is total household consumption of individual i at time t , imputed from the CEX as explained in Section 3.2, $E_{t,i}$ is labor earnings of individual i , $\delta_{t,s}$ are state-by-year fixed effects, which capture any variation that is common to all individuals within a state and year, and $x_{t,i}$ is a characteristic of the individual.¹² In order to include data from 1997 to 2015, when the PSID includes observations only for every other year, I consider two-year changes in both income and consumption. Using the estimated β_x , I can calculate the MPC for each individual as

$$\widehat{MPC}_{i,t} = \sum_x \widehat{\beta}_x x_{i,t}$$

Since many factors could simultaneously move income and consumption at the individual level, I identify the causal relationship between income and consumption in Equation 5 using an income shock as an instrument for $\Delta E_{t,i}$. In a general class of models, the MPC of the individual is a function of the

¹¹ In particular, a series of recent papers estimate that upon receiving tax rebates, workers, on average, spend more than half of the windfall within two quarters but that individuals with few financial resources spend more than those with more cash on hand (Johnson et al. (2006), Parker et al. (2013), Misra and Surico (2014)). See also Fagereng et al. (2021), Gelman (2020), Gross et al. (2020), Jappelli and Pistaferri (2014), Kaplan and Violante (2014), and Jappelli and Pistaferri (2010) for a comprehensive survey.

¹² In baseline specifications, x includes: five approximately equally-sized lagged earnings bins ($< \$22,000$, $\$22,000 - \$35,000$, $\$35,000 - \$48,000$, $\$48,000 - \$65,000$ and $> \$65,000$), a quadratic in age, female and black dummies, black interacted with age, and female interacted with black. See Appendix A.3 for additional details on sample restrictions. I include both household heads and spouses separately in the regression. Birinci (2019) documents using the PSID that the spousal response to the displacement of the family head is minimal, and thus the potential within-household spillovers are unlikely to meaningfully affect these estimates.

type of income shock. My baseline estimates are identified using unemployment as the shock to income, which is a large and persistent shock (Gruber (1997), Hendren (2017), Jacobson et al. (1993)). Specifically, I estimate Equation 5 for the set of workers employed in $t - 2$ and instrument $\sum_x (\beta_x \Delta E_{t,i} \times x_{t-1,i})$ with $\sum_x (\beta_x u_{t,i} \times x_{t-1,i})$, where $u_{t,i}$ is an indicator for whether the worker reports being unemployed at the time of the survey in year t . The coefficients on the reduced form of this regression capture the differential consumption change for different demographic groups for those workers who become unemployed relative to similar workers who did not lose their job over this time period. The first stage rescales this consumption drop for each demographic group by the income loss for the unemployed in that group, measured relative to those workers who remained employed. While it is possible that unemployment affects consumption directly by changing shopping behaviors (Kaplan and Menzio (2016)) or health (Sullivan and von Wachter (2009)), these behaviours would not violate the exclusion restriction if they are downstream of the large income shock.¹³ A key advantage of using unemployment is that it is a large, well-powered shock that captures the income variation that is most relevant for understanding recessions. Indeed, if all labor income shocks stemming from aggregate fluctuations have the same persistence as unemployment, then this is exactly the correct MPC to capture how the structure of the labor market affects the propagation of shocks. An MPC that is estimated using a purely transitory shock, such as a tax rebate, would fail to capture the consumption response to the type of income shocks that workers typically experience over the business cycle.¹⁴ In the online appendix, I further explore the importance of this shock persistence in a structural dynamic model and show that when combined with the 2-period multiplier in Section 2, these MPCs provide a good approximation for the dynamic multiplier in response to a persistent aggregate shock.

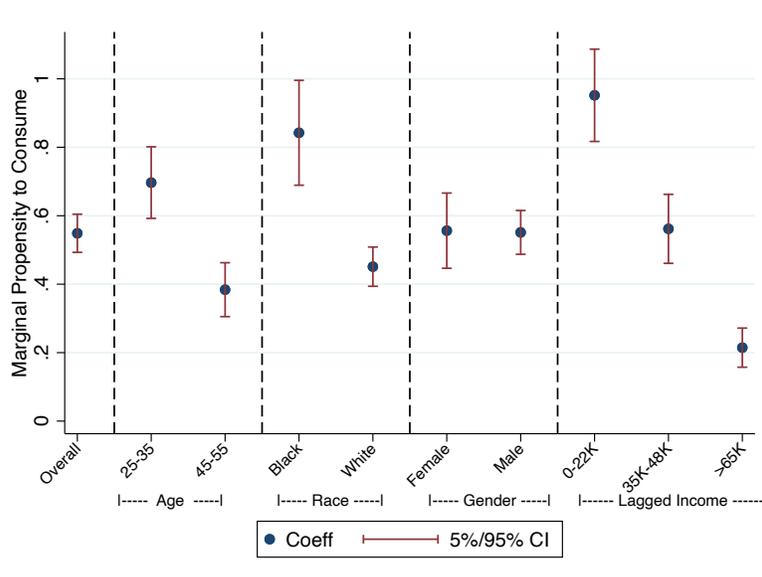
4.1.1 Description of MPC estimates

Before exploring the full distribution of MPCs that result from Equation 5, Figure 2 shows the MPC heterogeneity with bivariate regressions for select subgroups for the set of covariates that I include in the full estimation. The farthest-left estimate shows that the average MPC in the estimation sample is slightly

¹³ MPC estimates look similar when I restrict to the set of households that report being unexpectedly unemployed, suggesting that differential anticipation of the unemployment shock across demographic groups is not an important source of bias. See Appendix Table A5.

¹⁴Nonetheless, I explore the covariance using this alternative identifying shock and find similar patterns. See Section 5.1.

Figure 2: Heterogeneity in Marginal Propensity to Consume Estimates



Notes: Each estimate represents a separate regression including only the stated demographic group. In all cases, consumption is measured using total consumption, imputed using the method in Blundell et al. (2008). Income is measured using individual labor income. The instrument for income changes is an indicator for the worker reporting being unemployed in the survey in year t . The sample includes the set of workers who were employed two years before the current year. The sample in the PSID excludes observations with more than a 400 percent change in food consumption or income over a given two-year period. Lagged income is measured as the average labor market earnings of the individual in $t - 2$ and $t - 3$. All regressions include year-by-state fixed effects and observations from 1985 to 2013. Appendix Figure A4 shows the reduced form and first stage for these regressions. Standard errors are clustered at the individual level.

more than 0.5.¹⁵ Importantly, the coefficients to the right show that there is substantial variation in MPCs around this average – young, black, and low-income workers, on average, have a larger consumption drop per dollar of lost labor income. The lagged income measure along which I allow MPCs to vary is the average earnings of workers in years $t - 2$ and $t - 3$, which is meant to capture differences in permanent income across workers stemming from differences in education levels or other skills.¹⁶ In Figure 2, women and men have similar MPCs, on average, but women also earn less than men, and Appendix Table A5 shows that once I control for differences in earnings levels, women have lower MPCs than men. Putting this all together, Appendix Figure A3 shows the full distribution of MPCs in the PSID that result from Equation 5. There is a substantial amount of variation, with a large mass between 0 and 1, and a small number of

¹⁵ This average MPC is similar to others in the literature using similar variation. For example, using the Nielsen Consumer Panel, McKee and Verner (2015) estimate an MPC out of unemployment insurance benefits of between 0.6 and 0.9, and Jappelli and Pistaferri (2014) use survey data in Italy and find an average MPC out of transitory income of 0.48. See Jappelli and Pistaferri (2010) for a comprehensive review and Section 4.1.3 for an extended discussion.

¹⁶ I use the average earnings over the previous two years to balance capturing a more permanent measure of earnings capacity against a loss of sample that comes with the more stringent within-individual panel. However, I find that the patterns are similar when using income either lagged further, averaged over longer intervals, or fixed at a given age.

estimates above 1.

While I allow for heterogeneity in worker MPCs along only four demographic dimensions (race, age, gender, and earnings history), it is not necessary for the MPCs to be a direct function of those specific characteristics. Rather, it is likely the case that these characteristics are correlated with other underlying economic circumstances that directly affect MPCs. Specifically, in models of precautionary savings or credit constraints, a key source of heterogeneity in MPCs is the individual's cash on hand, and indeed, several studies find heterogeneity in MPCs along this margin (Jappelli and Pistaferri (2014), Kaplan et al. (2014), Parker et al. (2013), Fagereng et al. (2021)). The demographic patterns in Figure 2 are consistent with that source of heterogeneity, as differences in liquid assets across demographics alone explain about 20% of the variation in MPCs in Figure A3. The remaining heterogeneity likely reflects differences in risk preferences, differences in discount rates across demographic groups, or differences in access to credit.¹⁷ It is also possible that heterogeneity in $\hat{\beta}_x$ reflects differences in marriage rates across demographic groups, as married workers have lower MPCs due to risk-sharing within the household. Lastly, while heterogeneity in $\hat{\beta}_x$ could also reflect differences in the persistence of the unemployment shock across individuals, this does not seem to be first order in driving the cross-demographic patterns. While I find in the PSID that there is a positive correlation between the duration of unemployment and worker MPCs, I show in Appendix Figure A5 that MPCs estimated using tax rebates, which are clearly transitory for all groups, are also equally correlated with unemployment duration, suggesting that unemployment duration is correlated with other characteristics that drive MPC heterogeneity as opposed to heterogeneity in duration directly driving MPCs.¹⁸

¹⁷ There is increasing evidence suggesting that preference heterogeneity is important in explaining MPC heterogeneity in the population. In the context of stimulus payments, Parker (2017) finds that past income is highly predictive of household responses and that the patterns do not support identical households cycling in and out of high and low response states. Similarly, Gelman (2020) explores the degree to which MPC heterogeneity can be explained by persistent characteristics across individuals versus on transitory shocks and finds that about half of the variation in MPCs is driven by persistent characteristics. Finally, Aguiar et al. (2020) explore additional consumption patterns within the PSID to demonstrate that preferences play a prominent role in explaining differences in MPCs across households. Together, these findings suggest that some of both the between-demographic and within-demographic variation in MPCs stems from differences in preferences across households.

¹⁸ In addition, the calibrated dynamic model in the online appendix implies that heterogeneous unemployment durations are not the key driver of MPC heterogeneity in the model. Despite the model being calibrated to match heterogeneous unemployment durations across groups, the purely transitory MPCs in the model are nearly as heterogeneous across demographic groups as the unemployment-based MPCs, suggesting that the differences in unemployment duration are not the key drivers of the MPC heterogeneity in the model.

4.1.2 Assumptions for MPC imputation

While the above method for estimating MPCs closely follows existing methods in the literature, my subsequent imputation of these MPCs in the LEHD necessitates several important additional assumptions about the stability of the MPC estimates that warrant further discussion. It is important to note that these assumptions are not unique to my approach of combining the LEHD and the PSID, but will be needed to estimate the covariance term for *any* dataset in which MPCs are not directly observed.

First, in imposing the assumption that MPCs only vary by worker demographics, I assume that individual MPCs out of unemployment are similar to the MPCs out of business cycle income shocks of different signs and magnitude. This would imply both that the other business cycle shocks have a similar persistence to the unemployment shock and that the MPC does not depend on the magnitude of the shock. In a standard model in which agents maximize their expected utility subject to an intertemporal budget constraint, an individual's MPC depends on the persistence of the shock, not on the magnitude or sign. However, with liquidity constraints, the MPC out of small shocks may depend on the size of the shock as well as the sign (Kaplan and Violante (2014)). I explore the importance of the identifying shock by comparing MPC estimates that result from business cycle shocks of differing size and sign. Appendix Figure A6 shows that the average MPC and cross-demographic patterns are similar, although noisier, across these various identifying labor market shocks.

A second key stability assumption embedded in the imputation of MPCs in the LEHD is that conditional on demographics, the consumption response is constant over the business cycle. Existing empirical evidence on this cyclical stability of MPCs is scarce – Gross et al. (2020) find that the MPC out of liquidity is higher in recessions, but a calibration by Carroll et al. (2017) finds that MPCs are roughly constant over time. I explore this in my setting by adding an interaction of changes in income with the state unemployment rate, thereby allowing the MPC by demographic to vary over the business cycle, and find that differences, both on average and for each demographic group, are statistically and economically small.¹⁹

Third, I impose that at the individual level, the marginal propensity to consume is a function only of the characteristics that I include in x (i.e., age, earnings history, gender, and race). While this is obviously an approximation, this assumption would only upward bias my estimates of the covariance if there was sorting across jobs within a demographic group such that it was precisely the higher-MPC workers

¹⁹ See Appendix Figure A7 for a graph showing the estimates of MPCs at different points in the business cycle.

within the group who are at cyclically insensitive jobs. While my data do not allow me to fully address this, I explore the sensitivity of my MPC estimates in the PSID to including job-level characteristics that capture a dimension on which workers may sort. If sorting across jobs of different characteristics were important in explaining MPC heterogeneity within demographic group, these terms should have additional explanatory power. Appendix Table A6 shows that no additional job-level variables or regional controls meaningfully change the MPC estimates. While not ruling it out definitively, these estimates do suggest that this cross-job sorting within demographic group is small and unlikely to be a meaningful source of bias in my estimates.²⁰

4.1.3 Alternate MPC estimates

A key limitation of the PSID is that it includes only a restricted measure of consumption for most of the sample period, and therefore I must rely on an imputed measure of consumption. To assess the importance of this imputation for my estimates, I compare my baseline estimates to several alternatives. Column 1 of Table 1 reports the the baseline estimates from Section 4.1, which imputes total consumption within the PSID using data from the CEX following Blundell et al. (2008). In order to facilitate a comparison with other estimates in the literature, I split households into different groups based on their age, income, and place in the asset distribution, all defined in the pre-period ($t - 2$). Column 2 shows the same CEX-based imputation but for nondurable consumption only. The level of the MPCs fall, unsurprisingly, but the cross-demographic patterns remain consistent.

In columns 3 and 4, I compare these estimates to alternate MPC estimates constructed entirely within the PSID. One alternative is to use only food consumption, which is directly reported in the PSID over the full time series and does not require any imputation. A second alternative is to use the expanded consumption measures, which include spending on healthcare, education, childcare, transportation, and utilities, that are reported directly in the PSID beginning in 1999. Since limiting to the later years of the PSID severely limits the sample, I follow Attanasio and Pistaferri (2014) and impute consumption for households within the PSID for the full sample using the estimated relationship between expanded consumption and demographics/food consumption in the later years of sample (See Appendix A.2 for de-

²⁰I also explore the covariance using directly reported MPCs in the Italian Survey of Income and Wealth and find that the results look similar, further suggesting that unobserved heterogeneity is unlikely to be a major source of bias.

tails).²¹ Looking across columns 3 and 4, it is clear that the consumption measure affects the level of the MPCs, but the cross-demographic patterns are, again, reassuringly similar.

The final column of Table 1 shows comparable estimates from Ganong and Noel (2019), who use bank account data from the JP Morgan Chase Institute (JPMCI) to estimate the drop in nondurable consumption per dollar of lost income upon unemployment. Within this dataset, they are able to measure spending from debit and credit cards and cash withdrawals for 27 million households and identify unemployment using the receipt of UI payments. Comparing column 2 and 5, we see that the average level of the nondurable MPC estimates obtained from the CEX/PSID are very similar on average to those estimated in the JPMCI data. Moreover, the cross-sectional patterns are also very similar, with younger and poorer workers having substantially higher MPCs. The PSID/CEX estimates show somewhat more heterogeneity with income, possibly the result of the fact that the JPMCI data only includes those households with checking accounts and the unemployed who are eligible for and receive UI benefits. The consistency of these estimates across various measures of consumption demonstrates that the MPC patterns are unlikely to be an artifact of the particular measure of consumption used in the baseline estimates.

4.2 Heterogeneity in worker exposure

Using the baseline estimates of an individual’s MPC, I move to the LEHD and estimate the degree to which the earnings of workers of different MPCs are differentially sensitive to aggregate shocks. To do so, I estimate the following equation:

$$\Delta E_{i,t} = \alpha_1 \widehat{MPC}_{i,t-1} + \alpha_2 \widehat{MPC}_{i,t-1} \times \Delta \log Y_t + \delta_t + \epsilon_{i,t} \quad (6)$$

where $E_{i,t}$ is a measure of the individual’s fourth-quarter earnings; $\widehat{MPC}_{i,t-1}$ is the imputed MPC of individual i based on characteristics in $t - 1$; δ_t are year-fixed effects, soaking up any variation in earnings that is common across individuals in a given year; and $\Delta \log Y_t$ is the annual change in the log of national real GDP.²² The sample includes the set of all workers employed in year $t - 1$. The coefficient of interest

²¹ While Table 1 only shows coarse cross-demographic patterns, I can also estimate Equation 5 for each alternate measure of consumption in Columns 1 through 4. I find that the resulting MPCs are highly correlated with the baseline estimates (0.91 for food-only, 0.99 for nondurables, and 0.9 for PSID-based imputation).

²² For computational reasons, I estimate Equation 6 on a 5 percent random subsample of the data. Column 1 of Table 2 shows similar estimates with an aggregated estimation of Equation 6 that utilizes the full sample. I restrict attention to the set of employed workers for several reasons. First, the LEHD is a data set of employment and thus does not have complete coverage of the unemployed. I also earnings weight the regressions; thus, including the unemployed would necessitate an alternate weighting

Table 1: Comparing MPCs using Alternate Consumption Measures

	Total Consumption	Nondurable Consumption	Food Consumption	PSID Consumption	Ganong & Noel (2016)
Overall	0.55 (0.03)	0.33 (0.02)	0.048 (0.005)	0.23 (0.01)	0.40 (0.011)
Age < Median	0.61 (0.06)	0.39 (0.04)	0.057 (0.008)	0.23 (0.02)	0.467 (0.014)
Income < Median	0.87 (0.06)	0.54 (0.04)	0.068 (0.010)	0.32 (0.02)	0.467 (0.016)
Wealth in Bottom Quintile	0.781 (0.09)	0.47 (0.07)	0.045 (0.022)	0.37 (0.04)	0.51 (0.022)
Wealth in Top Quintile	0.28 (0.07)	0.14 (0.04)	0.006 (0.012)	0.15 (0.05)	0.106 (0.030)

Notes: Estimates from Ganong and Noel (2016) can be found in their Appendix Table 5. Consistent with the definition in Ganong and Noel (2016), each covariate in all columns is defined in the initial period ($t - 2$). In the PSID, income refers to individual labor income. In the PSID, wealth is defined as the sum of assets less the value of debt plus the value of home equity. Assets include farm or business worth, checking or savings accounts, real estate assets, stocks, vehicles, and individual retirement accounts. Debts include business debt, real estate debt, credit cards, student loans, medical debt, legal debt, and family loans. Wealth is defined in 1984, 1989, and every year from 1999 to 2015. Nondurable consumption is defined to include spending on food, alcohol, tobacco, utilities, transportation, personal care personal, apparel, shelter, entertainment and miscellaneous categories. Nondurable spending is imputed using the CEX using the same procedure as for total consumption. Food consumption is observed directly within the PSID. In column 4, PSID consumption is imputed for the extended PSID sample following Attanasio and Pistaferri (2014). All regressions in Columns 1 and 2 include data from 1985-2013 and are estimated using 2-year changes on the set of workers employed in $t - 2$. All regressions in column 3 and 4 include data from 1971-2013 and are estimated using 2-year changes on the set of workers employed in $t - 2$. Standard errors in columns 1 through 4 clustered at the individual level are included in parentheses.

is α_2 , which captures the degree to which the earnings of workers of different MPCs are differentially sensitive to movements in aggregate GDP.

I explore the different dimensions of earnings cyclicalities with various specifications of the outcome variable $E_{i,t}$. I capture the intensive margin of earnings elasticity by using the log of earnings, thus restricting the sample to the set of individuals who remain employed across years, and I capture the extensive margin of employment using an indicator for whether the individual is employed in time t . Lastly, I combine the intensive and extensive margin into one estimate using $\Delta E_{i,t} = \frac{E_{i,t} - E_{i,t-1}}{.5 \times E_{i,t} + .5 \times E_{i,t-1}}$. This transformation defines and bounds the earnings losses of those who lose their job between periods, thus providing an estimate for overall earnings elasticities. Importantly, each individual in the regression is weighted by strategy. Finally, a large fraction of earnings is earned by the employed rather than new hires. I explore the importance of the unemployed using the PSID and find that it has very little effect on the estimates. See Appendix A.5.

their share of overall earnings. What matters for the economy as a whole is not the differential elasticity of individuals but the differential elasticity of *dollars* earned in the economy.

In order to account for additional noise stemming from the imputed MPCs, I estimate Equation 6 using multiple imputation techniques from Rubin (1987). There are two imputations underlying Equation 6 – an imputation of total consumption in the PSID from the CEX, and an imputation of the MPC in the LEHD from the PSID. I implement the multiple imputation estimator by taking 100 draws of both the underlying CEX and PSID data. This produces 100 estimates of Equation 5, which estimates the MPC for each demographic group x . I then estimate Equation 6 for each imputation, which results in 100 estimates of α_1 and α_2 . Lastly, I combine these estimates following Rubin (1987), where the point estimate is the average across imputation draws, and the variance of the estimate is the combination of the average within-draw variance and the between-draw variance. In estimating Equation 6 for each draw, I twoway cluster standard errors by individual and year.

Before going directly to estimates of α_2 from Equation 6, recall Figure 1 from the introduction to demonstrate the variation underlying this key relationship. Each point represents a demographic group aggregated to the level of heterogeneity in the MPC imputation (for example, white men, ages 25 to 35 with lagged earnings between \$22,000 and \$35,000) and shows the relationship between group MPCs and earnings elasticities to GDP. This figure clearly shows that there is a tight, almost linear relationship between these two variables – higher-MPC demographic groups are much more exposed to recessions.²³

The slope of the fitted-value line in Figure 1 provides the estimate of α_2 from Equation 6 and is shown in column 1 of Table 2. Columns 2 through 4 move back to the individual level and explore the various dimensions of earnings heterogeneity. Column 2 shows the overall estimate, which is decomposed into the intensive and extensive margin of earnings in columns 3 and 4. Demographic groups that are cyclically sensitive on the intensive margin also tend to be cyclically sensitive on the extensive margin of employment, and thus, there is a strong positive relationship between the sensitivity of a worker’s income to GDP and a worker’s MPC along both the intensive and extensive margins. Both margins contribute similarly to the overall heterogeneity in exposure.²⁴ I construct an estimate of the covariance between earnings elas-

²³ In Appendix A.4, I replicate the findings in Guvenen et al. (2017) that the earnings of the very high and very low earnings are most exposed to GDP, as well as discuss how these patterns affect the relationship between MPCs and sensitivity to GDP.

²⁴ The extensive margin of adjustment within the LEHD includes shifts in labor market participation as well as true unemployment spells. However, I find in the PSID that high-MPC households are more likely to become unemployed during recessions (See Table A12).

Table 2: Earnings Elasticity to GDP by a Worker’s Marginal Propensity to Consume

	(1)	(2)	(3)	(4)	(5)	(6)
	Aggregated	Individual-level Specifications			Added Controls	
		Overall	Intensive Margin	Extensive Margin	Industry FE	Firm FE
$MPC_{i,t-1}$	-0.053 (0.010)	-0.229 (0.049)	-0.056 (0.008)	-0.103 (0.023)	-0.206 (0.053)	-0.149 (0.061)
$MPC_{i,t-1} * \Delta Y_t$	1.044 (0.167)	1.228 (0.383)	0.538 (0.141)	0.460 (0.176)	1.087 (0.387)	1.065 (0.445)
No. Observations	2,496	29.2 M	25.5 M	29.2 M	29.2 M	23.8 M
R-Squared	0.532	0.009	0.002	0.010	0.019	0.263
$Var(MPC_i)$	0.083	0.083	0.080	0.083	0.083	0.082
$Cov(MPC_i, \gamma_i)$	0.084	0.099	0.042	0.037	0.088	0.085
Fixed Effects						
Year	✓	✓	✓	✓		
Year by Industry					✓	
Year by Firm						✓

Notes: All observations are weighed by the individual’s earnings in $t - 1$. Earnings in each regression are defined as total quarterly earnings for each individual in the fourth-quarter of the year. In all columns, regressions are implemented using a multiple imputation approach explained in detail in Appendix A.4. The outcome variable in Columns 1, 2, 5 and 6 is $\Delta E_{i,t} = \frac{E_{i,t} - E_{i,t-1}}{.5 \times E_{i,t} + .5 \times E_{i,t-1}}$. The outcome variable in Column 2 is the change in the log of fourth-quarter earnings, and the outcome variable in Column 3 is an indicator for being employed in time t . Column 1 aggregates from the individual level to MPC bins (MPC rounded to the nearest 0.01) and runs a regression on the aggregated sample. All other columns are estimated on a 5 percent random subsample. $Var(MPC)$ is the earnings-weighted variance of MPCs in the sample, averaged across all 100 draws of the MPC estimates. The covariance reported for each specification is calculated as the regression coefficient on $MPC_{i,t-1} * \Delta GDP_t$ times $Var(MPC)$, averaged across all 100 draws of MPC estimates.

ticities and MPCs by multiplying α_2 by the variance of MPC. Using the baseline estimate in Column 2, I estimate a covariance of 0.10.

The final two columns of Table 2 show the estimated relationship with additional controls. One potential explanation for the observed covariance is that worker skills covary with MPCs and determine income exposure through cross-firm sorting. However, the results in Columns 5 and 6 suggest this is only a small part of the story. Column 5 adds industry-by-year fixed effects, thus isolating the correlation between MPCs and earnings elasticities within industries, and Column 6 includes firm-by-year fixed effects, isolating the correlation that remains within the firm. Even with firm-by-year fixed effects, the coefficient and associated variance only drops by about 11 percent, suggesting that a large majority of the aggregate relationship exists within firms.

This positive relationship between the cyclical sensitivity of earnings and the MPC of the worker is also very robust across empirical specifications. Appendix Table A7 shows that the estimated covariance

is robust to several decisions that define the baseline specification. Specifically, the covariance is slightly larger when using alternate functional form of the earnings outcome variable, slightly smaller when using state-level GDP as a cyclical indicator rather than aggregate GDP, and similar when using either annual income or a balanced panel of states. Looking across all of the columns in Table A9, the estimate of the covariance between MPCs and earnings elasticities ranges from 0.08 to 0.2.²⁵

Although the source of this relationship between MPCs and earnings elasticities is not immediately relevant for understanding its contribution to macroeconomic stability, it is an interesting object in its own right. The finding that the earnings of high-MPC workers are more exposed to aggregate conditions is at odds with a simple model of firm-worker risk-sharing, which would predict that workers for whom fluctuations in income are costly would sacrifice part of their expected earnings to enter into contracts in which their wages are less sensitive to aggregate demand fluctuations (Bailey (1974)). These patterns are not, at face value, inconsistent with the idea that workers sort across jobs according to their risk preferences, but that would require a strong negative correlation between risk aversion and worker MPCs (Schulhofer-Wohl (2011)). Additionally, the relative unimportance of firms suggest that this pattern is also consistent with models where workers' exposure to aggregate shocks and their MPCs are explicitly linked. For example, this covariance could result from a job ladder model as in Jarosch (2014), wherein as workers climb the job ladder, they sort into higher-paying and more-secure jobs. This model of worker flows would result in a pattern where within the firm, a worker with higher lagged earnings, both because she is in a higher-paying job further up the ladder and because she has not recently experienced an unemployment spell, is less exposed to shocks than the lower-earning recently-hired worker.

5 Matching Multiplier Estimates

The previous section measured the covariance between MPCs and earnings elasticities. In this section, I use the simplified framework in Section 2 to demonstrate that the magnitude of this covariance is large enough to have meaningful effects on aggregate consumption and output. Recall from Section 2 that the matching multiplier is defined as the difference in the Keynesian multiplier with the empirical incidence of aggregate shocks and the multiplier in a benchmark case in which all workers faced the same earnings

²⁵While Table 2 looks at the average covariance over the business cycle, Appendix Table A9 shows estimates using state-level GDP that separates expansions and contractions. While the estimates are noisier, the patterns suggest that the covariance is positive in both cases but is larger in recessions than expansions.

elasticity. In the framework used to derive Equation 4, all output is earned by workers in the form of wages. However, this is not true empirically, as workers earn income from other sources (e.g., capital and profits). Therefore, in the empirical estimates of the matching multiplier presented in Table 3, I make adjustments for 1) the MPC out of non-labor income (MPC_{nl}); 2) the share of overall output going to workers in the form of wages (α_l); and 3) the overall elasticity of income with respect to output ($\bar{\gamma}$). Specifically, the modified formulas for actual and benchmark MPCs with these empirical adjustments

$$MPC^b = \alpha_l \bar{\gamma} \overline{MPC} + (1 - \alpha_l \bar{\gamma}) MPC_{nl} \quad (7)$$

$$MPC^a = \alpha_l (\bar{\gamma} \overline{MPC} + Cov(MPC_i, \gamma_i)) + (1 - \alpha_l \bar{\gamma}) MPC_{nl} \quad (8)$$

where \overline{MPC} is the earnings-weighted average MPC out of labor income (See Appendix A.7 for the derivation). I estimate that the aggregate elasticity of labor earnings to GDP is 1.03.²⁶ I assume that the labor share is two-thirds, and I set the MPC out of nonlabor income to be 0.16, which I calculate using estimates from the literature on the MPC out of dividends, capital gains, and interest income.²⁷

Column 1 of Table 3 shows the baseline estimates with accompanying confidence intervals, which come from the estimation in Table 2. Panel A simply replicates the regression coefficient of interest from Table 2 and shows that the baseline regression implies a covariance between MPCs and earnings elasticities that is just shy of 0.1, with the 95% confidence interval spanning 0.04 to 0.16. Panel B demonstrates what this covariance means for the aggregate MPC. The top row of panel B shows the benchmark MPC. As in Equation 7, this is defined as the weighted average of the benchmark MPC out of labor income and the assumed MPC out of non-labor income, where the weight is the labor share times the average earnings elasticity. For labor income, the benchmark MPC is the earnings-weighted average individual MPC. This estimate captures the aggregate MPC in the benchmark economy in which all workers face the elasticity

²⁶Specifically, I regress the log change in real wages and salaries from BEA Table 2.2A and 2.2B against the change in real GDP from BEA Table 1.1.4 and 1.1.5 using data from 1947-2011. This captures the elasticity of the average dollar of labor earnings to aggregate output.

²⁷Specifically, in 2011, non-labor, non-transfer income was 33 percent proprietors income, 8 percent rental income, 15 percent dividends and 44 percent interest (See NIPA Table 2.1). Since Smith et al. (2019) demonstrate that three quarters of business profits can be classified as human capital income accruing primarily to top earners, I take the MPC out of proprietor's income to be the MPC that I estimate for labor income for the highest income group in Figure 2, which is 0.2. I take the MPC out of rental income to be 0.028, which is the MPC out of stock market income estimated in Chodorow-Reich et al. (2021). Finally, I use Swedish estimates from Di Maggio et al. (2020) and calibrate an average MPC of 0.49 for dividends and 0.04 for interest. Combining these estimates, I arrive at an estimate for the MPC out of non-labor income of 0.158.

of the average dollar. I estimate this benchmark MPC to be 0.33.

Row 7 in Panel B shows the actual MPC, which is defined in Equation 8 as the benchmark MPC plus the labor share times the estimated covariance. Using the baseline estimates, I find that the actual MPC is 0.07 higher than this benchmark, meaning that heterogeneity in worker exposure to aggregate fluctuations increases the aggregate MPC by 20 percent. Finally, Panel C of Table 3 translates the difference in the aggregate MPCs into a difference in the multiplier using the simple $\frac{1}{1-MPC}$ formula. Row 9 and 10 show that these estimates imply that the multiplier would be 17 percentage points smaller if the covariance between earnings elasticities and MPCs were zero.

The other columns of Table 3 show the estimated increase in the aggregate MPC and multiplier implied by the MPC estimates that use alternate measures of consumption from Table 1. Since each of these measures of consumption accounts for only a subset of consumption, in order to facilitate a direct comparison with the baseline estimates, I divide the estimates by the share of total consumption included in the MPC estimate, which is reported in row 5. This normalization assumes that the MPC out of unobserved consumption is the same as the MPC out of the observed component of consumption (i.e. the MPC for non-food expenditures is the same as the MPC for food, etc.).²⁸ Across columns 2 through 4, I find that benchmark aggregate MPC is slightly lower, but that the aggregate MPC increases by at least 0.04, or between 16 and 22 percent, for all consumption measures.

5.1 Alternate Matching Multiplier Estimates

My primary estimate uses the LEHD because it provides a large sample with highly precise income measures and detailed information on firms and geographies. In this section, I explore alternate estimates of the covariance term. Each of these alternate estimation strategies complements and extends the analysis in a different way, and they all support the main finding that there is positive amplification coming from this cyclical inequality channel.

Specifically, Figure 3 summarizes the headline results for three additional estimates, the details of which can be found in Appendix A.5. The light blue bars show the estimated covariance between MPCs

²⁸ While the assumption that the MPC out of unobserved income is similar to the MPC out of observed income is a reasonable benchmark, previous literature has found that the MPC for durable goods is in fact larger than the MPC for nondurable goods (See for example Parker et al. (2013), where the MPC for nondurable spending is less than a third of the MPC for total spending, and Lewis et al. (2019), where the average MPC for nondurable is 0.19 and for durable spending is 0.26. Since Columns 2 through 4 focus on nondurable spending, this implies that the MPC for unobserved spending is potentially higher than the MPC for observed spending. If that were the case, then these estimates would be even closer to the benchmark in Column 1.

Table 3: Benchmarking Amplification from the Matching Multiplier

	(1) Total Consumption	(2) Nondurable Consumption	(3) PSID Consumption	(4) Food Consumption
<u>Panel A: Covariance Calculation</u>				
(1) Coefficient on $MPC_{i,t-1} * \Delta GDP_t$	1.228 [0.477, 1.979]	1.755 [0.691, 2.819]	3.945 [1.495, 6.395]	12.47 [3.579, 21.361]
(2) Avg. MPC	0.420	0.340	0.307	0.207
(3) Var(MPC)	0.083	0.055	0.016	0.005
(4) $Cov(MPC_i, \gamma_i)$	0.099 [0.040, 0.164]	0.97 [0.038, 0.156]	0.063 [0.024, 0.102]	0.064 [0.018, 0.109]
(5) Consumption Share	1.000	0.735	0.550	0.140
<u>Panel B: Increase in Aggregate MPC</u>				
(6) Benchmark MPC (MPC^b)	0.338	0.283	0.261	0.192
(7) Actual MPC (MPC^a)	0.404 [0.364, 0.447]	0.331 [0.302, 0.360]	0.302 [0.276, 0.328]	0.235 [0.204, 0.264]
(8) Pct. Increase in MPC	19.6% [7.8%, 32.4%]	16.8% [6.6%, 27.1%]	16.0% [6.1%, 26.0%]	22.1% [6.3%, 37.9%]
<u>Panel C: Increase in Multiplier</u>				
(9) Benchmark Multiplier	1.510	1.395	1.352	1.237
(10) Actual Multiplier	1.678 [1.573, 1.810]	1.494 [1.432, 1.562]	1.433 [1.382, 1.489]	1.306 [1.256, 1.359]
(11) Pct. Increase in Multiplier	11.1% [4.2%, 19.8%]	7.1% [2.7%, 12.0%]	6.0% [2.2%, 10.1%]	5.5% [1.5%, 9.9%]

Notes: 95 percent confidence intervals are provided in brackets below the estimates and are calculated using the standard error on the regression coefficients in Table 2 and Table A8, all of which are calculated using multiple imputation techniques. In all columns, the average MPC is the earnings-weighted average MPC. In columns 2-4, I report the earnings-weighted average MPC and earnings-weighted variance of MPCs divided by the share of consumption. These shares are calculated within the PSID as the earnings-weighted share of consumption accounted for by each subcategory from 1990-2013. Full regressions for Panel A are reported in Table A8. $Var(MPC)$ is calculated as the earnings-weighted variance of MPC for each draw, then averaged across 100 imputation draws. In Panel B, MPC^b is defined as $\alpha_l \bar{\gamma} MPC + (1 - \alpha_l \bar{\gamma}) MPC_{nl}$ where $\alpha_l = \frac{2}{3}$, $\bar{\gamma} = 1.03$, and $MPC_{nl} = 0.158$. The actual MPC is defined as $MPC^b + \alpha_l * Cov(MPC_i, \gamma_i)$.

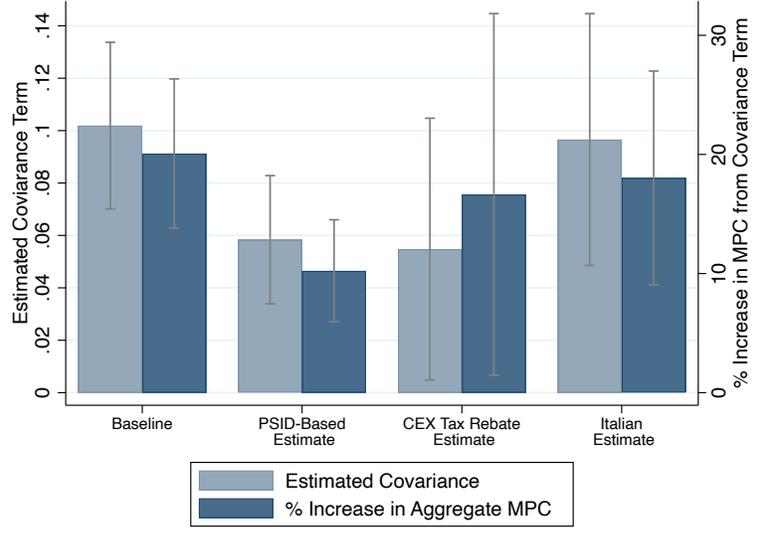
and earnings elasticities for each method and the navy bars show the percent increase in the aggregate MPC that comes from this covariance term (i.e. the equivalent of row 8 of Table 3). Gray ticks represent 1 standard deviation in each direction. The first set of bars reports the baseline estimates using the combination of the LEHD and PSID. The second set of bars shows the results of closely replicating the baseline analysis entirely within the PSID, restricting the sample to include data from 1997-2011 and including only the sample employed in $t - 2$. In order to calculate this moment, I estimate MPCs for each household in the PSID as described in Section 4.1 and then estimate Equation 6 within the PSID. The resulting covariance is smaller and substantially noisier than the baseline estimate, but it still implies a meaningful covari-

ance. While the PSID provides a noisier measure of income and a much smaller sample than the LEHD, an advantage to exploring the covariance within the PSID is that I am able to both explore additional dimensions of earning heterogeneity (e.g. distinguish between non-employment and unemployment) and utilize the longer time series to explore how the covariance has changed over time. In Appendix A.5, I show that households with higher MPCs are more likely to transition to unemployment during recessions, and that the covariance estimate is similar when including the government sector and unemployed workers, or when looking at total household income, rather than individual labor market earnings. I also find evidence within the PSID suggesting that the covariance has grown over time and is largest for recent recessions.

The third set of bars in Figure 3 shows the estimate of the covariance, again using the PSID to estimate earnings elasticities, but this time estimating marginal propensities within the CEX using tax rebates as the income shock. Following Johnson et al. (2006), I focus on nondurable spending, for which the MPCs are more tightly estimated. While shocks to labor market earnings identify the consumption response that is most applicable to the earnings inequality channel that is the focus of this paper, the covariance of labor market earnings with this alternate MPC is directly informative of this amplification in other heterogeneous agent models (Bilbiie (2020), Auclert (2019)). Using the CEX, I pool the 2001 and 2008 tax rebates and estimate MPCs for nondurable consumption by household demographics, then impute those MPCs in the PSID and estimate the elasticity of household labor market earnings by this alternate MPC measure. The covariance is lower since the MPC is defined for only nondurable spending, but the point estimate suggests that the percentage increase in the aggregate MPC is substantial. However, these estimates are substantially noisier than the baseline estimates, and the 95 percent confidence intervals cannot rule out a small negative covariance.

Lastly, the fourth set of bars show an estimate for the covariance between MPCs and earnings elasticities using the 1998-2018 waves of the Italian Survey on Household Income and Wealth (SHIW). A key advantage of this household-level survey is that respondents were explicitly asked what fraction of an unexpected windfall they would spend over the coming year in four waves of the sample period (see Jappelli and Pistaferri (2014) and Jappelli and Pistaferri (2020) for details). These self-reported MPC measures have been found to accurately capture average variation in MPCs and to be informative of MPC differences across households (Parker and Souleles (2019)). Importantly, these estimates are the *only* ones that do not

Figure 3: Alternate Estimates of the Covariance



Notes: Columns 2 and 3 use state GDP as the identifying shock in order to increase statistical power. Column 1 reproduces the baseline results in Column 1 of Table 3. Columns 2 and 3 use PSID data from 1997-2011 and Column 4 uses data from 1998-2018. For the PSID-based estimate in Column 2, I use the covariance estimate for overall individual labor earnings from the left panel of Figure A9. For the tax rebate estimates, I use household labor income, defined as the sum of labor income for the household head and spouse. The coefficient is from column 4 in Table A11. For all columns, I assume the average elasticity of labor income with respect to GDP is 1.03, which is the estimate used in Table 3. For the Italian estimate, I use household income and self-reported MPCs (see column 1 in Panel A of Table A10). Gray ticks represent plus or minus one standard deviation, calculated using the upper and lower bounds for the regression estimate on the interaction between MPCs and changes in log GDP. In columns 2 and 3, the standard errors are twoway clustered at the individual and year*state level. In column 4, standard errors are twoway clustered at the individual and year. In column 1, standard errors are calculated using multiple imputation as in Table 2. See Appendix A.5 for additional details on the estimates in Columns 2 through 4.

necessitate any imputation based on household characteristics. Figure 3 shows that the implied covariance between household income sensitivities and reported MPCs in Italy is very similar in magnitude to that in the baseline, amplifying the aggregate MPC by almost 20 percent. This is somewhat remarkable, given that the covariance between MPCs and earning elasticities is the result of the equilibrium distribution of workers across firms and could therefore vary across countries that differ in their labor market structure and degree of household insurance. These results therefore suggest that this amplification channel is not a unique feature of the United States.

Overall, while the baseline estimates provide the strongest evidence for this mechanism, and some of the other estimates are considerably noisier, the consistent direction and magnitude of the other estimates provides evidence that this result is not limited to a single dataset, approach, or even country.

6 Validating the Mechanism Using Local Labor Markets

The benchmark discussed in Section 5 shows that within the simplified framework in Section 2, the magnitude of the covariance between worker MPCs and earnings elasticities is large enough to have sizable effects on aggregate fluctuations. In this section, I empirically explore the importance of this covariance in determining the economy's response to shocks by utilizing the fine geographic detail in the LEHD. Under the assumption that demand in commuting zones is largely locally determined, the model in Section 2 would predict that output in areas with a higher measured covariance would be more sensitive to shocks than output in areas with a smaller covariance. If the mechanism were not important, whether because (a) consumption is not an amplification channel, (b) this covariance term does not matter, or (c) the assumptions underlying my empirical analysis do not hold, we should see no relationship between the local covariance and the size of local recessions.

I begin by separately estimating the earnings-weighted covariance between MPCs and earnings elasticities in each commuting zone.²⁹ My LEHD sample includes 270 commuting zones, and in each one of those areas, I separately estimate the covariance between MPCs and earnings elasticities. There is substantial variation in $\widehat{Cov}_c(MPC_i, \gamma_i)$ across local labor markets – the cross-commuting zone average covariance between worker MPCs and earnings elasticities is 0.076 and the standard deviation is 0.075. I also calculate the earnings-weighted average MPC in each commuting zone (\widehat{MPC}_c), which is the aggregate MPC in each area when the covariance between worker MPCs and earnings elasticities is zero.

Causally identifying the relationship between $\widehat{Cov}_c(MPC_i, \gamma_i)$ and local cyclicity is challenging, as the size of the covariance locally is likely correlated with many other features of the local labor market that also affect the cyclicity of the area. I address these identification challenges in two ways. First, I implement a case study of the Great Recession, where I estimate $\widehat{Cov}_c(MPC_i, \gamma_i)$ in a pre-period and use an array of controls for other local characteristics that would affect the cyclicity of the area and may be correlated with the local covariance. Second, I use the full sample period and test the additional prediction that the differential effects on employment across regions should be concentrated in nontradable industries, which are subject to local demand, rather than tradable industries, which are subject to a more national demand.

²⁹ I restrict attention to commuting zones where the 23 states in the subsample cover at least 90 percent of employment in that commuting zone, excluding an average of 12 percent of workers and earnings in each sample year.

First, I implement the case study of the Great Recession by estimating

$$\Delta_{2007-2010} \log L_{j,c} = \phi_1 \widehat{Cov}_c(MPC_i, \gamma_i) + \phi_2 \widehat{MPC}_c + X' \Phi + \delta_j + \epsilon_{j,c} \quad (9)$$

where $\Delta_{2007-2010} \log L_{j,c}$ is total change in log employment in 4-digit NAICS industry j in commuting zone c between 2007 and 2010, $\widehat{Cov}_c(MPC_i, \gamma_i)$ is the covariance in commuting zone c , \widehat{MPC}_c is the earnings-weighted average MPC in the area, X is a series of additional CZ-level controls, and δ_j are industry fixed effects, which allow for each industry to have a different employment trend in the Great Recession. The coefficient of interest is ϕ_1 , which the theory predicts is negative – all else equal, areas with a higher covariance should have seen larger employment declines in the Great Recession.

The left panel graphically illustrates the estimate of ϕ_1 . The relationship is strongly negative, and sectors in areas with a larger covariance experienced larger drops in employment in this period relative to the same sector in areas with smaller covariances. While the covariance does not depend on the employment cyclical of the area directly, the nature of the particular set of shocks in the sample period could jointly generate a big recession in the area and a high covariance. I explore this possibility by re-estimating $\widehat{Cov}_c(MPC_i, \gamma_i)$ in the pre-Great Recession period (i.e., 1995 to 2007) and using this as an instrument for the full-sample estimate of $\widehat{Cov}_c(MPC_i, \gamma_i)$.³⁰ The right panel plots the IV results, showing that the patterns are similar when using the pre-recession estimate, assuaging concerns that the simultaneous estimation of the two components is driving the results.

Second, I extend this case study by estimating the following extension of Equation 9

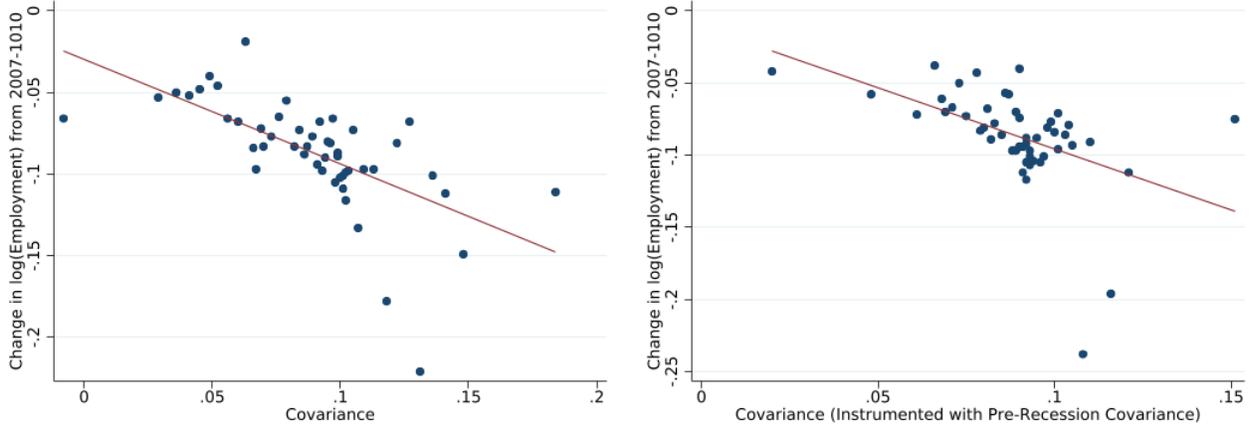
$$\Delta \log L_{j,c,t} = \omega_1 \widehat{Cov}_c(MPC_i, \gamma_i) \times \Delta \log Y_t + \omega_2 \widehat{MPC}_c \times \Delta \log Y_t + X' \Phi + \delta_{cj} + \delta_{jt} + \epsilon_{j,c,t} \quad (10)$$

where $\log Y_t$ is aggregate GDP, δ_{cj} are CZ-by-industry fixed effects and δ_{jt} are industry-by-year fixed effects. As in section 4.2, I estimate Equation 10 for 100 draws of the MPC estimates and report coefficients and standard errors in Table 4 calculated using Rubin's multiple imputation adjustment. The prediction of the theory is that ω_1 is positive – all else equal, employment in areas with a higher covariance is more sensitive to changes in aggregate output.

Column 1 of Table 4 first reports a specification that includes only the total MPC ($\widehat{Cov}_c + \widehat{MPC}_c$) and

³⁰While the estimates of pre-recession $\widehat{Cov}_c(MPC_i, \gamma_i)$ are somewhat noisier, the correlation with the full-sample $\widehat{Cov}_c(MPC_i, \gamma_i)$ is high at 0.33 and the first stage relationship is strong.

Figure 4: Employment in the Great Recession and the local $Cov(MPC_i, \gamma_i)$



Notes: Each point represents 2 percent of the sample, which includes 313 4-digit NAICS industries and 270 commuting zones. Each observation is weighted by its share of labor market earnings in 2007. Each plot includes industry fixed effects and controls for the earnings-weighted average MPC in the area, the average age and lagged earnings of the area, and the fraction of the commuting-zone that is female, black, and in the labor force. In the left panel, $\widehat{Cov}_c(MPC_i, \gamma_i)$ is calculated using data from 1995-2011. In the right panel, that covariance is instrumented with the covariance calculated using data from 1995-2007.

does not separately estimate the contribution of the covariance. I find that, as predicted, employment in areas with a higher overall aggregate MPC is more sensitive to business cycles. Column 2 then breaks apart the overall MPC into the two components that are the focus of this paper – \widehat{Cov}_c and \widehat{MPC}_c . In this specification, I include controls for local demographic variables, each on their own and interacted with aggregate GDP.³¹ I find that estimate of ω_1 is positive, statistically significant, and economically substantial. The magnitude of the coefficient in Column 2 implies that areas with the average covariance have an elasticity to aggregate output that is 0.44 percentage points higher than an area with a covariance of 0. Since the average elasticity of employment in an area to GDP is 0.80, the coefficient in Column 2 of Table 4 implies that the elasticity of an area with $\widehat{Cov}_c = 0$ is 56 percent lower than are area with an average covariance. Note that coefficient on the simple weighted-average MPC interacted with GDP is negative but statistically insignificant. Given the methodology for constructing MPCs, the average MPC of an area is strictly a function of the distribution of local demographics, and were I to be able to control flexibly enough for these demographics, this term would drop out of the regression entirely. Therefore, it is not surprising that after controlling for the demographics of the area, there is not enough variation left to identify the effect of the average MPC on local cyclicalities. The matching multiplier, however, is a function

³¹ Specifically, the demographic controls included are average worker age, the percentage of workers who are black, average lagged worker incomes, the fraction of workers who are female, and the fraction of the area that is employed.

Table 4: Employment sensitivity to output and the local $Cov(MPC_i, \gamma_i)$

	(1)	(2)	(3)	(4)	(5)
$(\widehat{Cov}_c + \widehat{MPC}_c) * \Delta \log Y_t$	4.542 (1.184)				
$\widehat{Cov}_c * \Delta \log Y_t$		5.873 (1.510)	2.682 (1.342)	4.160 (0.939)	0.069 (2.692)
$\widehat{MPC}_c * \Delta \log Y_t$		-0.590 (3.215)	-0.308 (2.459)	-0.615 (1.859)	0.980 (4.854)
Included Industries	All	All	T&N	N	T
Observations	487000	487000	130000	42000	88000
R-squared	0.345	0.345	0.318	0.340	0.281

Notes: In all columns, the dependent variable is the annual change in log employment in a four-digit NAICS within a commuting zone. All observations are weighted by the share of employment in $t - 1$. The row labeled “included industries” specifies whether the regression includes all industries (All), tradable industries (T), nontradable industries (N), or both (T&N). Tradable industries are those that show up in global trade data and non-tradable industries are retail- and restaurant-related (See Mian and Sufi (2014), Appendix Table A1). All regressions include demographic controls for average age and lagged earnings of the area; and the fraction of the CZ that is female, black, and in the labor force in $t - 1$, each included separately and interacted with $\Delta \log Y_t$. In all columns, coefficients and standard errors are calculated using multiple imputation techniques to account the imputation in the MPC estimates. See Appendix A.4 for additional details on this implementation. The average \widehat{Cov}_c across commuting zones is 0.076 and the average elasticity of employment in an area to GDP is 0.799. The number of observations is rounded to the nearest thousand.

not of the *level* of the MPCs but of the degree to which income changes are distributed across those groups, and therefore, these controls do not have the same effect on the covariance term.

Since the covariance affects the cyclicity of an area through its effect on local consumption, an additional prediction is that the relative employment effect of the local matching multiplier across regions should appear for nontradable industries, which are subject to local demand, rather than for tradable industries, which are subject to national demand. I explore this by estimating Equation 10 separately for tradable and nontradable industries. Following Mian and Sufi (2014), I define tradable industries as those in global trade data and nontradable as those related to retail and restaurants.

Columns 3 through 5 of Table 4 report the results. Column 3 restricts to only those industries classified as tradable or nontradable industries, respectively. Column 4 shows the relationship for employment in only nontradable industries, where the relationship is positive, large, and statistically significant. Conversely, Column 5 reports the relationship for employment in only tradable industries, showing that the relationship is small, negative, and statistically insignificant.³² These cross-industry patterns confirm the

³² Appendix Table A13 shows that the patterns are similar when adding an interaction between \widehat{Cov}_c , \widehat{MPC}_c , and an indicator for whether the industry is tradable to Equation 10 and are further robust to the inclusion of CZ-by-year fixed effects, which control for any employment movements common to industries within a commuting zone in a given year and identify only the degree to which the relative movement of tradable and nontradable employment co-varies with the matching multiplier, or

prediction that these cross-CZ patterns are driven by differences in nontradable industries, further validating the empirical relevance of this amplification mechanism.

7 Conclusion

This paper explores the link between inequality in the labor market and macroeconomic stability. I demonstrate that the aggregate MPC is higher when there is a positive covariance between worker MPCs and the elasticity of their earnings to aggregate output, a mechanism I call the matching multiplier. Empirically, it is precisely the high-MPC workers whose earnings are most exposed to recessions, and this relationship is large enough to have meaningful effects on the response of output to shocks. This moment directly informs the strength of the income heterogeneity channel in the heterogeneous agent new Keynesian models. Indeed, new quantitative work has incorporated this moment and continued to demonstrate its quantitative significance for amplifying shocks in HANK models (Slacalek et al. (2020), Alves et al. (2020)).

Uncovering the linkages between labor market inequality and the consumption multiplier has potentially important implications for macroeconomic stabilization policy. Indeed, policies can be made more effective in part by explicitly targeting this covariance between earnings heterogeneity and worker MPCs. For example, the government could consider this covariance when deciding how to target fiscal stimulus across industries or firms (Flynn et al. (2021)). While much of the covariance between MPCs and earnings sensitivities to GDP occurs within the firm, there is still some scope for targeting particular industries or firm types (e.g., young and small firms), where higher-MPC workers are more likely to be employed. Additionally, unemployment insurance that is targeted toward high-MPC workers could provide greater aggregate consumption stabilization benefits. Several countries, such as Germany, make unemployment insurance more generous for older workers, but the results in this paper suggest a rationale for making unemployment insurance benefits more generous for young workers, who have higher MPCs and more volatile earnings.

Lastly, these results may suggest another reason for policymakers to be alarmed by rising inequality in the economy. As wealth becomes more unequally distributed, MPCs in the population may become more dispersed, with a wide swath of consumption being greatly affected by aggregate shocks. A con-

controls for local financial conditions. Additionally, although noisier, results are similar when using a local Bartik-style shock rather than aggregate GDP as the demand shock. This alternative specification eliminates potential differences across commuting zones in their exposure to aggregate shocks, and thus isolates differences coming from local amplification channels.

current economic phenomenon of the past decade is that workers have become increasingly sorted across firms, occupations, and even types of employment contracts. These two economic forces could combine to further strengthen this mechanism and contribute to more economic instability in the future.

References

- Abowd, John M., Paul A. Lengeremann, and Kevin L. McKinney "The Measurement of Human Capital in the U.S. Economy," Technical Paper 2002-09, U.S. Census Bureau, 2003.
- Abowd, John M., Bryce E. Stephens, Lars Vilhuber, Fredrik Andersson, Kevin L. McKinney, Marc Roemer, and Simon Woodcock *The LEHD Infrastructure Files and the Creation of the Quarterly Workforce Indicators*, 149–230: University of Chicago Press, 2009.
- Aguiar, Mark A, Mark Bils, and Corina Boar "Who Are the Hand-to-Mouth?" Working Paper 26643, National Bureau of Economic Research, 2020.
- Alves, Felipe, Greg Kaplan, Benjamin Moll, and Giovanni L. Violante "A Further Look at the Propagation Mechanism of Monetary Policy Shocks in HANK," 2020.
- Attanasio, Orazio and Luigi Pistaferri "Consumption Inequality over the Last Half Century: Some Evidence Using the New PSID Consumption Measure," *American Economic Review*, 2014, 104 (5) , 122–26.
- Auclert, Adrien "Monetary Policy and the Redistribution Channel," *American Economic Review*, 2019, 109 (6) , 2333–67.
- Auclert, Adrien, Matthew Rognlie, and Ludwig Straub "The Intertemporal Keynesian Cross," Working Paper 25020, National Bureau of Economic Research, 2018.
- Bailey, Elizabeth E. "Innovation and regulation," *Journal of Public Economics*, 1974, 3 (3) , 285–295.
- Bilbiie, Florin "Limited Asset Markets Participation, Monetary Policy, and (Inverted) Aggregate Demand Logic," *Journal of Economic Theory*, 2008, 140 , 162–196.
- Bilbiie, Florin O. "The New Keynesian cross," *Journal of Monetary Economics*, 2020, 114 , 90–108.
- Bilbiie, Florin Ovidiu "Monetary policy and heterogeneity: An analytical framework," Available at SSRN 3106805, 2018.
- Birinci, Serdar "Spousal Labor Supply Response to Job Displacement and Implications for Optimal Transfers," *Working Paper*, 2019.
- Blundell, Richard, Luigi Pistaferri, and Ian Preston "Consumption Inequality and Partial Insurance," *American Economic Review*, 2008, 98 (5) , 1887–1921.
- Carroll, Christopher D., Jiri Slacalek, and Kiichi Tokuoka "Buffer-stock saving in a Krusell–Smith world," *Economics Letters*, 2015, 132 , 97 – 100.
- Carroll, Christopher, Jiri Slacalek, Kiichi Tokuoka, and Matthew N. White "The Distribution of Wealth and the Marginal Propensity to Consume," *Quantitative Economics*, 2017, 8 (3) , 977–1020.
- Chodorow-Reich, Gabriel, Plamen T. Nenov, and Alp Simsek "Stock Market Wealth and the Real Economy: A Local Labor Market Approach," *American Economic Review*, 2021, 111 (5) , 1613–57, URL: <https://www.aeaweb.org/articles?id=10.1257/aer.20200208>, DOI: <http://dx.doi.org/10.1257/aer.20200208>.
- Cloyne, James, Clodomiro Ferreira, and Paolo Surico "Monetary Policy when Households have Debt: New Evidence on the Transmission Mechanism," *The Review of Economic Studies*, 2019, 87 (1) , 102–129.

- Coibion, Olivier, Yuriy Gorodnichenko, Lorenz Kueng, and John Silvia "Innocent Bystanders? Monetary policy and inequality," *Journal of Monetary Economics*, 2017, 88 (C) , 70–89.
- Di Maggio, Marco, Amir Kermani, and Kaveh Majlesi "Stock market returns and consumption," *The Journal of Finance*, 2020, 75 (6) , 3175–3219.
- Fagereng, Andreas, Martin B. Holm, and Gisle J. Natvik "MPC Heterogeneity and Household Balance Sheets," *American Economic Journal: Macroeconomics*, 2021.
- Farber, Hank "Mobility and stability: The dynamics of job change in labor markets," *Handbook of Labor Economics*, 1999, 3 (B) , 2439–2483.
- Flynn, Joel, Christina Patterson, and John Sturm "Fiscal Policy in a Networked Economy," *Mimeo*, 2021.
- Fort, Teresa C., John Haltiwanger, Ron S. Jarmin, and Javier Miranda "How Firms Respond to Business Cycles: The Role of Firm Age and Firm Size," *IMF Economic Review*, 2013, 61 (3) , 520–559.
- Galí, Jordi, J. David López-Salido, and Javier Vallés "Understanding the Effects of Government Spending on Consumption," *Journal of the European Economic Association*, 2007, 5 (1) , 227–270.
- Ganong, Peter and Pascal Noel "Consumer Spending during Unemployment: Positive and Normative Implications," *American Economic Review*, 2019, 109 (7) , 2383–2424.
- Gelman, Michael "What drives heterogeneity in the marginal propensity to consume? Temporary shocks vs persistent characteristics," *Journal of Monetary Economics*, 2020.
- Glover, Andrew, Jonathan Heathcote, Dirk Krueger, and José-Víctor Ríos-Rull "Intergenerational Redistribution in the Great Recession," *Journal of Political Economy*, 2020, 128 (10) , 3730–3778, Publisher: The University of Chicago Press.
- Gross, Tal, Matthew J. Notowidigdo, and Jialan Wang "The Marginal Propensity to Consume over the Business Cycle," *American Economic Journal: Macroeconomics*, 2020, 12 (2) , 351–384.
- Gruber, Jonathan "The Consumption Smoothing Benefits of Unemployment Insurance," *American Economic Review*, 1997, 87 (1) , 192–205.
- Guvenen, Fatih, Sam Schulhofer-Wohl, Jae Song, and Motohiro Yogo "Worker Betas: Five Facts about Systemic Earnings Risk," *American Economic Review: Papers and Proceedings*, 2017, 107 (5) , 398–403.
- Guvenen, Fatih and Anthony A. Smith "Inferring Labor Income Risk and Partial Insurance from Economic Choices," *Econometrica*, 2014, 82 (6) , 2085–2129.
- Harrison, Glenn W., Morten I. Lau, and Melonie B. Williams "Estimating Individual Discount Rates in Denmark: A Field Experiment," *American Economic Review*, 2002, 92 (5) , 1606–1617.
- Hastings, Justine and Jesse M. Shapiro "How Are SNAP Benefits Spent? Evidence from a Retail Panel," *American Economic Review*, 2018, 108 (12) , 3493–3540.
- Heathcote, Jonathan and Fabrizio Perri "Wealth and Volatility," *The Review of Economic Studies*, 2018, 85 (4) , 2173–2213.
- Heathcote, Jonathan, Fabrizio Perri, and Giovanni L. Violante "Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States: 1967–2006," *Review of Economic Dynamics*, 2010, 13 (1) , 15–51.

- Hendren, Nathaniel “Knowledge of Future Job Loss and Implications for Unemployment Insurance,” *American Economic Review*, 2017, 107 (7) , 1778–1823.
- Hoynes, Hilary, Douglas L. Miller, and Jessamyn Schaller “Who Suffers During Recessions?” *The Journal of Economic Perspectives*, 2012, 26 (3) , 27–48.
- Jacobson, Louis S., Robert J. LaLonde, and Daniel G. Sullivan “Earnings Losses of Displaced Workers,” *The American Economic Review*, 1993, 83 (4) , 685–709.
- Jappelli, Tullio and Luigi Pistaferri “The Consumption Response to Income Changes,” *Annual Review of Economics*, 2010, 2 (1) , 479–506.
- “Fiscal Policy and MPC Heterogeneity,” *American Economic Journal: Macroeconomics*, 2014, 6 (4) , 107–36.
- “Reported MPC and Unobserved Heterogeneity,” *American Economic Journal: Economic Policy*, 2020, 12 (4) , 275–97.
- Jarosch, Gregor “Searching for Job Security and the Consequences of Job Loss,” working paper, 2014.
- Johnson, David S., Jonathan A. Parker, and Nicholas S. Souleles “Household Expenditure and the Income Tax Rebates of 2001,” *American Economic Review*, 2006, 96 (5) , 1589–1610.
- Kaplan, Greg and Guido Menzio “Shopping Externalities and Self-Fulfilling Unemployment Fluctuations,” *Journal of Political Economy*, 2016, 124 (3) , 771–825.
- Kaplan, Greg, Kurt Mitman, and Giovanni L. Violante “Non-durable consumption and housing net worth in the Great Recession: Evidence from easily accessible data,” *Journal of Public Economics*, 2020, 189 , 104–176.
- Kaplan, Greg, Benjamin Moll, and Giovanni L. Violante “Monetary Policy According to HANK,” *American Economic Review*, 2018, 108 (3) , 697–743.
- Kaplan, Greg and Giovanni L. Violante “A Model of the Consumption Response to Fiscal Stimulus Payments,” *Econometrica*, 2014, 82 (4) , 1199–1239.
- Kaplan, Greg, Giovanni L. Violante, and Justin. Weidner “The Wealthy Hand-to-Mouth,” *Brookings Papers on Economic Activity*, 2014, 2014 (1) , 77–138.
- Krusell, Per and Anthony A. Smith “Income and Wealth Heterogeneity in the Macroeconomy,” *Journal of Political Economy*, 1998, 106 (5) , 867–896.
- Lewis, Daniel J, Davide Melcangi, and Laura Pilossoph “Latent heterogeneity in the marginal propensity to consume,” *FRB of New York Staff Report*, 2019, (902) .
- McKay, Alisdair “Time-varying idiosyncratic risk and aggregate consumption dynamics,” *Journal of Monetary Economics*, 2017, 88 , 1 – 14.
- McKay, Alisdair, Emi Nakamura, and Jón Steinsson “The Power of Forward Guidance Revisited,” *American Economic Review*, 2016, 106 (10) , 3133–58.
- McKay, Alisdair and Ricardo Reis “The Role of Automatic Stabilizers in the U.S. Business Cycle,” *Econometrica*, 2016, 84 , 141–194.

- McKee, Graham and Emil Verner "The Consumption Response to Extended Unemployment Benefits in the Great Recession," Working Paper 1-056, Kilts Booth Marketing, 2015.
- Mian, Atif, Kamalesh Rao, and Amir Sufi "Household Balance Sheets, Consumption, and the Economic Slump," *Quarterly Journal of Economics*, 2013, 128 (4) , 1687–1726.
- Mian, Atif and Amir Sufi "What Explains the 2007–2009 Drop in Employment?" *Econometrica*, 2014, 82 (6) , 2197–2223.
- Misra, Kanishka and Paolo Surico "Consumption, Income Changes, and Heterogeneity: Evidence from Two Fiscal Stimulus Programs," *American Economic Journal: Macroeconomics*, 2014, 6 (4) , 84–106.
- Oh, Hyunseung and Ricardo Reis "Targeted transfers and the fiscal response to the great recession," *Journal of Monetary Economics*, 2012, 59 (S) , 50–64.
- Parker, Jonathan A. "Why Don't Households Smooth Consumption? Evidence from a \$25 Million Experiment," *American Economic Journal: Macroeconomics*, 2017, 9 (4) , 153–83.
- Parker, Jonathan A. and Nicholas S. Souleles "Reported Effects versus Revealed-Preference Estimates: Evidence from the Propensity to Spend Tax Rebates," *American Economic Review: Insights*, 2019, 1 (3) , 273–90.
- Parker, Jonathan A., Nicholas S. Souleles, David S. Johnson, and Robert McClelland "Consumer Spending and the Economic Stimulus Payments of 2008," *American Economic Review*, 2013, 103 (6) , 2530–53.
- Ravn, Morten O. and Vincent Sterk "Job uncertainty and deep recessions," *Journal of Monetary Economics*, 2017, 90 , 125 – 141.
- Ravn, Morten O and Vincent Sterk "Macroeconomic Fluctuations with HANK&SAM: an Analytical Approach," *Journal of the European Economic Association*, 2020.
- Rubin, D. B. *Multiple Imputation for Nonresponse in Surveys*: J. Wiley & Sons, 1987.
- Schulhofer-Wohl, Sam "Heterogeneity and Tests of Risk Sharing," *Journal of Political Economy*, 2011, 119 (5) , 925–958.
- Slacalek, Jiri, Oreste Tristani, and Giovanni L. Violante "Household balance sheet channels of monetary policy: A back of the envelope calculation for the euro area," *Journal of Economic Dynamics and Control*, 2020, 115 , 103879.
- Smith, Matthew, Danny Yagan, Owen Zidar, and Eric Zwick "Capitalists in the Twenty-first Century," *The Quarterly Journal of Economics*, 2019, 134 (4) , 1675–1745.
- Sorkin, Isaac "Ranking Firms Using Revealed Preference," *The Quarterly Journal of Economics*, 2018, 133 (3) , 1331–1393.
- Spletzer, James R. "Inequality Statistics from the LEHD," *CES Working Paper Series*, 2014.
- Sullivan, Daniel and Till von Wachter "Job Displacement and Mortality: An Analysis Using Administrative Data*," *The Quarterly Journal of Economics*, 2009, 124 (3) , 1265–1306.
- Vilhuber, Lars and Kevin McKinney "LEHD Infrastructure Files in the Census RDC," working papers, Centre for Economic Studies (CES), 2014.

Werning, Iván "Incomplete Markets and Aggregate Demand," Working Paper 21448, National Bureau of Economic Research, 2015.

A Data appendix

A.1 LEHD data

This project uses a 23-state subset of the LEHD. The LEHD is an unbalanced panel of states, and Table A2 reports the years for which each state is included in this analysis. I exclude individuals for whom I am missing industry and firm information. For each state, I drop the first two years of available individual-level data, as lagged incomes are not well-defined for workers in those years. In addition, I drop the first two years in which an individual appears in the entire sample. For the majority of individuals, this is the same restriction as dropping the first two years of the state. However, for workers who appear in the sample in the later years (e.g., workers moving into the sample, young workers, etc.), this is an additional restriction that ensures that the lagged earnings, and thus MPCs, are well-defined for all workers. This two-year lag restriction excludes about 20 percent of the sample in each year. I exclude another 12 percent of the original sample who are not between the ages of 25 and 62.

One potential concern with the rolling-panel structure of the LEHD is that the magnitude of the measurement error in constructing the total earnings series is changing over time. Both in constructing lagged incomes for the MPC estimation and in defining income processes over the business cycle, I rely on the full set of earnings across all states for the individual. As states enter the sample over time, total earnings for individuals employed across multiple states may jump artificially. In order to address these issues, I supplement the baseline analysis with analysis run on the subset of states with data available by 1993. These states, listed in bold in Table A2, represent 70 percent of the workers in the full sample. While I only report the robustness for the key findings, this subsample produces very similar patterns throughout the entire analysis.

The LEHD provides a comprehensive snapshot of employment in each quarter, but it does not provide information on labor market activity for workers in periods when they are not employed within this sample. Therefore, I must take a stance on the labor market activity of workers who leave my sample. This assumption enters in both my measurement of labor market outcomes and my calculation of an individual's MPC, in so far as it affects the level of an individual's average earnings in the two previous years.

Throughout the main analysis, I assume that prime-age workers who leave employment in my sample make no labor market earnings in those quarters. This assumption would be violated if individuals moved to a job either outside the LEHD coverage (i.e., to the military, federal employment, or self-employment) or to a state that is not in my sample. Using the ACS subsample, I find that among workers who leave my LEHD sample between $t - 1$ and t , and are in the ACS in year t , 24 percent report being employed elsewhere, suggesting that this margin is potentially non-negligible. However, this would be a serious problem for my covariance estimate only if workers of different characteristics were differentially likely to move outside the LEHD sample over the business cycle. For example, if younger workers were disproportionately likely to move to states outside of my sample during recessions, then I would overstate the unemployment of young workers in recessions, and thus erroneously conclude that the earnings of young workers were more sensitive to recessions, when in fact they are not. In addition to overstating the sensitivity of these workers' earnings to GDP, I may exaggerate the difference in MPCs between workers of different ages, as the lagged earnings of younger workers would be biased downward (because they had positive earnings in other states, rather than the zero earnings that I assume). Note that since I mostly focus on the cross-section of employment, and since I remove the initial employment periods, this bias in the measurement of lagged earnings would only appear for workers who moved out of my sample to noncovered employment and then returned to my sample in a future period. However, both of these patterns together would lead me to potentially overstate the differential sensitivity of workers of different

MPCs to business cycles.

While I cannot fully address these concerns, I explore their importance in two ways. First, using the matched monthly basic Current Population Survey (CPS), I explore the probabilities that workers of different MPCs transition from private sector employment into the military, self-employment, or the federal government, all of which are sectors beyond the scope of the LEHD. The CPS features a rolling panel structure wherein individuals are interviewed for four months, have eight months off, and then are interviewed again for four months. I flag an individual as making a transition to noncovered employment at time t if they are newly self-employed, in the military, or in federal government at time t and were in the private sector in any previous survey. Table A14 shows how the probability of moving to noncovered employment varies over the business cycle. Column 1 shows that on average, high-MPC workers are less likely to move to employment outside the LEHD sample; Column 2 shows that on average, transitions from LEHD employment to noncovered employment are less likely during recessions; while Column 3 shows that there is no differential sensitivity in mobility by a worker's MPC.³³

Second, within the PSID, I calculate the baseline LEHD regressions in Table 2 within the PSID. Since the PSID includes all states, I can directly compare the estimates using the actual observed change in earnings and then replacing worker earnings with 0 if they moved from one of the states in the LEHD sample to a state that is not in the LEHD sample. The second of these outcomes is the equivalent to what I assume within the LEHD, while the first includes the effects of mobility. The gap between these estimates should demonstrate the importance of these cross-state moves for my LEHD estimates. Table A15 shows the estimates for total earnings (columns 1 and 2), the extensive margin of employment (columns 3 and 4) and the intensive margin of earnings (columns 5 and 6). As you can see from comparing the estimates in column 1 to column 2, column 3 to column 4, and column 5 to column 6, I find that this assumption makes only a very small difference overall.

A.2 PSID data

The PSID surveys households annually from 1968 to 1997 and every other year from 1997 to 2015. Each household in the PSID is interviewed once a year, primarily between March and June. In each year of the survey, households are asked about the demographics and labor market status of the household head and spouse. For the length of the survey from 1968 to 2015, households are also asked about how much their household spent on food in the average week. While this is not specific to the week of the interview, it likely refers to the recent period. The income measures, however, refer to the annual income for the previous tax year (i.e., 2007 income is collected in the 2008 survey), and I use the panel structure of the data to get a measure of annual labor market earnings in the year of the interview. This means that the income measure will capture all income in that calendar year, while the consumption and labor market variables will refer to the survey month.

Table A3 shows the summary statistics for the PSID sample. Columns 1 and 2 show the snapshot of employed workers in the PSID in the LEHD sample states and nationally, respectively. These two samples are very similar to each other, showing that the LEHD sample is largely nationally representative. A comparison of Column 1 in Table A3 and Column 3 of Table A1 in the main text also shows that the PSID sample is similar to the LEHD sample on all demographics. The only exception is that the average income in the PSID is higher than the average income in the LEHD. This is likely because the PSID is restricted to household heads and spouses who are employed full time, while the LEHD includes all workers who had any covered earnings in that quarter. Lastly, Column 3 shows summary statistics for the sample used in the

³³ Since the CPS does not record the lagged income over the previous two years, I impute the MPC of the individual using only demographic characteristics (race, age, and gender). Within the PSID, the MPCs that result from this modified imputation are similar to those that result from the baseline estimation used in the LEHD.

baseline estimation of MPCs. This sample differs from that in Column 2 along several dimensions. Most significantly, it includes workers who were employed in $t - 2$ but makes no restrictions on employment status in time t . It also adds the SEO sample, which oversamples low-income households.

A key limitation of the PSID data is that the main measure of expenditure is food. Figure A1 shows that the fraction of total expenditures that is spent on food is changing across the income distribution. The upward slope at the low end of the income distribution reflects the phase out of food stamps, which subsidize the consumption of food for lower-income households.

The main method that I use to impute total consumption in the PSID closely follows the methodology laid out in Blundell et al. (2008). They propose a method to impute expenditures in the PSID using information in the Consumer Expenditure Survey (CEX).³⁴ The approach involves estimating a demand system for food as a function of nondurable consumption, demographic variables, and relative prices. Under the assumption that food demands are monotonic, this demand function can be inverted to get an estimate of total consumption in the PSID. In order to deal with measurement error in expenditures, Blundell et al. (2008) instrument the nondurable consumption with the average (by cohort, year, and education) hourly earnings of the husband and wife.

I modify and simplify the Blundell et al. (2008) analysis along several dimensions. First, in order to be consistent with the second method described below, I estimate the reverse of the equation (i.e., estimate total consumption as a function of food consumption). This decision does not affect the estimates, as the resulting consumption series is nearly identical when using the reverse specification. Second, I estimate this relationship for both durable and nondurable consumption combined.³⁵ Third, I include households with single household heads. I do this because I am including those households within the PSID and LEHD. Fourth, I estimate the relationship using OLS, rather than using the instruments as the original paper does. Fourth, I use an updated sample, extending their sample to 2013. Table A4 shows the summary statistics for the CEX sample used in the imputation. The sample looks very similar to the PSID sample on demographics, except that on average, the workers have slightly lower incomes. Using this sample in the CEX, I estimate the following equation

$$\ln C_{ht} = Z_{ht}\beta + p_t\gamma + g(f_{ht}, X_{ht})\theta + u_{ht} \quad (\text{A1})$$

where C_{ht} is total household consumption, Z_{ht} are household demographics, p_t are relative prices (i.e., Consumer Price Index, or CPI, series), f_{ht} is food consumption, and X_{ht} are demographic characteristics and time dummies that shift the relationship between food consumption and overall consumption. These time dummies allow the food share to shift over time and can be used because the PSID and CEX have overlapping time frames.³⁶ Figure A2 shows that imputed total consumption in the PSID closely captures the relationship between food consumption and total consumption across the income distribution in the CEX.

An alternative methodology for imputing total consumption in the PSID is to follow Attanasio and Pistaferri (2014) and use the relationship between food consumption and overall consumption in the later years of the sample to impute the total consumption in the previous years of the sample.³⁷ This imputation

³⁴ The CEX includes much more comprehensive measures of consumption. Indeed, it covers around 95 percent of all expenditures, excluding housekeeping, personal care products, and nonprescription drugs. In the interview survey data, consumption is recorded for each month in the three months preceding the interview. This is then aggregated to create measures of total quarterly consumption for each household in an array of spending categories.

³⁵ Specifically, I estimate this equation using total consumption, which is the sum of nondurable and durable consumption. However, results are very similar when I estimate it separately for nondurable and durable consumption and then aggregate.

³⁶ I follow Blundell et al. (2008) in my choice of controls. These include dummies for the number of children in the household; three education bins; a quadratic in age; region of residence dummies; an indicator for being white; and education, year, and children dummies. All of these are interacted with food consumption.

³⁷ The expanded consumption measure within the PSID includes home insurance, rent, electricity, heating, water and miscella-

approximates total consumption less food (i.e., net consumption) as a function of demographics and food consumption

$$\ln n_{ht} = Z_{ht}\beta + p_t\gamma + g(f_{ht})\theta + u_{ht} \quad (\text{A2})$$

where n_{ht} is the net consumption of the household, Z_{ht} are various socioeconomic variables, p_t are prices, and f_{ht} is food consumption. I estimate this equation restricting the sample to one observation per household and include controls individually for the demographics of the head and spouse. The implicit assumption in this imputation is that the preferences of individuals are stable over time, and thus, the relationship between overall consumption and food consumption remains stable. This contrasts with the assumption in the CEX-based imputation that uses the same time period from two different samples.³⁸ Using the β , γ , and θ that result from estimating Equation A2 on the 1999 to 2013 subsample, I recover an estimate of total household consumption in each year using

$$\widehat{c}_{ht} = f_{ht} + e^{Z_{ht}\widehat{\beta} + p_t\widehat{\gamma} + g(f_{ht})\widehat{\theta}}$$

A.3 Additional estimates of marginal propensities to consume

A.3.1 Robustness of baseline MPC estimates

In the baseline estimation of MPCs, I restrict the sample to those individuals who are employed in year $t - 2$. I restrict to $t - 2$ rather than $t - 1$ so that I can include the later years of the PSID sample when the survey is collected every two years.³⁹ The changes in both income and consumption are also defined over a two-year period. From the entire PSID sample, I exclude observations that do not meet the panel structure necessary to define two-year changes in income and consumption, restrict attention to those between ages 25 and 62 in year t , and drop observations with missing race or education. In addition, in each regression, I exclude observations where the two-year change in log consumption or log income is more than 4.⁴⁰ I define an individual's lagged income as the labor market earnings for the individual in years $t - 1$ and $t - 2$. I group this average into five approximately equally sized bins: $< \$22,000$, $\$22,000 - \$35,000$, $\$35,000 - \$48,000$, $\$48,000 - \$65,000$, and $> \$65,000$. The measure of lagged income is intended to capture differences in permanent earnings capacity across groups. However, I find that patterns across lagged incomes are not sensitive to the particulars of how lagged earnings are defined; the same patterns for estimated MPCs appear when using additional income lags or fixing earnings at a given age, which may capture a more permanent measure of income.

Table A5 and Figure A4 display supporting statistics for the baseline estimates discussed in detail in the main text. Specifically, Figure A4 show the first-stage and reduced-form estimates associated with Figure 2 in the main text. The left panel shows substantial variation in the effect of unemployment on the level of labor income, with the largest falls, unsurprisingly, being among the highest earners. The right

neous utilities, car insurance, car repairs, gasoline, parking, bus fares, taxi fares and other transportation, school tuition and other school expenses, child care, health insurance and out-of-pocket health costs, and food.

³⁸ I closely follow Attanasio and Pistaferri (2014) in parameterizing controls in Z . Like them, I include a third-degree polynomial in total food consumption; dummies for age, education, marital status, race, state, and employment status; the hours worked by the household head; homeownership status; family size and the number of children in the household, and consumer price indices to capture relative prices (overall CPI, CPI for food at home, CPI for food away from home, and CPI for rent). I also include household income as a consumption shifter and the spouses' labor market variables as controls in Z .

³⁹ These later years are particularly important both because they overlap with the time period of the LEHD and because they represent a significant fraction of the years for which the CEX and the PSID overlap – and thus the years for which I have the expanded consumption measure.

⁴⁰ This restriction on outliers is similar to that in Hendren (2017), who excludes individuals with more than a threefold change in food consumption, and Gruber (1997), who excludes observations with a greater than 1.1 log change in food consumption.

panel shows that there is less, although still substantial, variation in the level of the consumption drop across households. Table A5 shows the regression estimates for Equation 5 that produce the distribution of MPCs shown in Figure A3 in the main text. The left column reports regression coefficients using food consumption only; the middle panel shows estimates using the PSID-based imputation measure, which is described in Appendix A.2; and the third panel shows the baseline estimates using the CEX-based imputation of total consumption, which is used as the baseline consumption measure throughout this analysis. The fourth panel shows results using unexpected job loss as the instrument. Unsurprisingly, these multivariate estimates echo the patterns displayed in Figure 2, in which black, lower-income, and young workers have higher MPCs.

The left panel of Figure A6 shows the overall estimates of the marginal propensity to consume that result from re-estimating Equation 5 using different identifying income shocks. First, the left-most point shows the OLS version of Equation 5. The coefficient is close to 0, suggesting a substantial downward bias and the need for an instrument to identify the causal relationship between consumption and income movements. However, while the use of an instrument matters critically, across the x-axis, estimates of the MPC are relatively stable to the type of income shock used as the instrument. For comparison, the second point shows the baseline MPC estimated using the unemployment shock. The next four estimates show the MPC estimated using either the change in state GDP or the national unemployment rate of the worker’s industry. For an individual worker, these aggregate changes are plausibly exogenous to their own earnings and affect their earnings both positively and negatively and on both the intensive and extensive margins. While noisier, the average MPC estimates are only slightly smaller. This is true whether I include all workers (as in points 3 and 4) or restrict to those workers who remain employed (as in points 5 and 6). Those who remain employed across years experience a smaller income change yet a similar MPC.⁴¹ Lastly, the farthest-right point shows the average MPC estimated using an indicator for whether the worker becomes employed between $t - 2$ and t .⁴² The average MPC is slightly higher with the positive income “shock,” but this is an artifact of the different estimation samples – the hires estimation restricts to the nonemployed, who, on average, have higher MPCs than the employed. When averaged on the same sample, the estimates are similar. The right panel of Figure A6 shows that not only are the averages similar for these different shocks, but the alternate MPC estimates are also highly correlated at the individual level.

Figure A7 explores the stability of MPCs over the business cycle and plots the bivariate versions of Equation 5, modified to allow the MPC to vary with the state unemployment rate. The blue circles plot the MPC at the average unemployment rate in the state, and the red squares show the implied MPC at 4 percentage points above the average unemployment rate. Generally, the MPC is somewhat lower in recessions, but the differences are economically and statistically insignificant.

Table A6 shows the correlation between the baseline MPC estimates and those estimates including various job-level characteristics. In the first specification, I include the individual’s tenure with their current employer. This variable is intended to capture some amount of private information on the riskiness of the individual’s job, as workers with longer job tenures are less likely to lose their jobs (Farber (1999)). Column 3 adds the lagged variance of an individual’s earnings. Columns 4 and 5 include the variance of an individual’s lagged earnings, capturing the fact that individuals with a higher earnings variance may differ in their MPCs. This variable is calculated using the matched monthly basic CPS from 1976 to 2013 and is the sample average of the change in log earnings between interview 4 and interview 8, which are

⁴¹ I find that those who remain employed across surveys experience a drop in total hours worked when the unemployment rate is high, suggesting moves to part-time employment.

⁴² This specification includes only those who were not employed in $t - 2$; thus the control group is the set of individuals who remain nonemployed between $t - 2$ and t . Patterns are robust to including only those who are unemployed, rather than nonemployed, in $t - 2$.

a year apart. This variable is intended to capture the expected variability of earnings of the job. Lastly, I include dummies for the census region of residence, allowing MPCs to vary geographically. As shown in Table A6, the resulting MPC heterogeneity does not change meaningfully when these variables are added.

A.4 Additional results for $Cov(MPC_i, \gamma_i)$

A.4.1 Adjusted Standard Errors

Clustered standard errors for Equation 6 do not take into account the additional noise imposed by the imputation of a worker’s MPC. Indeed, the worker MPC estimates rely on two imputations. I first impute total consumption in the PSID using the CEX, and then I impute the MPC in the LEHD using the MPC estimates from the PSID. In order to account for the additional noise injected at each of these steps, I implement multiple imputation techniques, as in Rubin (1987). Specifically, I take 100 draws in which I randomly sample with replacement both the CEX and the PSID. This produces 100 estimates of main-text Equation 5, which estimates the MPC for each demographic group x . In the LEHD, I then estimate main-text Equation 6 for each imputation, which results in 500 estimates of the degree to which workers of different MPCs are exposed to recessions. I combine these various estimates following the formulas derived in Rubin (1987):

$$\bar{\alpha}_2 = \sum_{i=1}^{100} \frac{\hat{\alpha}_{2,i}}{100}$$

$$var(\alpha_2) = \sum_{i=1}^{100} \frac{var(\hat{\alpha}_{2,i})}{100} + \sum_{i=1}^{100} \frac{(\hat{\alpha}_{2,i} - \bar{\alpha}_2)^2}{99} + \sum_{i=1}^{100} \frac{(\hat{\alpha}_{2,i} - \bar{\alpha}_2)^2}{99 * 100}$$

The point estimate is the average across imputation draws, while the variance of the estimate is the combination of the average within-draw variance and the between-draw variance.

A.4.2 Robustness of LEHD-based estimates

Figure 1 clearly demonstrates that there is a positive relationship between the average earnings cyclicality of a demographic group and the average marginal propensity to consume of that group. The additional results presented in this section support the robustness of this pattern. First, Figure A8 shows this pattern separately for the intensive margin of earnings (i.e., earnings conditional on remaining employed between $t - 1$ and t) and the extensive margin of employment. The figure clearly shows that higher-MPC workers are more likely to become unemployed during recessions and earn less conditional on remaining employed. Indeed, the same demographic groups that are exposed on the intensive margin are also exposed on the extensive margin of earnings.

Tables A7 and A8 probe the robustness of the positive relationship between a worker’s MPC and the exposure of their earnings to recessions to several data decisions. Table A8 shows that the estimated relationship is robust to various methods of imputing MPCs. Since the magnitude of the MPCs changes with the imputation method, so does the magnitude of the coefficient, but the proportional relationship is fairly stable (See Table 3 in the main text for implied amplification for each measure).

Table A7 show that the covariance is relatively stable across specifications. When estimating the relationship between earnings sensitivities and MPCs at the individual level, taking the log of earnings will restrict to those who remain employed. Therefore, in the baseline analysis, I estimate an overall earnings elasticity at the individual level by replacing the change in log earnings with $\Delta E_{i,t} = \frac{E_{i,t} - E_{i,t-1}}{.5 * E_{i,t} + .5 * E_{i,t-1}}$, which

bounds the earnings loss of the nonemployed at negative 2. However, column 2 shows that the patterns are similar when using a log transformation (i.e. $\Delta E_{i,t} = \log(E_{i,t} + 100) - \log(E_{i,t-1} + 100)$).

The following columns of Table A7 show that the estimated relationship is robust to various other modifications. Baseline estimates consider movements in fourth-quarter earnings, but Column 3 shows that patterns are still present when considering annual incomes. Column 4 restricts to a balanced panel from 1995 to 2011 and finds that in this subset, the estimate is similar but slightly larger. Column 5 replaces aggregate GDP with state-level GDP and shows that heterogeneity patterns are similar, suggesting that these patterns hold not only across states but also within states.

A.4.3 Relationship to Guvenen et al. (2017)

Using individual earnings data from the U.S Social Security Administration, Guvenen et al. (2017) document that the earnings of both the lowest and the very highest earners are more sensitive to aggregate income fluctuations. In Figure A9, I largely replicate this finding using my sample within the LEHD. The left panel shows the elasticity of worker earnings to aggregate GDP by the income decile of the worker. As in Figure 1 of Guvenen et al. (2017), there is a U-shaped relationship; the sensitivity of worker earnings is decreasing through much of the earnings distribution, but it spikes again at the very top of the income distribution. The LEHD misses the income of top earners along two dimensions. First, it is topcoded, generally at the 99th percentile within the state. Second, it includes only UI-covered income, and therefore misses the high earning self-employed. I benchmark the degree to which this missing income may bias my result by noting the amount of potential missing income and reestimating Equation 6 assigning to that share of income 1) the lowest MPC estimate and 2) the highest earnings elasticity estimate. I find that this extreme case attenuates the relationship by around 20 percent, suggesting that this omission is potentially meaningful but not large enough to negate the mechanism.

How does this relate to the relationship between worker MPCs and earnings sensitivity documented in Figure 1? This U-shaped relationship in earnings does not directly imply any relationship between worker MPCs and earnings sensitivity – the bottom of the income distribution has high estimated MPCs but a small share of overall earnings; the top of the income distribution has both low MPCs and the majority of the earnings in the economy; and in my estimation, lagged income only explains 40 percent of the overall variation in MPCs. For a direct comparison, the right panel of Figure A9 shows earnings elasticities by the MPC decile of the worker. There are two important observations. First, the overall pattern is upward-sloping, meaning that workers in the top decile have the highest income sensitivity. Second, there is a nonlinear pattern in the bottom deciles, likely reflecting the increased earnings sensitivity of the very higher earners.

A.5 Alternate Estimates of $Cov(\gamma_i, MPC_i)$

A.5.1 PSID-Based Estimates

I complement the baseline analysis with estimates entirely contained within the PSID. While the sample is substantially smaller and it does not have the benefit of high-quality administrative data, the PSID allows me to extend the analysis along several dimensions. First, it includes more comprehensive coverage of the labor force, namely the inclusion of the federal government and the unemployed. Second, it provides a longer time period, allowing me to explore changes in this moment over time. Third, it includes a more comprehensive measure of income (i.e. total income as opposed to labor market earnings and household income as opposed to individual-level income). This allows me to supplement the labor market analysis, which is the main focus of the paper, with an analysis of how other margins of income adjust as well. In

Figure A10 and A11, I report the estimates using the PSID and explore the importance of these various dimensions. For the purposes of facilitating a direct comparison with the results in LEHD, I use the baseline MPC estimates that utilize the CEX imputation. In the left panel, I closely replicate the baseline analysis in the PSID by restricting to the set of incumbent workers in the private sector between 1997-2011. In the right panel, I extend the sample to include government workers and unemployed workers (top estimates), and earlier years in the sample (bottom estimates). I find that extending the sample to include other workers does not meaningfully affect the covariance estimate, but that the estimate shrinks substantially when including the full PSID sample from 1968-2011. While the noise in the sample makes it challenging to make a strong claim, this suggests that the covariance between MPCs and earnings elasticities has grown over time.

Figure A11 shows estimates of the covariance for household income, rather than individual income. There are two main takeaways. First, results are similar when moving from the individual to the household level, suggesting that intra-household dynamics are not important for driving this moment. Second, the covariance is still large and positive when considering total household income, as opposed to just labor income. There is a slight negative covariance for other income (i.e. transfers go up more for high-MPC households while capital income falls for low-MPC households) but this is comparatively small and the overwhelming effect comes from labor income.

A.5.2 Italian Survey on Household Income and Wealth (SHIW)

I use data from the 1998-2018 waves of the survey, when the survey is available every other year. The survey has a rotating panel dimension – each year, about 50% of respondents were there in the previous year while about 50% are new. This means that some households are in the sample for many years while others are only in the survey once. Directly-reported MPC information is available in 4 years within the sample:

- In 2016, respondents were asked *“Imagine you unexpectedly receive a refund equal to the household’s monthly income. How much of the sum would you save and how much would you spend?”*
- In 2012, respondents were asked *“Imagine you receive an unexpected inheritance equal to your household’s income for a year. Over the next 12 months, how would you use this windfall?”*
- In 2010, respondents were asked *“Imagine you unexpectedly receive a reimbursement equal to the amount your household earns in a month. How much of it would you save and how much would you spend?”*
- In 2000, respondents were asked *“If you were informed that you had unexpectedly won the sum of 10 million lire, payable immediately, by how much would your consumption increase during 2001?”*

Using this information on self-reported MPCs, I construct several possible measures of household MPCs:

- Raw reported MPC: This is the simplest measure since it requires no imputation. This is available for 30 percent of the sample.
- 2010 MPC: I assume here that household MPCs are fixed over time and assign the 2010 reported MPC to all observations for the household that are in the sample in 2010. This provides an MPC estimate for 39 percent of the sample. This is unlikely to be a good assumption, since the MPCs within household over time are not highly correlated.⁴³

⁴³The correlation between the 2010 MPC estimates and the 2000, 2012, and 2016 MPC estimates are 0.04, 0.12, and 0.17, respectively. See also Jappelli and Pistaferri (2020).

- **Fitted MPC:** I project the self-reported MPCs from each year on cash-on-hand deciles and 4 age bins. I choose these variables since they are the variables shown to be most important in Jappelli and Pistaferri (2020). I then impute the MPC for all households based on these estimated relationships. This is available for 100 percent of the sample and allows MPCs to change over time within household.

I also supplement these self-reported MPCs with an MPC that I estimate following the methodology outlined in Section 4.1 using the panel-structure of the SHIW and the unemployment shock. Specifically, I estimate

$$\Delta C_{i,t} = \sum_x \beta_x \Delta E_{i,t} \times x_{i,t} + \gamma_t + \epsilon_{i,t}$$

where C is reported total household consumption, E is labor income for individual i , and x is a set of household characteristics. As in US analysis, I restrict the estimation to those employed in $t - 2$ and I instrument the change in income with an indicator for whether the individual is unemployed in time t . The left panel of Figure A12 shows the resulting MPC estimates – interestingly, the patterns look very similar to those that I uncover in the United States, with younger and poorer workers having higher MPCs.

Lastly, using these various MPC estimates, I calculate the following regression:

$$\Delta \log E_{h,t} = \beta_0 + \beta_1 MPC_{h,t-2} + \beta_2 MPC_{h,t-2} \times \Delta \log(Y_t) + \epsilon_{h,t}$$

where $E_{h,t}$ is household income and $MPC_{h,t-2}$ is the MPC of the household in the previous survey. Households are weighted by their income in $t - 2$. I use household income as the baseline because, with the exception of the unemployment-based MPCs, the MPCs are at the household level. However, I also show results for individual payroll income. Table A10 shows that the estimates of β_2 are generally large, positive, and statistically significant. The magnitudes are also meaningful – the standard deviation of the reported MPC is 0.34. The average elasticity of total household earnings to GDP is 0.58 and the coefficient in panel A, column 1 suggests that households with 1 standard deviation higher MPCs have household earnings elasticities that are 0.28, or just about 50 percent, higher than average.

A.5.3 CEX Analysis using Tax Rebates

For this section, I closely follow Misra and Surico (2014), Johnson et al. (2006) and Parker et al. (2013). Specifically, I pool observations across the 2001 and 2008 tax rebates and estimate MPCs as:

$$C_{h,t+1} - C_{h,t} = \sum_s \beta_{0_s} \times month_{s,h} + \beta_1 X_{h,t} + \beta_2 \text{Rebate}_{h,t+1} + \epsilon_{h,t+1}$$

where h is the household, s is the month, and C is nondurable household consumption. X represents a set of controls, including worker age, changes in the number of adults and children in the household, and a dummy for whether the data is from the 2001 or 2008 experiment. Following the literature, I instrument the size of the tax rebate with an indicator for whether the rebate is positive. The right panel of Figure A12 shows the resulting patterns. While the standard errors are large, the patterns are familiar, with lowest income workers having the largest MPCs. These broad patterns are consistent with those in Johnson et al. (2006) and Parker et al. (2013).

Since these MPCs are only directly observable for the households that are in the sample in 2001 and 2008 and the CEX only includes a short time series, I must project these MPCs onto household characteristics in order to explore the covariance between these MPCs and earnings elasticities for workers. Since the CEX has only a short panel, in order to do the analysis entirely within the CEX, I would need to use pseudo-cohorts. However, this would mean that I would be constrained to using the set of characteristics

that are largely invariant over the business cycle to avoid endogenous movements of workers across cohorts. Therefore, I instead impute these MPC estimates for households in the PSID, where I do observe the same household over time. In order to facilitate a comparison with the unemployment-based MPCs in the main body of the paper, I project these tax-rebate MPCs onto the following household characteristics: 4 age bins, 2 race bins and 5 income bins. Note, again, that unlike the unemployment-based MPCs, these MPC estimates are at the household level, and therefore, within the PSID, I explore the elasticity of *household* earnings by the MPCs estimated using the 2001 and 2008 tax rebates. Table A11 shows the results. Columns 1 and 4 show the results using the labor income of the household head and spouse as the measure of household earnings, columns 2 and 5 use total household income as the dependent variable, and columns 3 and 6 use individual labor earnings. In all cases, the interaction term between MPCs and earnings elasticities is positive, although it is not precisely estimated. The magnitude is also meaningful – the estimates in column 1 show that being 1 standard deviation higher in the MPC distribution is associated with an elasticity of household income to aggregate GDP that is 0.17 higher than someone with the average MPC.

Comparing MPCs out of tax rebates to MPCs out of unemployment There are potentially many reasons that the MPCs in response to tax rebates and those in response to unemployment shocks could be different both in levels and in their distribution across the population. Tax rebates are positive income shocks while unemployment shocks are negative on average, tax rebates are temporary income shocks while unemployment shocks are more persistent, and tax rebates are generally small shares of household balance sheets while unemployment shocks have very large effects on annual incomes (See Figure A4). These differences are not only true on average, but they also potentially vary across demographic groups in important ways (e.g. unemployment is more persistent for some groups). Despite these many potential differences, I find that resulting MPC estimates are similar, both on average and across the distribution.

Specifically, to facilitate a direct comparison, I focus on the MPC for nondurable consumption. For MPCs upon unemployment, I estimate MPCs using the set of demographics characteristics as described in the main text (See Column 2 of Table 3 in the main paper). For MPCs out of tax rebates, I impute MPCs within the PSID using the households characteristics described above and used in Table A11. Despite the many differences, I find that the *average* MPCs within the PSID are somewhat different, with an average estimate of .25 for tax rebates and .41 for unemployment. However, the MPC estimates are positively correlated across demographic groups with a correlation of 0.34.

One particular concern with MPC estimates using unemployment is that the heterogeneity could be driven by differences in the duration of unemployment. If duration were driving the MPC heterogeneity, that would change the interpretation of the results. Indeed, within the PSID, I find that there is a positive correlation between the average duration of unemployment for a demographic group and the MPC estimate for that demographic group. However, I find a very similar correlation between MPCs calculated using tax rebate heterogeneity and the average unemployment duration of the demographic group as well. Specifically, Figure A5 plots estimates of the MPC versus average unemployment duration for income and age groups. There is a positive relationship for *both* estimates, with the slope of the relationship being nearly identical. Since tax rebates are clearly transitory, the correlation between the MPC and unemployment duration likely reflects other cross-sectional differences that correlates with average unemployment duration and affect MPCs such as time preferences, credit constraints or liquidity, rather than being a direct function of unemployment duration.

A.6 Details of commuting zone analysis

A.6.1 Additional data definitions

Local control variables: I closely follow Kaplan et al. (2020) in defining household wealth measures in each local labor market. Specifically, I define housing wealth as the total number of housing units in a county, which are published by the U.S. Census Bureau, multiplied by the Zillow Home Value Index for All Homes.⁴⁴ Data on household debt come from the Federal Reserve Bank of New York Consumer Credit Panel (CCP), which provides the dollar values of mortgage, auto, and revolving credit debt annually in each county from 1999 to 2011. I define household debt as the total value of both mortgage and non-mortgage debt. I construct data on financial assets by allocating total financial assets in a quarter from the Flow of Funds Balance Sheet of Households and Nonprofit Organizations to counties using the fraction of total financial assets in that county from the quarterly IRS Statistics of Income. Lastly, I aggregate these county-level measures to the commuting zone level, restrict attention to fourth-quarter estimates, and divide by population estimates to obtain per-capita values.⁴⁵

The fraction of the commuting zone that is employed comes from the ACS. All other control variables – namely, demographic controls for the area and the average size and age of establishments by commuting zone – are calculated within the LEHD in each year.

Bartik shock: I construct a Bartik-style shock at the commuting zone level using

$$Shock_{c,t} = \sum_i \frac{L_{i,c,t_0}}{L_{c,t_0}} \Delta \log E_{i,t,-c}$$

where $t_0 = 1999$, L is employment, and $\Delta \log E_{i,t,-c}$ is the change in the log of total earnings in industry i within the states that are in the LEHD subsample in year $t - 1$ but excluding earnings in commuting zone c . I exclude own commuting zones, since my LEHD sample is not national and thus any given commuting zone may represent a non-negligible fraction of the industry's total earnings. Additionally, because the LEHD is not balanced across states, the aggregated time series for any given industry are inconsistent due to the entry of states. Therefore, to be consistent over time, I define the change in earnings in an industry using only the incumbent states in each year. Unsurprisingly, this shock is very highly predictive of both changes in overall earnings in a commuting zone and movements in the GDP of the commuting zone's state.

In exploring the role of the local labor market multiplier in affecting local cyclicity using this Bartik shock, I *both* re-estimate MM_c using this shock and then re-estimate Equation 10 replacing aggregate GDP with this shock.

A.7 The Matching Multiplier and the labor share

While the matching multiplier is derived in a setting in which all output is earned by labor, in order to provide *empirical* estimates of the matching multiplier, I need to take into account the fact that in reality,

⁴⁴ This housing index is publicly available monthly for each U.S. county beginning in 1996. Housing units are available annually at the county level back until 2000. Prior to 2000, these counts are only released at the state-by-year level. I interpolate county housing units prior to 2000 by assuming that the fraction of houses in each county in the state is constant and by assigning total state housing units in each year to counties based on the 2000 distribution.

⁴⁵ CCP data are only released for counties with an estimated population of at least 10,000 consumers with credit reports in the fourth quarter of 2010. This restriction excludes 20 percent of counties. Since these are predominately small counties, I ignore these missing values in the aggregation from counties to commuting zones.

not all output goes to worker wages. Since the focus of this paper is on quantifying a particular mechanism within the labor market, I do not explore potentially important heterogeneity in MPCs out of nonlabor income. Rather, I assume the covariance of MPCs out of non-labor income with the elasticity of non-labor income to the aggregate is zero and I make modest adjustments to the simple framework to rescale the contribution of this particular mechanism. Consider the case where output is given by $Y = E + K$, where E are labor market earnings and K are earnings from nonlabor income (e.g., profits, return on capital, etc). In this case, the aggregate MPC in the economy is given by:

$$MPC = \frac{dC}{dY} = \frac{\partial C}{\partial E} \frac{dE}{dY} + \frac{\partial C}{\partial K} \frac{dK}{dY}$$

Using $dY = dE + dK$, defining the labor share $\frac{E}{Y} = \alpha_l$, and $\gamma = \frac{dE}{dY} Y E$,

$$MPC = \frac{\partial C}{\partial E} \frac{dE}{dY} + \frac{\partial C}{\partial K} \frac{dY - dE}{dY}$$

$$MPC = \frac{\partial C}{\partial E} \frac{dE}{dY} + \frac{\partial C}{\partial K} \left(1 - \frac{dE}{dY}\right)$$

$$MPC = \frac{\partial C}{\partial E} \frac{dE}{dY} + \frac{\partial C}{\partial K} (1 - \alpha_l \gamma)$$

Summing across individuals in the labor earnings term, we get:

$$MPC = \sum_i \frac{\partial C_i}{\partial E_i} \frac{dE_i}{dY_i} + \frac{\partial C}{\partial K} (1 - \alpha_l \gamma) = \frac{E}{Y} \sum_i MPC_i \frac{E_i}{E} \gamma_i + \frac{\partial C}{\partial K} (1 - \alpha_l \gamma)$$

$$MPC = \frac{dC}{dY} = \alpha_l \left(\sum_i MPC_i \frac{E_i}{E} \gamma_i \right) + \frac{\partial C}{\partial K} (1 - \alpha_l \gamma)$$

This simple total derivative highlights the importance of three terms that were not in the simple framework – the consumption response from changes in nonlabor income ($\frac{dC}{dK} = MPC_{nl}$), the labor share ($\alpha_l = \frac{E}{Y}$) and the average elasticity of labor earnings with respect to the aggregate $\bar{\gamma}$. The importance of the labor share is intuitive – a mechanism affecting labor market income matters more for the total economy when labor earns a higher share of total income.

The benchmark case is one in which all workers have the average MPC (i.e. $\frac{\partial C}{\partial E} \frac{dE}{dY} = \bar{\gamma} \overline{MPC}$)

$$MPC^b = \alpha_l \bar{\gamma} \overline{MPC} + (1 - \alpha_l \bar{\gamma}) MPC_{nl}$$

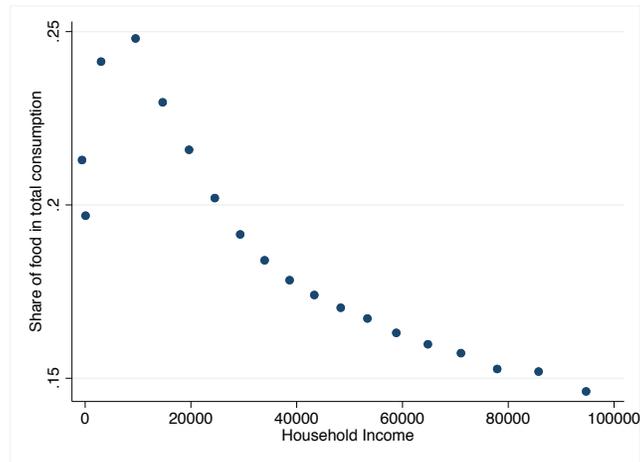
The actual MPC is one with the empirical covariance (i.e. $\frac{\partial C}{\partial E} \frac{dE}{dY} = \bar{\gamma} \overline{MPC} + Cov(\gamma_i, MPC_i)$)

$$MPC^a = \alpha_l (\bar{\gamma} \overline{MPC} + Cov(\gamma_i, MPC_i)) + (1 - \alpha_l \bar{\gamma}) MPC_{nl}$$

These final two expressions are Equations 7 and 8 in the main text.

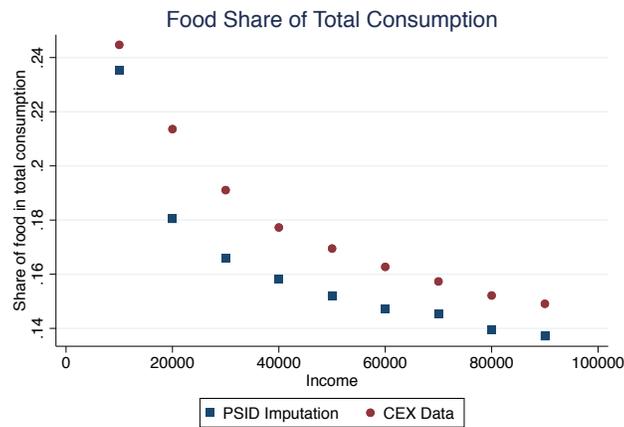
A.8 Appendix figures

Figure A1: The Fraction of Food in Total Spending



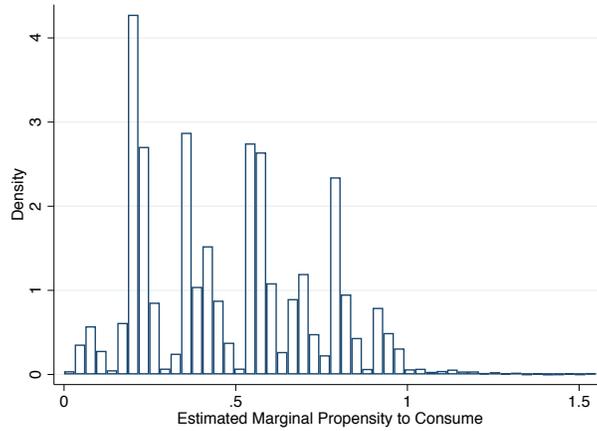
Notes: Data are from the Consumer Expenditure Survey and are pooled across 1984 to 2014 for all households with a head between the ages of 25 and 62. Household income is adjusted to 2010 dollars.

Figure A2: Imputed Total Consumption Using CEX-Based Imputation



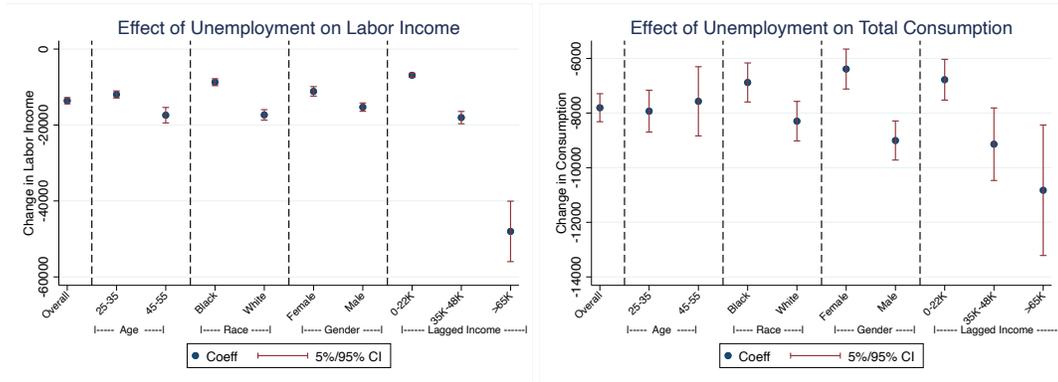
Notes: Average total consumption in the PSID is imputed using Equation A1. Averages are calculated using sample weights in the CEX and based on the nationally representative subsample in the PSID. Income refers to household income, adjusted to 2010 dollars.

Figure A3: The Distribution of Estimated Marginal Propensities to Consume



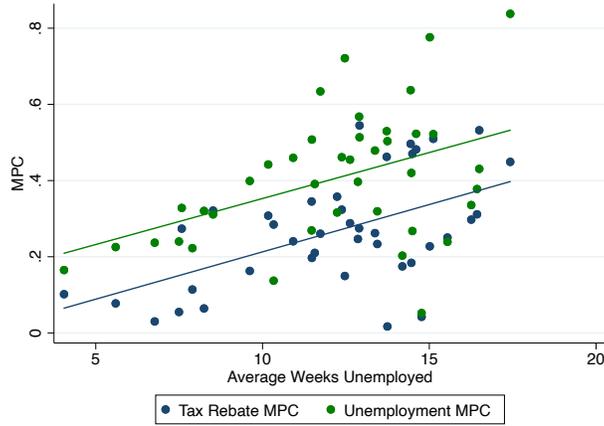
Notes: See Appendix Table A5 for the coefficients that underlie this imputation. Negative imputed MPCs are set to 0. Consumption is measured using total consumption, imputed using the method in Blundell et al. (2008). Income is measured using individual labor income. The instrument for income changes is unemployment. The sample includes the set of workers who were employed two years before the current year. The sample in the PSID excludes observations with more than a 400 percent change in food consumption or income over a given two-year period. Lagged income is measured as the average labor market earnings of the individual in $t - 2$ and $t - 3$. The regression includes year-by-state fixed effects and observations from 1985 to 2013.

Figure A4: Heterogeneity in the Marginal Propensity to Consume: First Stage and Reduced Form



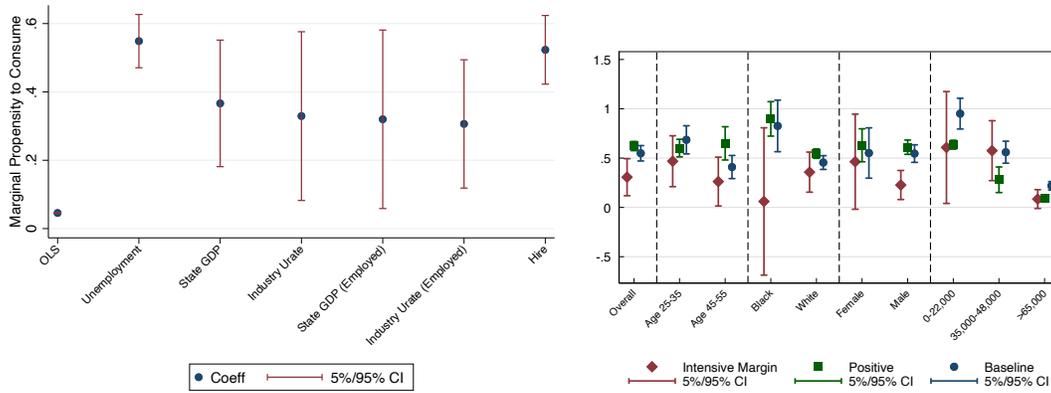
Notes: The left panel shows the first stage of unemployment on the level of labor earnings and the right panel shows the reduced form of unemployment on the level of consumption. These correspond to the instrumented regressions in Figure 2. Consumption is measured using total consumption, imputed using the method in Blundell et al. (2008). Income is measured using individual labor income. The sample includes the set of workers who were employed two years before the current month. The sample in the PSID excludes observations with more than a 400 percent change in food consumption or income over a given two-year period. All regressions include year-by-state fixed effects and observations from 1981 to 2013. Standard errors are clustered at the individual level.

Figure A5: The importance of unemployment duration for MPC heterogeneity



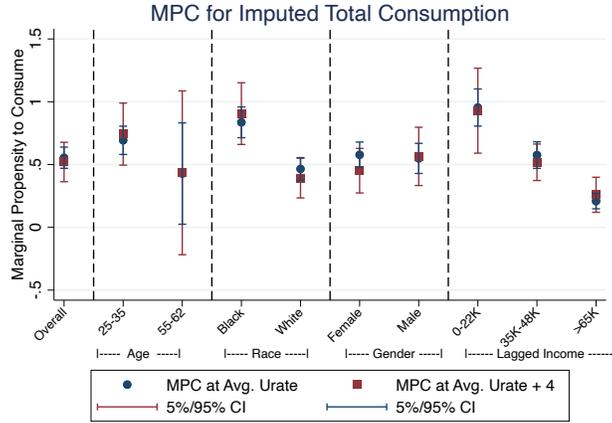
Notes: Each dot is the combination of five household income groups, 2 race groups, and four age groups. Average unemployment duration for each group is calculated as the average number of weeks of unemployment conditional on being unemployed. All averages are unweighted. Sample includes all individuals within the PSID, including both employed and unemployed workers. MPCs for tax rebates are estimated at the household level using five household income bins, four age bins and two race categories (as in Table A11). MPCs for unemployment are calculated as in the main text using individual level income and five lagged income bins, a quadratic in age, female and black indicators, black interacted with age, and female interacted with black.

Figure A6: MPCs Using Different Identifying Income Shocks



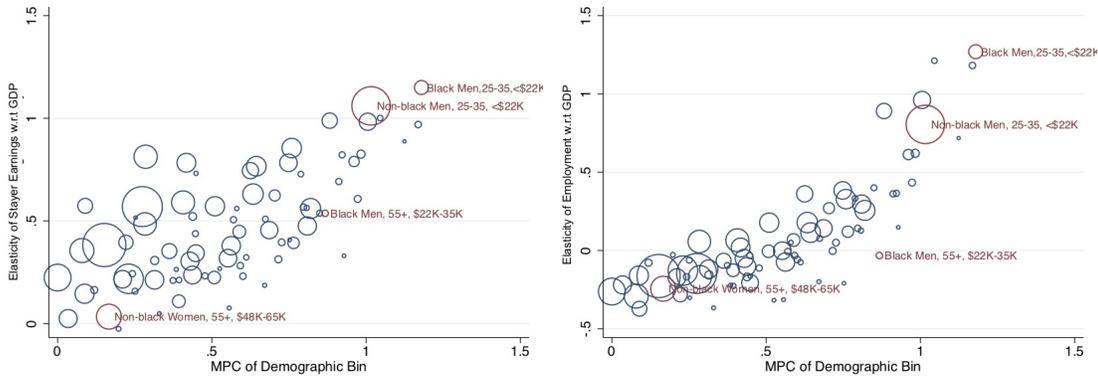
Notes: In the left panel, the instrument labeled “State GDP” is defined as the percentage change in state GDP, defined by the Bureau of Economics Analysis. The industry unemployment rate is calculated from the Basic Monthly Current Population Survey, pooled over months within the year, and defined using time-consistent 1990 census industry codes. In Columns 1, 3, and 4, the sample includes the entire sample (no work restriction); in Column 2, the sample includes those employed in $t - 2$; in Columns 5 and 6, the sample is restricted to those who are employed in $t - 2$ and t ; and in Column 7, the sample is restricted to those who are not employed in $t - 2$. In the right panel, the “baseline” shock is unemployment, the “intensive” series instruments the change in earnings with the unemployment rate in the industry in which the individual worked in $t - 2$, and the “positive” series instruments earnings using hiring between $t - 2$ and t . MPCs estimated using hires restrict the estimation sample to those not employed in $t - 2$. MPCs estimated using unemployment restrict the sample to those who were employed in $t - 2$. MPCs estimated using the industry unemployment rate are estimated without a restriction on employment status in $t - 2$.

Figure A7: The Stability of Marginal Propensity to Consume Estimates Over the Business Cycle



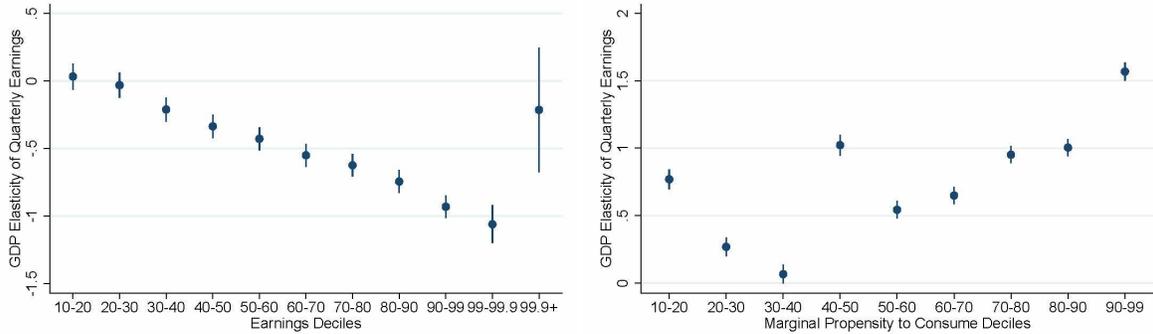
Notes: The unemployment rate is defined as the unemployment rate in the state in which the individual was employed in $t - 2$. Blue dots show the average MPC for the specified bin at the average unemployment rate in the sample. The red squares show the average MPC calculated at the average unemployment rate for each subsample at the average unemployment rate plus 4 percentage points. Regressions are based on two-year periods.

Figure A8: Earnings Sensitivity to GDP and MPCs: Intensive and Extensive Earnings Margins



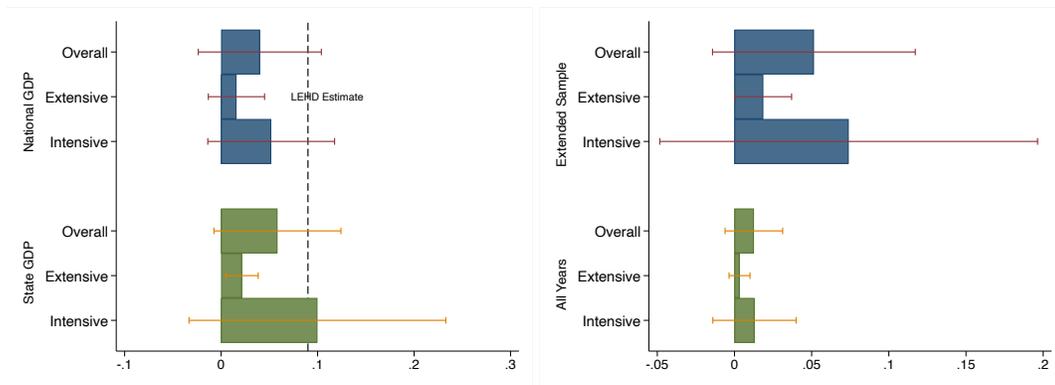
Notes: Sample includes the set of all workers employed in a sample state in year $t - 1$ between 1995 and 2011. The dependent variable in the regression producing the y-axis estimates on the left graph is $\log(E_{i,t}) - \log(E_{i,t-1})$. The dependent variable in the regression producing the y-axis estimates in the right subplot is $L_{i,t}$, where $L_{i,t}$ is an indicator for being employed in time t . The size of each bubble represents the earnings share of that demographic group. The demographic groups are defined as in Figure 1.

Figure A9: Earnings Sensitivity to GDP by Decile of the MPC and Income Distribution



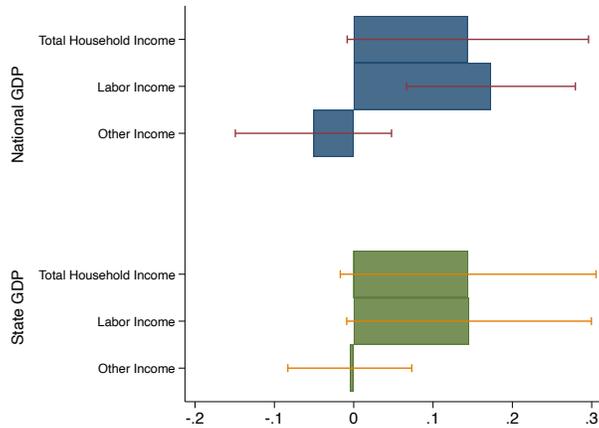
Notes: Sample includes a 5 percent random subset of all workers employed in a sample state in year $t - 1$ between 1995 and 2011. The dependent variable in the regression is $\frac{E_{i,t} - E_{i,t-1}}{.5 * E_{i,t} + .5 * E_{i,t-1}}$. MPC decile bin cutoffs are defined on a sample pooled across all years and each bin represents an equal number of dollars of earnings, rather than individuals. Income deciles are also defined on a sample pooled across all years. Regressions include year fixed effects. Standard errors are clustered at the individual level. Blue bars reflect 95 percent confidence intervals.

Figure A10: PSID-Based Estimates of Covariance



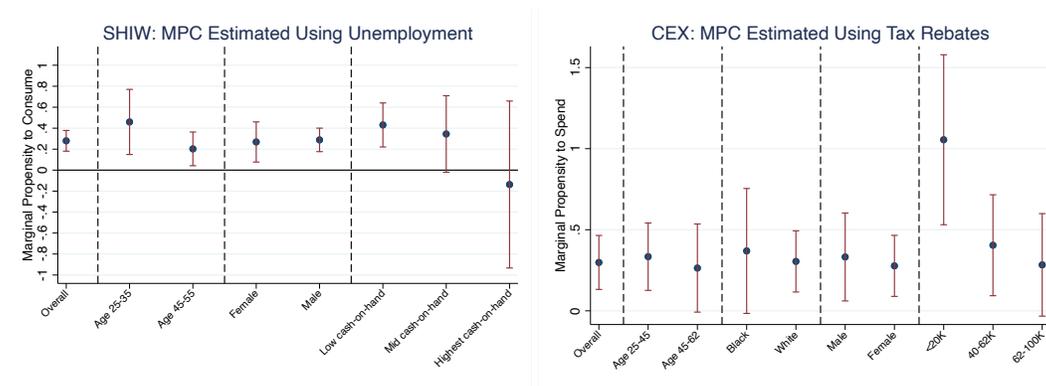
Notes: In all regressions, I restrict attention to the civilian population and drop those serving in the military. All regressions in the right panel use state GDP as the measure of output. In all regressions, observations are weighted by their share of labor income in $t - 2$. In the left panel and the top bars of the right panel, the sample includes all years from 1997-2011. In the bottom of the right panel, the sample includes all years from 1971-2013. In the left panel, the sample is restricted to include those who are employed in year $t - 2$. Standard errors are twoway clustered at the individual and year. Bars reflect 95% confidence intervals.

Figure A11: PSID-Based Estimates of Covariance: Beyond Labor Income



Notes: In all regressions, I restrict attention to the civilian population and drop those households where one of the primary members is serving in the military. In all regressions, households are weighted by the level of household income in $t - 2$. The sample includes all households from 1997-2011. Household labor income is defined as the sum of the labor income of the household head and spouse. Other income is the difference between total household income and the labor income of the household head and spouse. The dependent variable is the change in the log of the income measure in all cases. Standard errors are twoway clustered at the individual and year. Bars reflect 95% confidence intervals.

Figure A12: Cross-demographic Heterogeneity MPCs: SHIW and Tax Rebates



Notes: The left panel shows MPCs estimated upon unemployment using the SHIW. The sample is restricted to those employed in $t - 2$ and consumption is total household consumption, reported directly within the survey. The right panel shows estimates of MPCs for nondurable consumption out of tax rebates using the 2001 and 2008 waves of the CEX.

A.9 Appendix tables

Table A1: Summary Statistics for the LEHD Sample

	1995 Snapshot	2011 Snapshot	Full Sample
Number of workers	22,680,000	45,380,000	38,078,235
Number of establishments	1,015,000	3,464,000	2,772,529
<i>Worker Characteristics</i>			
Fraction Male	0.53	0.51	0.52
Average Worker Age	40.61	43.05	42.10
Average 2-year Lagged Income	36,600	42,210	40,373
Fraction Black	0.09	0.11	0.10
Fraction College Educated	0.30	0.28	0.29
<i>Job Characteristics</i>			
Average Total Quarterly Earnings	10,650	11,920	11,794
Average Annual Earnings	40,140	45,430	44,119
Average No. Jobs per quarter per worker	1.16	1.14	1.15
Fraction with multiple jobs	0.13	0.11	0.12

Notes: Sample includes all individuals in the baseline sample. Averages are unweighted, and nominal values are expressed in 2010 dollars. Column 1 shows the data in 1995, Column 2 shows the data in 2011, and Column 3 shows the sample averaged from 1995 to 2011. Counts for the number of workers and the number of establishments are rounded to comply with U.S. Census disclosure requirements.

Table A2: Years in the Estimation Sample by State

State	Sample Years
Arkansas	2004-2011
Arizona	2006-2011
California	1993-2011
Colorado	1995-2011
Washington DC	2007-2011
Delaware	2000-2011
Florida	1994-2011
Iowa	2000-2011
Illinois	1992-2011
Indiana	2000-2011
Kansas	1995-2011
Maryland	1992-2011
Maine	1998-2011
Montana	1995-2011
New Mexico	1997-2011
Nevada	2000-2011
Oklahoma	2002-2011
Oregon	1993-2011
Pennsylvania	1999-2011
South Carolina	2000-2011
Tennessee	2000-2011
Washington	1992-2011
West Virginia	1999-2011

Notes: Sample years exclude the first two years for which there are individual-level data available. Bold states are those included in the balanced-panel subset of the data.

Table A3: Summary Statistics for the Panel Study of Income Dynamics

	LEHD Comparison Sample	National Sample of Employed	MPC Estimation Sample
<i>Worker Characteristics</i>			
Fraction Male	0.548	0.54	0.514
Average Worker Age	42.7	42.3	41.3
Average 2-year Lagged Income	57,295	57,458	51,153
Fraction College Educated	0.365	0.373	0.308
Fraction Black	0.0469	0.053	0.265
Average Income	60,692	60,312	52,683
Change in Income	3,013	2,734	1,552
<i>Household Characteristics</i>			
Household Size	2.995	3.06	3.134
Food Consumption	9,033	9,163	8,795
Total Consumption	63,052	63,266	58,431
Change in Food Consumption	36.13	59.42	84.12
Change in Consumption	1,020	1,097	1,600
Number of Individuals	12,189	32,832	77,876

Notes: This table shows summary statistics for the PSID sample used in the analysis. The sample in all columns restricts to individuals ages 25 to 62 who are also observed in $t - 2$ and $t + 1$, who have nonmissing changes in food or income over two years, and whose consumption and income change over two years is less than 400 percent. The third column restricts to the set of individuals used in the estimation of MPCs and thus restricts to the set of workers employed in $t - 2$. Column 2 instead restricts to those currently employed who are in the nationally representative subsample of the PSID. Column 1 further restricts to those living in the set of state-years available in the LEHD sample.

Table A4: Summary Statistics for the Consumer Expenditure Survey

	CEX Sample	PSID Sample
<i>Worker Characteristics</i>		
Fraction Male	0.47	0.48
Average Worker Age	43.31	42.30
Fraction College Educated	0.34	0.37
Fraction Black	0.10	0.06
Average Income	45,620	50,775
<i>Household Characteristics</i>		
Household Size	3.16	3.03
Food Consumption	8,366	8,884
Total Consumption	41,725	43,181
Number of Individuals	127,165	97,204

Notes: The first column shows summary statistics for the CEX sample used to impute total consumption. The second column shows the same statistics for the similarly constructed PSID sample. All nominal variables are adjusted to 2010 dollars.

Table A5: Coefficient Estimates for Individual Marginal Propensities to Consume

Dep. Var.	Food Consumption	PSID Imputation	CEX Imputation	CEX Imputation
< 22000*Labor Income	0.0273 (0.1173)	0.2223 (0.2125)	0.9171 (0.7844)	0.9173 (0.7841)
(22,000-35,000)*Labor Income	0.0148 (0.1166)	0.1235 (0.2117)	0.6824 (0.7858)	0.6825 (0.7857)
(35,000-48,000) *Labor Income	-0.0157 (0.1173)	0.0635 (0.2127)	0.5572 (0.7901)	0.5573 (0.7898)
(48,000-65,000) *Labor Income	-0.0309 (0.1191)	0.0047 (0.2155)	0.3471 (0.7878)	0.3473 (0.7875)
> 65,000 *Labor Income	-0.0348 (0.1187)	-0.0305 (0.2143)	0.1887 (0.7936)	0.1888 (0.7933)
Age*Labor Income	0.0027 (0.0059)	0.0045 (0.0106)	-0.0013 (0.0381)	-0.0014 (0.0381)
Age ² *Labor Income	-0.0334 (0.0701)	-0.0223 (0.1264)	0.0415 (0.4464)	0.0416 (0.4462)
Female*Labor Income	-0.0224 (0.0139)	-0.0873 (0.0234)	-0.1347 (0.0706)	-0.1346 (0.0704)
Black*Labor Income	0.0417 (0.0589)	0.2073 (0.1036)	1.2314 (0.4141)	1.2310 (0.4115)
Black*Age*Labor Income	-0.0008 (0.0014)	-0.0042 (0.0025)	-0.0237 (0.0095)	-0.0236 (0.0094)
Black*Female*Labor Income	-0.0052 (0.0296)	0.0128 (0.0515)	-0.2532 (0.1717)	-0.2531 (0.1712)
No. Observations	123439	110403	69788	69788
Year FEs	X	X	X	X
Identifying Shock	Unemployment	Unemployment	Unemployment	Job Loss

Notes: The table shows the regression estimates from PSID imputations. The dependent variable in the first column is total food consumption. The dependent variable in the second column is extended consumption imputed from the later years of the PSID. The dependent variable in the third and fourth columns is total consumption imputed using the CEX data. The fitted values of the regression in Column 3 are plotted in Figure A3. Income is measured using individual labor income. The instrument for income changes is unemployment in Columns 1 through 3 and unexpected job loss in Column 4. The sample includes the set of workers who were employed two years before the current year. The sample in the PSID excludes observations with more than a 400 percent change in food consumption or income over a given two-year period. Lagged income is measured as the average labor market earnings of the individual in t-2 and t-3. All regressions include state-by-year fixed effects. Columns 1 and 2 include years from 1971 to 2013 while Column 3 and 4 includes data from 1985 to 2013. All standard errors are clustered at the individual level.

Table A6: MPCs Estimated Using Job-Level and Geographic Variables

Variables	Baseline	w/ Tenure	w/ Lagged Variance	w/ Ind. Variance	w/ Occ. Variance	w/ Region
Baseline	1.000					
w/ Tenure	0.978	1.000				
w/ Lagged Variance	0.974	0.956	1.000			
w/ Ind. Variance	0.995	0.971	0.968	1.000		
w/ Occ. Variance	0.972	0.953	0.955	0.968	1.000	
w/ Region	0.999	0.977	0.974	0.994	0.972	1.000

Notes: The table shows the pairwise correlation between the baseline MPC estimate and estimates including the stated additional characteristic. The dependent variable is total consumption imputed using the CEX data. All regressions include state-by-year fixed effects and include individuals employed in time $t - 2$. The expected variance of earnings at the industry or occupation level are calculated using the matched CPS monthly data and averaged over the sample period. These capture the average within-individual variance in earnings over a one-year period. Tenure is defined as the number of months the worker has been with the firm in which they were employed in time $t - 2$. The lagged variance is the variance of an individual's earnings between $t - 3$ and $t - 2$. See Appendix text for more details on data construction.

Table A7: Robustness of Relationship Between MPCs and Earnings Elasticities: Alternate Specifications

	(1)	(2)	(3)	(4)	(5)
	Baseline Estimate	log(Eit +100)	Annual Income	1993 State Subsample	State GDP
$MPC_{i,t-1}$	-0.229 (0.049)	-0.362 (0.116)	-0.080 (0.013)	-0.246 (0.049)	-0.221 (0.048)
$MPC_{i,t-1} * \Delta Y_t$	1.228 (0.383)	2.550 (0.829)	1.091 (0.249)	1.540 (0.459)	0.837 (0.219)
No. Observations	29,204,700	29,204,700	29,204,700	20,859,200	29,204,700
R-Squared	0.009	0.004	0.003	0.010	0.012
$Var(MPC_i)$	0.083	0.080	0.080	0.084	0.083
$Cov(MPC_i, \gamma_i)$	0.099	0.200	0.084	0.126	0.068
<i>Fixed Effects</i>					
Year	✓	✓	✓	✓	
Year by State					✓

Notes: Column 1 shows regression estimates from Table 2, estimated on a 5 percent subsample of the data set. Column 2 specified the outcome variable as the change in the log of earnings plus 100. Column 3 defines income as the annual income in the calendar year, rather than quarterly income in the fourth-quarter. This analysis, however, still restricts the sample to the set of individuals employed in the fourth quarter of the previous year. Column 4 restricts to the subsample of states that are present in 1993, meaning that there is a balanced panel of states over time. Column 5 replaces aggregate GDP with state-level GDP. This specification includes state-by-year fixed effects, rather than simply year fixed effects. Across all columns, the number of observations is rounded to comply with U.S. Census Bureau disclosure requirements. In all columns, coefficients and standard errors are calculated using multiple imputation techniques, as explained in the main text.

Table A8: Robustness of Relationship Between MPCs and Earnings Elasticities: Alternate MPC Estimates

<i>MPC Definition:</i>	(1) Baseline Estimate	(2) Food	(3) Non-durables	(4) PSID-Based Consumption	(5) Hires Instrument
$MPC_{i,t-1}$	-0.229 (0.049)	-2.060 (0.750)	-0.306 (0.072)	-0.821 (0.168)	-0.260 (0.110)
$MPC_{i,t-1} * \Delta \log Y_t$	1.228 (0.383)	12.470 (4.536)	1.755 (0.543)	3.945 (1.250)	1.086 (0.509)
No. Observations	29,204,700	29,204,700	29,204,700	29,204,700	29,204,700
R-Squared	0.009	0.006	0.008	0.013	0.002
Avg. MPC_i	0.431	0.029	0.210	0.169	0.246
$Var(MPC_i)$	0.083	0.001	0.041	0.009	0.051
$Cov(MPC_i, \gamma_i)$	0.099	0.009	0.070	0.034	0.051
<i>Fixed Effects</i>					
Year	✓	✓	✓	✓	✓

Notes: Each column defines the marginal propensity to consume in a different way. Column 1 is the baseline MPC estimate. Column 2 uses the MPC defined only using food consumption. Column 3 uses the CEX imputation but includes only nondurable consumption. Column 4 uses the MPC calculated using the PSID-based imputation of expanded consumption. Column 5 uses the MPC estimated using hires, rather than unemployment, as the instrument for changes in income. The outcome variable in all regressions is the annual change in quarterly earnings across all jobs. All regressions include year fixed effects, and all regressions are estimated using multiple imputation techniques, as defined in the main text. Across all columns, the number of observations is rounded to comply with U.S. Census Bureau disclosure requirements.

Table A9: Amplification in Recessions versus Expansions

	(1) MPC from unemp.	(2)	(3) MPC from hires
$MPC_{i,t-1}$	-0.221 (0.048)	-0.215 (0.048)	-0.247 (0.109)
$MPC_{i,t-1} * \Delta GDP_{s,t}$	0.837 (0.219)	0.663 (0.243)	0.660 (0.354)
$MPC_{i,t-1} * \Delta GDP_{s,t} * Rec_{s,t}$		0.576 (0.398)	0.309 (0.465)
No. Observations	29.2 M	29.2 M	29.2 M
\overline{MPC}_i	0.43	0.43	0.246
$Var(MPC_i)$	0.083	0.083	0.051

Notes: Regressions include year*state fixed effects and are estimated on a 5% subsample of the data. All coefficients and standard errors are calculated using multiple imputation techniques as described in the main text.

Table A10: Earnings Elasticities by MPCs: Evidence from Italy

<u>Panel A. Household Income</u>				
	Reported MPC	2010 MPC	Fitted MPC	Unemployment MPC
MPC * Change in GDP	0.822 (0.345)	0.648 (0.661)	22.987 (5.854)	3.067 (0.804)
MPC	-0.048 (0.011)	-0.023 (0.018)	0.761 (0.310)	0.052 (0.050)
Change in GDP	0.092 (0.394)	0.461 (0.511)	-9.816 (2.596)	0.074 (0.170)
Average MPC	.45	.45	.45	.12
Std. Dev. MPC	.34	.35	.03	.21
Observations	7497	12720	20095	20095
R-Squared	0.002	0.004	0.016	0.012
<u>Panel B. Individual Labor Income</u>				
	Reported MPC	2010 MPC	Fitted MPC	Unemployment MPC
MPC * Change in GDP	0.669 (0.382)	0.876 (0.239)	7.937 (3.201)	0.309 (0.443)
MPC	-0.037 (0.019)	-0.033 (0.012)	0.158 (0.222)	0.019 (0.029)
Change in GDP	0.428 (0.353)	0.290 (0.190)	-3.013 (1.389)	0.542 (0.064)
Average MPC	.45	.46	.46	.17
Std. Dev. MPC	.34	.35	.03	.22
Observations	7419	13060	19705	19705
R-Squared	0.005	0.005	0.005	0.004

Notes: All observations are weighted by their income in the beginning of the period. For panel A, this means weighting by total household income in t-2 and in panel B, this means weighting by individual payroll in t-2. MPCs range from 0 to 1. With the exception of the MPC in column 1, all MPCs refer to the MPC in the initial period. All standard errors are twoway clustered at the household and year.

Table A11: Elasticity of Income by MPS Estimated Using Tax Rebates

	(1)	(2)	(3)	(4)	(5)	(6)
MPC	0.128 (0.107)	0.111 (0.129)	-0.093 (0.103)	0.044 (0.134)	-0.009 (0.134)	-0.119 (0.084)
MPC * Change in Aggregate GDP	1.351 (1.243)	0.488 (1.042)	0.951 (0.852)			
MPC * Change in State GDP				2.845 (2.594)	1.968 (2.478)	1.468 (1.319)
Earnings Definition	HH Labor	Total HH	Ind. Labor	HH Labor	Total HH	Ind. Labor
Average MPC	.191	.192	.192	.191	.192	.192
Std. Dev. MPC	.122	.123	.122	.122	.123	.122
Observations	25476	25182	39171	22702	22154	34755
R-Squared	0.282	0.272	0.090	0.294	0.287	0.098

Notes: Households are weighted by their household's income in t-2. All regressions include year*state fixed effects. The dependent variable in all columns is the 2-year change in the log of the stated income variable. For Columns 1 and 4, this income variable is the sum of labor income for the head an spouse. For Columns 2 and 5, the income variable is total household income. For columns 3 and 6, the income variable is individual labor income. Columns using household income measures use only 1 observation per household while columns using individual income use 1 observation per head and one per spouse, if present. Sample is restricted to include 1997-2011. Standard errors are twoway clustered at the individual and year level in columns 1 through 3 and at the individual and state*year level in columns 4 through 6.

Table A12: Unemployment Probabilities over the Business Cycle

	LEHD Sample	Extended Sample	Non-employment	All Years
MPC	0.084 (0.010)	0.160 (0.032)	0.566 (0.039)	0.123 (0.017)
MPC * Change in Aggregate GDP	-0.539 (0.180)	-0.731 (0.281)	-0.781 (0.325)	-0.236 (0.114)
Observations	29422	34977	42137	108663
R-Squared	0.034	0.040	0.127	0.042

Notes: All regressions use MPCs estimated using CEX imputation. All regressions include state*year fixed effects. Changes are defined across 2-year periods. The LEHD sample restricts the sample to 1997-2011 and restricts attention to the set of workers employed in year t-2 in non-government industries. Since unemployment is measured in the month of the survey, which is mostly concentrated in March, the change in GDP is taken from Q1 in each year. Standard errors are twoway clustered at the individual and year level.

Table A13: Employment Cyclical and the Local Covariance: Robustness

	Baseline		Financial Controls	Bartik Shock	CZ-level Variation
$\widehat{Cov}_c * \Delta \log Y_t$	4.858 (1.098)		3.966 (1.203)	2.171 (1.543)	1.887 (1.119)
$\widehat{Cov}_c * \Delta \log Y_t * Tradable_i$	-5.710 (2.297)	-4.741 (2.114)	-5.431 (2.589)		
$\widehat{MPC}_c * \Delta \log Y_t$	-2.100 (2.407)		-1.304 (2.893)	3.736 (1.409)	0.033 (2.479)
$\widehat{MPC}_c * \Delta \log Y_t * Tradable_i$	5.080 (1.271)	5.080 (1.271)	5.489 (1.444)		
Demographic Controls	X	X	X	X	X
CZ*Year FE		X			
Financial Controls			X		
Included Industries	T&N	T&N	T&N	All	-
No. Observations	130000	130000	96000	332000	2000
R-Squared	0.319	0.339	0.348	0.370	0.746

Notes: Each regression is calculated using multiple imputation techniques using 100 draws of MPC estimates. Regressions in columns (1)-(4) include an unbalanced panel of 270 commuting zones and 313 NAICS codes from 2001 to 2011. Column (5) includes an unbalanced panel of 270 commuting zones from 2001 to 2011. In columns 1 through 3, all controls are included separately, interacted with GDP, and triple interacted with GDP and a tradable indicator. Column 4 shows results using the Bartik shock rather than GDP. See Appendix A.6 for details on the construction of the shock. In each regression, observations are weighted by the share of employment in $t - 1$. Observations are rounded to the nearest 1000.

Table A14: Mobility Patterns: Cross-Sector

Marginal Propensity to Consume (MPC)	-0.0036 (0.0002)		-0.0032 (0.0004)
$\Delta \log GDP$		0.0087 (0.0043)	
MPC * $\Delta \log GDP$			-0.0412 (0.0301)
No. Observations	11104884	11104884	11104884

Notes: Data are from the Basic Monthly Current Population Survey from 1990 to 2011. The dependent variable is an indicator for moving from employment in the private sector to employment in self-employment, the military, or federal employment. MPC is imputed using PSID estimates based on age, gender, and race. The sample includes all adjacent periods in which an individual is employed. Standard errors are clustered at the individual level. Column 1 includes quarter fixed effects, and all other columns include year-by-month fixed effects.

Table A15: Earnings elasticities in the PSID: sensitivity to state-level moves

	Combined Earnings		Extensive Margin		Intensive Margin	
	Actual	LEHD Sample	Actual	LEHD Sample	Actual	LEHD Sample
MPC	-0.124 (0.076)	-0.122 (0.076)	-0.073 (0.031)	-0.073 (0.031)	-0.102 (0.130)	-0.104 (0.132)
MPC * Annual change in GDP	1.288 (0.674)	1.271 (0.668)			1.448 (1.026)	1.463 (1.043)
MPC * Q1-Q1 change in GDP			0.550 (0.236)	0.550 (0.235)		
Observations	11617	11641	12563	12563	11166	10982
R-Squared	0.078	0.234	0.039	0.320	0.094	0.087

Notes: All regressions use MPCs estimated using CEX imputation. All regressions include state*year fixed effects. Changes are defined across 2-year periods. The LEHD sample restricts the sample to 1997-2011 and restricts attention to the set of workers employed in year t-2 in non-government industries. The sample is further restricted to the workers living in the LEHD sample states in period t-2. Since employment is measured in the month of the survey, which is mostly concentrated in March, the change in GDP for columns 3 and 4 is taken from Q1 in each year. Since income is measured for the year, the change in GDP in columns 1,2,4 and 5 is calculated as the annual average. All observations are weighted by their share of period t-2 earnings. Standard errors are twoway clustered at the individual and year*state.

B Online Appendix: Matching Multiplier in a calibrated model

This section explores the matching multiplier mechanism in a dynamic setting. This is particularly important as I estimate MPCs using unemployment as the identifying income shock. Unemployment is a persistent shock, and therefore, the MPCs that I estimate include both the response of consumption today to income today, but also the response of consumption today to expected income shocks in all future periods. The empirical matching multiplier accounts for these dynamics only insofar as they are embedded in the empirical MPC estimate, but does not formally capture the role of these dynamics in the multiplier. Therefore, the empirical matching multiplier is a reduced-form approximation to a more complicated dynamic process.⁴⁶

I explore this mechanism within a standard Bewley-Huggett-Aiyagari model augmented along three dimensions. First, I introduce endogenous labor supply and rich consumer heterogeneity. Second, I pair this model of aggregate demand with fixed wages in the short run, which capture the role that this mechanism plays in demand-driven amplifications. Third, I introduce an exogenous labor rationing process that generates labor income fluctuations in the presence of fixed wages. Within the context of this model, I clarify the role that the persistence of the unemployment shock plays in shaping my estimates for the importance of the matching multiplier mechanism. I show that the persistence of the unemployment shock is important in determining the level of the multiplier and matching multiplier, but is not the driving force for the amplification coming from the covariance. Additionally, I show that the empirical approximation from Section 2 closely captures the dynamic estimates under the assumption that the aggregate shock has the persistence of the average unemployment spell.

B.1 Environment

The following setting is similar to the generalized setting in Auclert et al. (2018) and is a simplified version of the multisector model in Flynn et al. (2021). Since the focus of this exercise is to understand the role of heterogeneity in the demand side, I allow for rich heterogeneity among consumers and keep the supply side intentionally simple. Consider an economy in discrete time with T periods. The economy is populated by a continuum of agents of I types, where each type i has a mass ν_i of individuals such that $\sum_i \nu_i = 1$. Households within each group are ex ante homogenous, but households face idiosyncratic risk in their productivity or labor supply $e(s)$, a process that may vary across demographic groups I . Households across groups may differ both in the income process they face and in their discount rates (β_i). Households have preferences over both consumption and leisure, and agents have the ability to borrow and save into a real asset ($a_{i,t}$) to smooth consumption but are subject to a borrowing constraint that $a_{i,t} \geq b$. The household's problem therefore is to choose paths for their consumption $c_{i,t}$ and labor supply $l_{i,t}$ to maximize their utility taking wages, prices, and the real interest rate as given:

$$\max_{c_{i,t}, l_{i,t}} \sum_{t=0}^T \beta_i^t E[u(c_{i,t}) - v(l_{i,t})] \quad (\text{B1})$$

subject to

$$w_t l_{i,t} e(s) + r_t a_{i,t-1} - \tau_{i,t} = p_t c_{i,t} + a_{i,t} \quad , \quad a_{i,t} \geq b \quad (\text{B2})$$

where $\tau_{i,t}$ are lump sum taxes, b is the borrowing constraint, and p_t is the price of the final good at time t . I assume that $u(c) = \frac{c^{1-\omega}}{1-\omega}$ and $v(l) = \frac{l^{1+\psi}}{1+\psi}$, where ω is the intertemporal elasticity of substitution and

⁴⁶Alternatively, the empirical matching multiplier could be interpreted as the multiplier in which period 1 is the short run and period 2 is the long run.

ψ is the Frisch labor supply elasticity.⁴⁷

While the demand side of the model features rich heterogeneity, the supply side of the model is simple.⁴⁸ All workers are employed by a representative competitive firm, which produces with constant returns to scale technology and takes labor as the only input:

$$Y_t = L_t \quad \forall t \quad (\text{B3})$$

Firm profit maximization implies that

$$w_t = p_t \quad \forall t \quad (\text{B4})$$

The government sets potentially individual-specific lump sum taxes $\tau_{i,t}$ to finance government spending G_t . Rather than balance its budget strictly between periods, the government can issue bonds B_t to smooth fluctuations across periods and therefore is subject to an intertemporal budget constraint:

$$\sum_t \frac{\tau_{i,t}}{\prod_{i \leq t} (1 + r_i)} = \sum_t \frac{p_t G_t}{\prod_{i \leq t} (1 + r_i)} \quad (\text{B5})$$

I assume that government spending preferences are given exogenously by θ_G , such that $G_t = G(r_t, \tau_t, \theta_G)$. Even though this fiscal rule is specified exogenously, government spending still responds to interest rate changes in order to maintain the budget constraint in Equation B5.

B.2 Equilibrium and the output multiplier

Consider first the case where all prices are fully flexible. The household problem in Equation B1 results in a demand for consumption and a labor supply function given by

$$c_{i,t} = c_i(\{\lambda_t\}_{t \in T}, \{\tau_{i,t}\}_{t \in T}, \beta_i, b) \quad (\text{B6})$$

$$l_{i,t} = l_i(\{\lambda_t\}_{t \in T}, \{\tau_{i,t}\}_{t \in T}, \beta_i, b) \quad (\text{B7})$$

where $\lambda_t = \{r_t, w_t, p_t\}$ is the vector of prices. These are Marshallian demands, and thus these functions only depend directly on exogenous parameters $(\beta_i, b, \tau_{i,t})$ and prices (λ_t) . The goods and labor market clearing condition are given, respectively, by

$$Y_t = C_t + G_t = \sum_i \nu_i c_{i,t} + G_t \quad \forall t \quad (\text{B8})$$

$$L_t = \sum_i \nu_i l_{i,t} e(s) \quad \forall t \quad (\text{B9})$$

An allocation of $\{c_{i,t}, l_{i,t}, \tau_{i,t}, r_t, p_t, w_t, G_t, Y_t\}$ that satisfies Equations B4, B5, B6, B7, B8, and B9 characterizes the flexible price equilibrium.

⁴⁷ These functional forms are used for the quantitative exercise in this section. However, the results in this section apply to a broader set of preferences. See Auclert (2019) or Flynn et al. (2021) for a discussion using more general preferences. These preferences have the advantage that they guarantee that the resulting labor supply and Marshallian demand functions are continuous and differentiable in r_t . The assumed CRRA utility function also exhibits sufficient diminishing marginal utility of consumption to guarantee the existence of an equilibrium.

⁴⁸ In a similar framework, Auclert et al. (2018) explore the importance of worker MPCs in a model with an enriched supply side that includes capital, sticky prices, and a Taylor rule for monetary policy. They show that while these modifications reduce the overall size of the multiplier, worker MPCs still remain crucial in determining the output response to fiscal policy. Specifically, they show that in a model that matches the empirical estimates of intertemporal MPCs and with deficit-financed spending, impact multipliers can be above 1, even with active monetary policy, distortionary taxation, and investment crowd out.

The exercise in this paper will be to consider the response of an economy, initially at this flexible price equilibrium, to an unanticipated demand shock when wages are fixed for $k + 1$ periods.⁴⁹ Equation B4 immediately implies that prices are also fixed over this period. With fixed wages, the interest rate does not adjust to clear the labor market, and thus, workers are off their labor supply curves for the first k periods (i.e., in response to a negative shock, there are workers who would like to work more but cannot because a firm is not willing to hire them). Rather, in those periods, labor supply is rationed, and a worker's labor supply is imposed exogenously as

$$l_{i,t} = n_{i,t}(Y_t) \quad (\text{B10})$$

such that $\sum_i \nu_i n_{i,t} = L_t$.⁵⁰ This rationing function takes as inputs the aggregate change in output and the amount workers would like to work given the wage $l_{i,t}^*$ and returns a labor supply $n_{i,t}$ for each individual. This function is what determines the change in the worker's earnings in response to an aggregate demand shock. This reduced form specification, similar to Werning (2015), captures the notion that, for example, in response to a negative demand shock, workers are not able to work as much as they would like. In order to capture the relationship between the exposure of worker earnings to aggregate shocks and worker MPCs documented in Section 4.2, I parametrize the rationing function $N(Y_t) = \{n_{i,t}\}$ as

$$n_{it} = \frac{Y_t}{L_t^*} (1 - \chi \overline{MPC} + \chi MPC_i) l_{it}^* \quad (\text{B11})$$

where $L_t^* = \sum_i \nu_i l_{i,t}^*$, \overline{MPC} is the earnings-weighted average MPC in the economy and χ is the slope of this incidence function with respect to GDP. This formulation of the rationing function implies that the elasticity of worker i 's earnings to the aggregate is linear in the worker's MPC and is given by $\gamma_i = 1 - \chi \overline{MPC} + \chi MPC_i$.

Since the interest rate is not pinned down by the labor market clearing condition, it must be set by monetary policy. Assume for simplicity that the central bank targets a fixed real interest rate, $r_t = \bar{r}$.⁵¹ In this rationing equilibrium, consumer demand becomes

$$c_{i,t} = c_i(\{y_{i,t}\}_{t \leq k}, \{\lambda_t\}_{t \in T}, \{\tau_{i,t}\}_{t \in T}, \beta_i, b) \quad (\text{B12})$$

This is similar to the flexible price condition (Equation B6), except that now it is a function of incomes in periods 1 through k , as these are now *exogenously* given by the rationing function.

Definition *The rationing equilibrium is defined as the set of $\{c_{i,t}, l_{i,t}, \tau_{i,t}, r_t, p_t, w_t, G_t, Y_t\}$ such that firms optimize as in Equation B4, consumers optimize consumption according to Equation B12 and supply labor according to $N(Y_t, l_{i,t-1})$ for periods 1 through k and according to Equation B7 for $t \geq k$, and goods and labor markets clear in each period as in Equations B8 and B9.*

Using bold variables to represent vectors of aggregate variables (i.e. $\mathbf{Y} = \{Y_t\}$), I derive the response of the economy to shocks in Proposition 1. I define the partial equilibrium effect of the shock on output as the response of the economy to a shock to any of the parameters of the model, before accounting for any

⁴⁹ See Auclert (2019) for the case where wages are sticky indefinitely.

⁵⁰ A more complete alternative to this rationing function is to explicitly model heterogeneity in the labor market through search frictions. See Ravn and Sterk (2020) for a recent example. Additionally, while the resulting mechanism is similar, the endogenous redistribution mechanism here differs from that in Bilbiie (2008) and Bilbiie (2020). In his setting, cyclical inequality comes from the redistribution of firm profits – the government taxes firm profits (held by the unconstrained agents) and rebates them lumpsum to the constrained agents. I allow for a rationing function that is reduced form but disciplined by labor market data.

⁵¹ Since my focus is on quantifying the importance of heterogeneity in the labor market, I abstract from potential offsetting effects coming from countervailing monetary policy.

of the general equilibrium responses of the economy. In this economy, this amounts to all effects on the economy *before* incomes or interest rates change. Let $\partial \mathbf{Y}$ be the vector of the partial equilibrium change in output in each period. Define C_Y to be a matrix where the k, j entry is given by $\frac{dc_k}{dY_j} = \sum_i \frac{dc_{i,k}}{dy_{i,j}} \gamma_i \frac{l_{i,j}}{L_j}$, which is the aggregate response of consumption at time k to income in time j .

Proposition 1. *Under the assumption that wages are sticky for $k+1$ periods, for any shock to parameters $(\beta_i, \tau, \theta_G)$, the total change in output from an initial flexible price allocation is given to first order by:*

$$d\mathbf{Y} = (I - \mathbf{C}_Y J_k - (\mathbf{C}_r + \mathbf{G}_r) J_{T-k} (\mathbf{L}_r)^{-1})^{-1} \partial \mathbf{Y} \quad (\text{B13})$$

where subscripts denote partial derivatives (i.e. \mathbf{C}_r is the partial derivative of consumption with respect to r), and J_k and J_{T-k} are diagonal matrices with 1s in the first k or the last $T-k$ entries, respectively.

Proof. Begin by totally differentiating the good market clearing condition (Equation B8) in each period t :

$$dY^t = \sum_{j=1}^k C_{t,y_j} dy_j + \sum_{j=1}^T (C_{r_j} + G_{r_j}) dr_j + \sum_{j=1}^T C_{t,\tau} d\tau_t + C_{t,\beta} d\beta + \sum_{j=1}^T G_{t,\tau_t} d\tau_t + G_{t,\theta_G} d\theta_G \quad (\text{B14})$$

where $C_{t,x}$ is an $(1 \times I)$ vector across individuals where each entry is the partial derivative of the individual consumption function $c_i(\{y_{i,t}\}_{t \leq k}, \{\tau_{i,t}\}_{t \in T}, \{r_t\}_{t \in T}, \beta_i, b)$ with respect to the variable x . Similarly, derivatives $dy_j, dr_j, d\beta, d\tau$ and $d\theta_G$ are $(I \times 1)$ vectors that capture the change in the exogenous parameters for each individual i . Recall that in the rationed equilibrium, income is exogenous in all rationed periods, and thus enters the consumption function. Note that the first sum is only across periods 1 through k , the periods in which there is labor market rationing. Beyond that, the workers are back on their labor supply curves and their income is endogenously given by their decisions and prices. By the definition of the income process imposed by the rationing function in Equation B11,

$$dy_{i,t} = n_{i,t} - l_{i,t-1} = \gamma_i \frac{l_{i,t-1}}{L_t} dY_t$$

Denote N_y as the I vector where the i entry is $dy_{i,t} = \gamma_i \frac{l_{i,t-1}}{L_t}$. Plugging this in, we get:

$$dY^t = \sum_{j=1}^k (C'_{t,y_j} N_y dY_j) + \sum_{j=1}^T (C_{r_j} + G_{r_j}) dr_j + \sum_{j=1}^T C_{t,\tau} d\tau_t + C_{t,\beta} d\beta + \sum_{j=1}^T G_{t,\tau_t} d\tau_t + G_{t,\theta_G} d\theta_G \quad (\text{B15})$$

Note that $C'_{t,y_j} N_y = \frac{dc_t}{dy_j} = C_{t,j}$, which is the aggregate response in time t to a change in income at time j . Equation B15 holds for all periods t , and stacking equations, this becomes

$$dY = C_Y J_k dY + (C_r + G_r) dr + C_\tau d\tau + C_\beta d\beta + G_\tau d\tau + G_{\theta_G} d\theta_G \quad (\text{B16})$$

where C_Y is a matrix where the m, n entry is the aggregate consumption response at time m to an income shock in time n and J_k is a diagonal matrix with 1s only for the first k periods. For the first k periods, the interest rate is pinned down by the monetary policy rule (rather than by labor market clearing). Given the rule specified,

$$dr_t = 0 \quad \forall t \leq k$$

After k periods, the interest rate goes back to the flexible price scenario and the change in the interest rate is pinned down by the labor market clearing condition. The total derivative of the labor market clearing condition is given by:

$$dY_t = \sum_{j=0}^T L_{t,r_j} dr_j + \sum_{j=1}^T L_{t,\tau} d\tau_t + L_{t,\beta} d\beta \quad \forall t > k$$

Stacking across periods and solving for dr , we get the expression for the change in the interest rate at time $t > k$:

$$dr = (L_r)^{-1} (dY - L_\tau d\tau - L_\theta d\theta)$$

Stacking these equations over time and defining ∂Y as a $(1 \times T)$ vector that is the change in output in each period before incomes and interest rates have been allowed to adjust:

$$\partial Y = C_\tau d\tau + C_\beta d\beta + G_\tau d\tau + G_{\theta_G} d\theta_G + (C_r + G_r) J_{T-k} (L_r)^{-1} (L_\tau d\tau - L_\beta d\beta)$$

we can rewrite Equation B16 as

$$dY = C_Y J_k dY + (C_r + G_r) J_{T-k} (L_r)^{-1} dY + \partial Y \quad (\text{B17})$$

where dY and ∂Y are $T \times 1$ matrices and C_Y is a $T \times T$. J_k is just a diagonal matrix with ones along the diagonal for the first k entries, and J_{T-k} is a diagonal matrix with 1s for the last $T - k$ entries. C_r, G_r are $T \times T$ matrices. r_y are $T \times T$ matrices. \square

The first term (i.e. $C_y J_k$) captures the heterogeneous agent intertemporal version of the traditional Keynesian multiplier and embeds the mechanism that is the key focus of this paper. This matrix of intertemporal MPCs features prominently in Auclert et al. (2018), who argue in a similar setting that these moments are essential for determining general equilibrium effects in heterogeneous agent models. Due to the forward-looking nature of the consumer's problem, what matters for the total consumption response today is not just the change in today's income but also the change in future period incomes. The matrix J_k simply captures the fact that wages are only fixed for some period; thus, the consumer only responds directly to income changes in those periods. The heterogeneous incidence of labor shocks directly affects the magnitude of C_Y – when high-MPC workers have higher γ_i , the components of the C_Y matrix are larger.

The second term (i.e. $(C_r + G_r) J_{T-k} (L_r)^{-1}$) captures the movements in the interest rate in periods after k . In those periods, the interest rate will adjust to bring workers back onto their labor supply curves and clear the labor market. Workers anticipate the future adjustment, and thus, consumption today will depend on the future change in the interest rate. Proposition 1 shows that this matrix is a sufficient statistic for characterizing the first-order general equilibrium effect on output of a demand shock.

Using Proposition 1, I define the matching multiplier in Corollary 1 as a special case of the model that abstracts from the potentially offsetting price effects far in the future.

Corollary 1. *The Matching Multiplier is defined as the difference in the output response to any small unanticipated shock to parameters (β, τ, θ_G) between the actual case where γ_i varies across i and the case where $\gamma_i = 1$. Under the assumption that $C_r + G_r = 0$ and under the assumption that wages are sticky for $k + 1$ periods, this is given by:*

$$MM = \left[(I - C_Y(\gamma_i)) J_k \right]^{-1} - \left[(I - C_Y(\gamma_i = 1)) J_k \right]^{-1} \quad (\text{B18})$$

where subscripts denote partial derivatives (i.e. C_Y is the partial derivative of consumption with respect to y) and J_k is a diagonal matrices with 1s in the first k entries. $C_Y(\gamma_i)$ denotes the aggregate intertemporal MPC matrix when γ_i varies by individual and $C_Y(\gamma_i = 1)$ reflects the case where $\gamma_i = 1$ for all individuals.

A comparison of Corollary 1 with the empirical matching multiplier derived in Section 2 demonstrates that the empirical moment from Equation 4 is a reduced form approximation of the multiplier in this dynamic setting. Recall that the k, j entry is given by $\frac{dc_k}{dy_j} = \sum_i \frac{dc_{i,k}}{dy_{i,j}} \gamma_i \frac{l_{i,j}}{L_j}$, which is the aggregate response of consumption in period k to an income shock in period j . Were the MPC estimates (\widehat{MPC}_i) in Section 4.1 based on a purely transitory shock, then $\widehat{MPC}_i = \frac{dc_{i,1}}{dy_{i,1}}$ and the empirical matching multiplier from Section 2 would be exactly the sufficient statistic in Proposition 1 in the case that wages were only sticky for 1 period. However, as discussed above, the MPCs that I estimate capture the consumption response to unemployment, which is a more durable shock, and therefore $\widehat{MPC}_i \neq \frac{dc_{i,1}}{dy_{i,1}}$. In the dynamic setting, there is no direct analytical mapping between the empirical covariance and the matching multiplier defined in Corollary 1. Therefore, the exercise in the following sections will be to calibrate the dynamic model below using the average empirical MPC and the estimated covariance between MPCs and earnings elasticities as targeted moments and explore the matching multiplier in the dynamic setting.

B.3 Model calibration

In order to match the empirical exercise, where I consider heterogeneity in MPCs and earnings sensitivities across both demographic and income groups, I calibrate an economy with eight demographic groups characterized by the combination of two genders, two education bins, and two race bins. Within each demographic group, agents are ex ante homogenous, but across demographic groups, agents differ along several ex ante dimensions.

First, and most importantly, I allow each demographic group to have a different income process that captures differences in the overall riskiness of income, features earnings shocks that mirrors the unemployment shock that I use to estimate the MPC in Section 4.1, and allows for heterogeneity in the persistence of that unemployment shock. I adopt a 2-part income process, where total earnings are given by

$$Y_{it} = e_{it} * w_{it}$$

where e_{it} is scalar for the number of months employed and w_{it} are the monthly earnings among the employed. I assume that these two processes are independent and calibrate them separately for each group.

I model the unemployment process as a discrete-time Markov process with two states – employed for the whole year ($e_{it} = 1$) and unemployed for some fraction of the year ($e_{it} = X$, where X is the average fraction of earnings lost in the year of unemployment). The transition probabilities between these two states are determined by the job-finding (f) and job-separation (s) rates, which I calculate for each demographic group using the measured job flows in the monthly basic Currently Population Survey (CPS).⁵² Panel A of Table B1 reports the estimates. Consistent with prior literature, I find that, on average, job finding rates are higher for the more educated and for men. Table B2 demonstrates that the unemployment rate and unemployment durations implied by these labor market flows closely capture the empirical heterogeneity in these variables across demographic groups. Under the assumption that one earns nothing in

⁵² In calculating the labor market flow rates, I abstract for non-participation and define $f = \frac{UE}{UE+UU}$ and $s = \frac{EU}{EE+EU}$ where UE is the number of people transitioning from unemployment to employment across months, UU is the number from unemployment to unemployment, and so on. I restrict the sample to only include 1994-2018, those who report being either unemployed or employed, and those between the ages of 25 and 62. From that sample, I calculate the probability that across two consecutive months, someone transitions from unemployment to employment and from employment to unemployment. I then average these monthly rates across the entire sample period.

the periods of unemployment, the differences in the job-finding and job-separation rates fully characterize the heterogeneity in the persistence of unemployment across demographic groups.

I assume that the earnings of the employed (w_{it}) follow an AR(1) process with persistent and transitory shocks, the parameters of which are the persistence parameter, the variance of the persistent shock, the variance of the transitory shock, an initial variance in earnings, and an average earnings level. I estimate these parameters following the methodology in Heathcote et al. (2010) and using the earnings of those who report being employed in the Panel Study of Income Dynamics.⁵³ Panel B of Table B1 reports the level and variance of earnings produced by these estimates for each demographic group. I capture the well-documented differences in the average earnings by gender, race, and education.⁵⁴

For each group, I set the intertemporal elasticity of substitution to be 1.5, assume that agents cannot borrow (i.e., $b = 0$), and set the interest rate to be 1.02. Taking the earnings process for each demographic group as given, I choose the discount rate such that I match the average “empirical MPC” for each demographic group. In other words, I choose the discount factor such that the earnings-weighted average consumption drop per dollar lost upon transitioning to the unemployment state equals the earnings-weighted average MPC that I estimate in the PSID using the unemployment shock. The magnitude of this “empirical MPC” will in part reflect the persistence of the unemployment shock – all else equal, when the unemployment shock is more persistent, the average MPC is higher. Column 6 of Table B1 shows the resulting annual discount rates, which vary from 0.85 to 0.98. Empirically, groups with higher MPCs (Column 6) and more volatile incomes (Column 5) have lower discount rates (Column 7). Since agents are risk-averse, agents facing more income volatility want to accumulate assets, which pushes them away from their borrowing constraint and brings down the average MPC of that group. However, since it is precisely those groups that have higher measured MPCs, a higher degree of impatience is needed to match the MPC of the group in the model. The discount rates in the model need not reflect pure differences in time preferences across individuals and may also capture unmodeled differences in access to the banking sector or costs of borrowing. Even so, the estimates of the discount factor are actually in line with empirical estimates, both on average and in their patterns across demographics. For example, using experimental evidence in Denmark, Harrison et al. (2002) estimate higher discount factors for the rich, skilled, and educated. The last column of Table B1 shows the MPC in the model out of an unanticipated and transitory income shock. Unsurprisingly, for all demographic groups, this MPC is lower than the empirical MPC, but the two are highly correlated ($\rho = 0.97$), demonstrating that heterogeneity in the persistence of the unemployment shock across groups is not a key driver of MPC heterogeneity in the model.

Lastly, having matched the average MPC with the discount factor, I match the covariance between MPCs and earnings elasticities with the parameter χ in the rationing function. I estimate this parameter such that the $Cov(\widehat{MPC}_i, \gamma_i)$ is equal to 0.09, which is the empirical estimate from Section 5. I get that $\chi = 1.95$, which is similar to the empirical estimate in Table 2 as well.

B.4 Model estimates

I begin exploring the model-based estimates of the matching multiplier mechanism by looking at the importance of the incidence of the shock on the response of aggregate consumption. The left panel of Figure B1 shows the aggregate MPC in two scenarios. In the “actual scenario,” I consider an aggregate shock

⁵³ I restrict attention to those individuals who report being employed at the time of the PSID survey, are between the ages of 30 and 40, and for whom I can impute an MPC (i.e., observations with at least two lags of earnings).

⁵⁴ These parameter estimates are similar to, although slightly higher than, comparable estimates in the literature. For example, Heathcote et al. (2010) estimate an average permanent component of around 0.015 and a transitory component of around 0.1. Carroll et al. (2015) review the literature and show that estimates in the literature for the variance of the permanent component range from 0.01 to 0.054 and that estimates of the transitory variance range from 0.01 to 0.2.

to incomes that is distributed to workers such that the earnings elasticity of each worker is exactly equal to the γ_i , calibrated as described above. In the “benchmark scenario,” I consider an aggregate shock to income of the same size that is distributed such that the earnings elasticity of all workers is equal to 1. I derive the aggregate MPC by aggregating the response of consumption and dividing by the size of the aggregate change in incomes in period 1.

Before reporting those estimates, the first column in the left panel of Figure B1 reports, for comparison, the increase in the aggregate MPC coming from the covariance that I estimated in Section 4. The benchmark scenario in this case is simply the earnings-weighted average MPC in the data. As in Section 2, the aggregate MPC in the actual scenario is the benchmark MPC plus the estimated covariance between the empirical MPC and the earnings elasticity, which is 0.09.⁵⁵ With this specification, the estimated covariance increases the aggregate MPC by 24 percent.

The last two columns in the left panel of Figure B1 show how this empirical estimate compares to those from the model and demonstrate the key role played by the persistence of the unemployment shock in determining this mechanism. The second column shows first the increase in the aggregate MPC that results from a 1-period aggregate shock. Since this is a 1-time purely transitory cash drop, the aggregate MPC is determined not by the *empirical* MPC but rather by the *transitory* MPC, which Table B1 showed was far below the empirical estimates. However, Table B1 also shows that these transitory MPCs are highly correlated with the empirical MPCs, and thus, the aggregate MPC in the actual case is larger than the aggregate MPC in the empirical case. Indeed, while the levels are lower than the empirical estimates, the percent increase in the aggregate MPC is even larger at 30 percent. In this model, while the persistence of the shock matters for the level of the aggregate MPC, it is not critical for the amount of amplification implied by the heterogeneous incidence of the shock.

Lastly, the final columns in the left panel of Figure B1 shows the aggregate MPC in response to a persistent aggregate shock. In period 1, agents receive their period 1 shock, but also anticipate that the shock will decay slowly and therefore may respond today to expected future changes in income. I choose the persistence of the aggregate shock to match the persistence of the average unemployment spell, which is given by $\rho = 1 - 2\bar{f} = 0.34$. With this persistent shock, the aggregate benchmark MPC is much larger than with the transitory shock and very close to the empirical estimate in column 1. Importantly, the increase in the aggregate MPC across scenarios is large, similar in magnitude to the empirical estimate, and larger than in the transitory case in Column 2. When the shock is persistent and unevenly distributed in multiple periods, the same workers who face a large shock today also anticipate facing a larger shock in the next period, and thus the drop in their consumption today is even larger. The heterogeneous exposure of the future shock amplifies the heterogeneity today, and thus, the overall strength of this mechanism increases.⁵⁶

The right panel of Figure B1 uses the simplified structure of the model and the derived expressions for the multiplier in Proposition 1 to explore how well the empirical matching multiplier from Section 2 captures the general-equilibrium effects on *output* in this dynamic setting. First, column 1 in the right panel reproduces the empirical estimate, which uses the empirical MPCs and the simplified 2-period formula for

⁵⁵ These numbers differ from those in Table 3 because they do not account for the existence of non-labor income. Therefore, I do not scale the benchmark MPC by the empirical average elasticity of income to GDP and I do not rescale the covariance by a measure of the labor share. This results in a larger baseline MPC and a larger level increase between the actual and benchmark scenarios. However, the percentage increase in the MPC in Figure B1 is similar to that in Table 3.

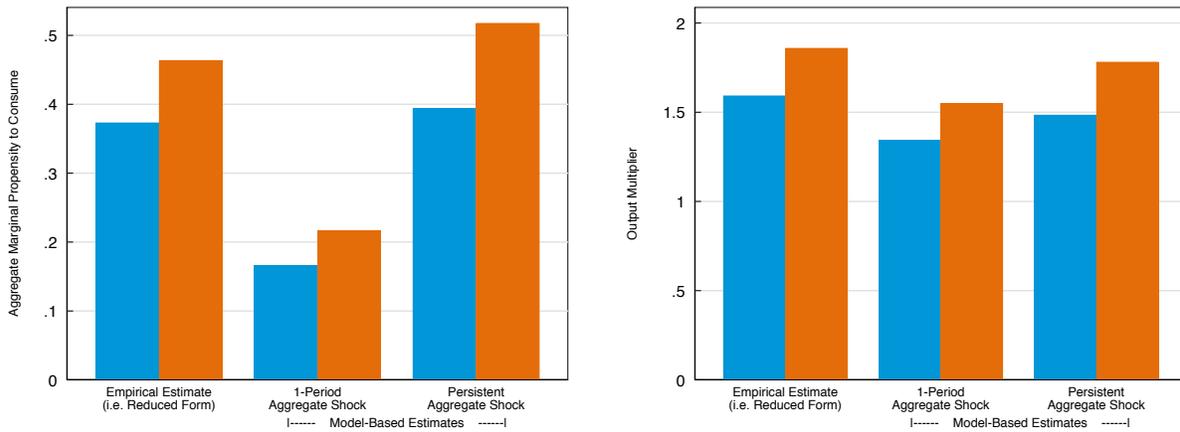
⁵⁶ It is important to note that this experiment captures the average persistence of the unemployment shock as in the data, but does not capture the heterogeneous persistence of the unemployment shock across individuals. If a large part of the heterogeneity in empirical MPCs across groups were driven by heterogeneity in the persistence of the unemployment shock, the amplification of the consumption response in the actual scenario in the model would be smaller than in the data. However, the amplification from the heterogeneous incidence is slightly larger in the model than in the data, suggesting, if anything, that the heterogeneous persistence of unemployment across individuals dampens the matching multiplier mechanism.

the multiplier. The following columns compare this to the model-based multipliers that take the dynamics of the model seriously, and are therefore based on the model-implied estimates of the intertemporal MPC matrix as in Auclert et al. (2018). Column 2 in the right panel corresponds to the scenario in column 2 in the left panel and shows the multiplier in response to a 1-period partial-equilibrium shock.⁵⁷ Since the level of the MPC is much lower, the level of the dynamic multiplier is also much lower in both the baseline and actual scenarios. However, still, the amplification implied by the heterogeneous incidence is meaningful, increasing the baseline multiplier by 21 percentage points, or 15 percent. Lastly, Column 3 in the right panel of Figure B1 shows the multiplier in response to a shock that has the persistence of the average unemployment spell. This column looks very similar to column 1, with a baseline multiplier of about 1.5 that is amplified by 20 percent with the heterogeneous incidence of the aggregate shock.

The analysis in both panels of Figure B1 demonstrate that the empirical matching multiplier used throughout the analysis, which was derived using MPCs out of persistent unemployment and a simplified 2-period multiplier closely approximates a dynamic multiplier of a persistent shock. Additionally, while the persistence of the shock is important for determining the level of the amplification, it is not the driver of the matching multiplier mechanism.

⁵⁷For the estimates in both columns 2 and 3 of the right panel of Figure B1, I assume that wages are rationed (and thus prices are fixed) for 10 periods.

Figure B1: The Matching Multiplier in the Model



Notes: The blue bars refer to the "benchmark scenario" in which the shock is distributed such that all workers have an elasticity of 1. The red bars refer to the "actual scenario", where the shock is distributed such that the each worker has an elasticity determined by the empirical covariance between elasticities and empirical MPCs. The left panel shows the aggregate MPC across the two scenarios and the right panel shows the output multiplier across the two scenarios.

Table B1: Model Calibration by Demographic Group

	A: Unemployment		B: Earnings		C: Model Calibration			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	f	s	$\overline{\log(w)}$	Std. $\log(w)$	Std. $\log(y)$	Empirical MPC	β	Transitory MPC
<i>High School or Less</i>								
Non-Black Men	0.35	0.02	10.76	0.75	0.75	0.46	0.96	0.23
Black Men	0.26	0.03	10.44	0.79	0.80	0.79	0.85	0.57
Non-Black Women	0.31	0.01	10.12	0.87	0.87	0.55	0.94	0.29
Black Women	0.23	0.02	10.07	0.89	0.90	0.59	0.93	0.29
<i>Some College or More</i>								
Non-Black Men	0.32	0.01	11.17	0.84	0.85	0.31	0.98	0.11
Black Men	0.29	0.01	10.77	0.73	0.73	0.73	0.89	0.45
Non-Black Women	0.34	0.01	10.58	0.98	0.98	0.36	0.96	0.17
Black Women	0.27	0.01	10.44	0.84	0.85	0.47	0.96	0.20

Notes: Panel A reports the job-finding (f) and job-separation (s) rates calculated from the monthly basic Current Population Survey from 1994-2018 and defined as $f = \frac{UE}{UU+UE}$ and $s = \frac{EU}{EE+EU}$, where UE is the number of people transitioning across months from unemployment to employment, UU is the number of people from unemployment to unemployment, etc. Panel B reports the average log of earnings and the standard deviation of that log of earnings from the model's income process for the employed. This AR(1) income process is estimated to match the PSID data following the methodology in Heathcote et al. (2010). I include both the nationally representative sample and SEO subsample of the PSID and use individual labor rather than household earnings. Panel C shows parameters for the steady-state of the model. Column 5 reports the standard deviation of the overall earnings process for each group in the model. The empirical MPC in Column 6 is the earnings-weighted average MPC estimated in the PSID as described in Section 4.1. Column 7 reports the discount factor that is needed to match the earnings-weighted average MPC given the other parameters. Column 8 reports the earnings-weighted average transitory MPC in the model, which is defined as the increase in consumption per dollar increase in income.

Table B2: Additional Statistics for Unemployment Process

	Unemployment Rate		Unemployment Duration			Earnings Recovery (X)			
	f	s	Model	CPS	Model	CPS	PSID	Model	PSID
<i>High School or Less</i>									
Non-Black Men	0.35	0.35	0.05	0.05	2.85	5.16	3.32	0.68	0.66
Black Men	0.26	0.26	0.09	0.09	3.78	6.66	3.84	0.60	0.80
Non-Black Women	0.31	0.31	0.04	0.04	3.19	5.19	2.44	0.65	0.57
Black Women	0.23	0.23	0.08	0.08	4.42	6.51	3.37	0.55	0.55
<i>Some College or More</i>									
Non-Black Men	0.32	0.32	0.02	0.03	3.16	5.55	2.36	0.65	0.63
Black Men	0.29	0.29	0.05	0.05	3.44	6.30	2.61	0.62	0.64
Non-Black Women	0.34	0.34	0.02	0.03	2.92	4.93	2.15	0.67	0.49
Black Women	0.27	0.27	0.04	0.05	3.72	6.19	2.80	0.60	0.63

Notes: Columns 1 and 2 report the average monthly job finding (f) and job separation (s) rates, respectively, calculated from the matched monthly basic CPS. Column 3 reports the unemployment rate implied by the job-finding and job separation rate ($u = \frac{s}{s+f}$) while Column 4 reports the average monthly unemployment rate in the CPS. Column 5 reports the average unemployment duration implied by the job-finding rate ($\frac{1}{f}$) while column 6 reports the average unemployment duration (in months) in the CPS. Column 8 reports the average number of months in a year that someone who is unemployed at the time of the survey will be unemployed for in the PSID. Unlike the estimate in Column 7, this measure is not restricted to a continuous unemployment spell and is truncated at 52 weeks. Lastly, Column 8 reports the model-based fraction of earnings captured by those in the unemployment state, which is the result of a simulated earnings process with 10,000 individuals and 30 years of earnings. Column 9 reports the percentage difference in annual earnings for the unemployed and employed in the PSID. This is estimated on the sample used for the MPC estimation and described in detail in Section 4.1.