NBER WORKING PAPER SERIES

STOCK MARKET WEALTH AND THE REAL ECONOMY: A LOCAL LABOR MARKET APPROACH

Gabriel Chodorow-Reich
Plamen T. Nenov
Alp Simsek

Working Paper 25959
http://www.nber.org/papers/w25959

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2019, Revised February 2020

We would like to thank George-Marios Angeletos, Ricardo Caballero, Anthony DeFusco (discussant), Paul Goldsmith-Pinkham, Annette Vissing-Jorgensen, Kairong Xiao, and numerous seminar participants for helpful comments. Joel Flynn and Katherine Silva provided excellent research assistance. Chodorow-Reich acknowledges support from the Molly and Dominic Ferrante Economics Research Fund. Nenov would like to thank Harvard University and the NBER for their hospitality during the initial stages of the project. Simsek acknowledges support from the National Science Foundation (NSF) under Grant Number SES-1455319. Any opinions, findings, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the NSF or the National Bureau of Economic Research.

Stock Market Wealth and the Real Economy: A Local Labor Market Approach
Gabriel Chodorow-Reich, Plamen T. Nenov, and Alp Simsek
NBER Working Paper No. 25959
June 2019, Revised February 2020
JEL No. E21,E32,E44

ABSTRACT

We provide evidence of the stock market wealth effect on consumption by using a local labor market analysis and regional heterogeneity in stock market wealth. An increase in local stock wealth driven by aggregate stock prices increases local employment and payroll in nontradable industries and in total, while having no effect on employment in tradable industries. In a model with consumption wealth effects and geographic heterogeneity, these responses imply a marginal propensity to consume out of a dollar of stock wealth of 3.2 cents per year. We also use the model to quantify the aggregate effects of a stock market wealth shock when monetary policy is passive. A 20% increase in stock valuations, unless countered by monetary policy, increases the aggregate labor bill by at least 1.7% and aggregate hours by at least 0.75% two years after the shock.

Gabriel Chodorow-Reich
Department of Economics
Harvard University
1805 Littauer Center
Cambridge, MA 02138
and NBER
chodorowreich@fas.harvard.edu

Plamen T. Nenov
Norwegian Business School (BI)
Nydalsveien 37 0484 Oslo, Norway
plamen.nenov@bi.no

Alp Simsek
Department of Economics, E52-552
MIT
50 Memorial Drive
Cambridge, MA 02142
and NBER
asimsek@mit.edu

A online appendix is available at http://www.nber.org/data-appendix/w25959
A Dynamic link to the most recent draft: is available at https://www.dropbox.com/s/7q1zzxkkr184rjy/crns_stock

1 Introduction

According to a recent textual analysis of FOMC transcripts by Cieslak and Vissing-Jorgensen (2017), many U.S. policymakers believe that stock market fluctuations affect the labor market through a consumption wealth effect. In this view, a decline in stock prices reduces the wealth of stock-owning households, causing a reduction in spending and hence in employment. While apparently an important driver of U.S. monetary policy, this channel has proved difficult to establish empirically. The main challenge arises because stock prices are forward-looking. Therefore, an anticipated decline in future economic fundamentals could also lead to both a negative stock return and a subsequent decline in household spending and employment.

We use a local labor market analysis to address this empirical challenge and provide quantitative evidence on the stock market consumption wealth effect. Our empirical strategy combines regional heterogeneity in stock market wealth with aggregate movements in stock prices to identify the causal effect of stock wealth changes on regional labor market outcomes. We then present a model that relates the regional outcomes to the household-level propensity to consume out of stock wealth as well as to the aggregate labor market effects of stock wealth changes. Our empirical estimates map into a household-level annual marginal consumption propensity of 3.2 cents per dollar of stock wealth and imply that annual aggregate payroll increases by 1.7% following a yearly standard deviation increase in the stock market, unless countered by monetary policy.

To frame the regional analysis, it helps to begin by describing the consumption wealth effect in our model setting. The environment features a continuum of areas, a tradable good and a nontradable good, and two factors of production, capital and labor. The only heterogeneity across regions is in their ownership of capital, which also equates to stock wealth. The aggregate price of capital is endogenous and fluctuates due to changes in households' beliefs about the expected *future* productivity of capital. An increase in stock wealth increases local spending on nontradable goods, and more so in areas with greater capital ownership. Higher spending drives up the labor bill and increases employment in the nontradable sector and in total. Local wages increase (weakly) more in high wealth areas, which induces a (weak) fall in tradable employment.

In the data, we measure changes in county-level stock market wealth in three steps. In the first step, we capitalize dividend income reported on tax returns aggregated to the county level to arrive at a county-level measure of taxable stock wealth. Our capitalization method improves on existing work such as in Saez and Zucman (2016) by allowing for heterogeneity in dividend yields by wealth, which we obtain using a sample of account-level portfolio holdings from a large discount broker. In the second step, we adjust this measure of taxable stock wealth to account for non-taxable (e.g., retirement) stock wealth, using information on the relationship between taxable and total stock wealth and demographics in the Survey of Consumer Finances. In the final step, we multiply the total county stock wealth with the return on the market (CRSP value-weighted) portfolio and a county-specific portfolio beta constructed from county demographic information and variation in betas across the age distribution in the data from the discount broker. This provides a measure of the change in county stock wealth driven by the aggregate stock return. Motivated by our theoretical analysis, we then divide this change by the county labor bill to arrive at our main regressor.

Our empirical specification identifies the effect of changes in stock wealth on local labor market outcomes by exploiting the substantial variation in the aggregate stock return that occurs independent of other macroeconomic variables. In particular, we allow high wealth areas to exhibit greater sensitivity to changes in aggregate bond wealth, aggregate housing wealth, and aggregate labor income and non-corporate business income, and also control for county fixed effects, state-by-quarter fixed effects, and a Bartik-type industry employment shift-share. Our identifying assumption is that, conditional on these controls, areas with high stock market wealth do not experience unusually rapid employment or payroll growth following a positive aggregate stock return for reasons other than the stock market wealth effect on local spending.

An increase in local stock wealth induced by a positive stock return increases total local employment and payroll. Seven quarters after an increase in stock market wealth equivalent to 1% of local labor market income, local employment is 0.77 basis points higher and local payroll is 2.18 basis points higher. Because stock returns are nearly i.i.d., these responses reflect the short-run effect of a permanent change in stock market wealth. Motivated by the theory, we also investigate the effect on employment and the labor bill in the nontradable and tradable industries, following the sectoral classifications in Mian and Sufi (2014). Consistent with the theory, the employment response in nontradable industries exceeds the overall response, while employment in tradable industries does not increase. We also report a large response in the residential construction sector, consistent with a household demand channel.

The main threat to a causal interpretation of these findings is that high wealth areas respond differently to other aggregate variables that co-move with the stock market. This concern motivates the variables included in our baseline specification. The absence of "pre-trend" differences in outcomes in the quarters before a positive stock return and the non-response of employment in the tradable sector support a causal interpretation of our findings. We report additional robustness along a number of dimensions, including: using a more parsimonious specification that excludes the parametric controls; including interactions of stock market wealth with TFP growth to allow wealthier counties to have different loadings on this variable; controlling for local house prices; using only within commuting zone variation in stock market wealth; subsample analysis including dropping the wealthiest counties and the quarters with the most volatile stock returns; and not weighting the regression. A decomposition along the lines of Andrews et al. (2017) shows that no single state drives the results. We also report a quantitatively similar response using cross-state variation and state-level consumption expenditure from the Bureau of Economic Analysis.

We combine our empirical results with the theoretical model to calibrate two key parameters: the strength of the household-level stock wealth effect and the degree of local wage adjustment. To calibrate the stock wealth effect, we provide a separation result from our model that decomposes the empirical coefficient on the *nontradable* labor bill into the product of three terms: the partial equilibrium marginal propensity to consume out of stock market wealth, the local Keynesian multiplier (equivalent to the multiplier on local government spending), and the labor share of income.[1] This decomposition applies to more general changes in local consumption demand and therefore may be of use outside our particular setting. We use standard values from previous literature to calibrate the labor share of income and the local Keynesian multiplier. Given these values, the empirical response of the nontradable labor bill implies that in partial equilibrium a one dollar increase in stock-market wealth increases annual consumption expenditure by about 3.2 cents two years after the shock. For the degree of wage adjustment, comparing the response of total employment with the response of the total labor bill suggests that a 1 percent increase in labor (total hours worked) is associated with a 0.9 percent increase in wages at a two year horizon.

Finally, we use the model to quantify the *aggregate* effects that stock price shocks would generate if monetary policy (or other demand-stabilization policies) did not respond to the shock. We first show that a one dollar increase in stock market wealth has the same *proportional* effect on the *local nontradable* and *aggregate total* labor bills, up to an adjustment for the difference in the local and aggregate spending multipliers. Homothetic preferences and production across sectors underlies this theoretical result, and we provide evidence of such homotheticity in the data. Next, we show how the local response of wages informs about the aggregate wage Phillips curve in our model. Since labor markets are local, the aggregate wage response is similar to the local wage response, with an adjustment due to the fact that demand shocks impact aggregate inflation and local inflation differently. We then consider a 20% positive shock to stock valuations—approximately the yearly standard deviation of stock returns. Using our empirical estimate for the nontradable labor bill, and applying a

[1] In general, there may be an additional term reflecting the response of output in the tradable sector when relative prices change across areas. This term disappears in our benchmark calibration, which features Cobb-Douglas preferences across tradable goods produced in different regions. Allowing for a non-unitary elasticity of substitution across regions does not meaningfully change our conclusions.

bounding argument for moving from local to aggregate effects similar to that in Chodorow-Reich (2019), this shock would increase the aggregate labor bill by at least 1.7% two years after the shock. Combining this effect with the degree of aggregate wage adjustment implied by our local estimates, the shock would also increase aggregate hours by at least 0.75%.

The rest of the paper is organized as follows. After discussing the related literature, we start by presenting the empirical analysis. Section 2 describes the data sets and the construction of our main variables. Section 3 presents the baseline empirical specification and discusses conditions for causal inference. Section 4 contains the empirical results. We then turn to the theoretical analysis and the structural interpretation. Section 5 describes our model. Section 6 uses the empirical results to calibrate the model and derive the partial equilibrium wealth effect. Section 7 calculates the implied aggregate wealth effects, and Section 8 concludes.

Related literature. Our paper contributes to a large literature that investigates the relationship between stock market wealth, consumption, and the real economy. A major challenge is to disentangle whether the stock market has an effect on consumption over a relatively short horizon (the direct wealth effect), or whether it simply predicts future changes in productivity, income, and consumption (the leading indicator effect). The challenge is compounded by the scarcity of data sets that contain information on household consumption and financial wealth. The recent literature has tried to address these challenges in various ways (see Poterba (2000) for a survey of the earlier literature).

The literature using aggregate time series data finds mixed evidence (see e.g. Poterba and Samwick, 1995; Davis and Palumbo, 2001; Lettau et al., 2002; Lettau and Ludvigson, 2004; Carroll et al., 2011). Davis and Palumbo (2001) and Carroll et al. (2011) estimate a wealth effect of up to around 6 cents. On the other hand, Lettau and Ludvigson (2004) argue for more limited wealth effects. However, an aggregate time series approach introduces two complications: First, in an environment in which monetary policy effectively stabilizes aggregate demand fluctuations, as in our model, there can be strong wealth effects and yet no relationship between asset price shocks and aggregate consumption. Second, stock market fluctuations may affect aggregate demand via an investment channel (see Cooper and Dynan (2016) for other issues with using aggregate time series in this context).

Another strand of the literature uses household level data and exploits the heterogeneity in household wealth to isolate the stock wealth effect. Dynan and Maki (2001) use Consumer Expenditure Survey (CE) data to compare the consumption response of stockholders with non-stockholders. They find a relatively large marginal propensity to consume (MPC) out of stock wealth—around 5 to 15 cents per dollar per year. However, Dynan (2010) re-examines the evidence by extending the CE sample to 2008 and finds weaker effects. More recently, Di Maggio et al. (forthcoming) use detailed individual-level administrative wealth data for Sweden to identify the stock wealth effect from variation in individual-level portfolio returns. They find substantial effects: the top 50% of the income distribution, who own most of the stocks, have an estimated MPC of around 5 cents per dollar per year.[2]

We complement these studies by focusing on *regional* heterogeneity in stock wealth. We show how the regional empirical analysis can be combined with a model to estimate the household-level stock wealth effect. The MPC implied by our analysis (3.2 cents per dollar per year) is close to estimates from the recent literature. An important advantage of our approach is that it directly estimates the local *general equilibrium* effect. In particular, by examining the labor market response, we provide direct evidence on the margin most important to monetary policymakers.

Case et al. (2005) and Zhou and Carroll (2012) also use regional variation to estimate financial wealth effects. Case et al. (2005) overcome the absence of geographic data on financial wealth by using state-level mutual fund holdings data from the Investment Company Institute (ICI) and measure state consumption using retail sales data from the Regional Financial Associates. Zhou and Carroll (2012) criticize the data construction and empirical specification in Case et al. (2005) and construct their own data set using proprietary data on state-level financial wealth and retail sales taxes as a proxy for consumption. Both papers find negligible stock wealth effects and a sizable housing wealth effect. Relative to these papers, we exploit the much greater variation in financial wealth across counties than across states and provide evidence on the labor market margin directly. Other recent papers use regional variation but focus only on estimating housing wealth effects (Mian et al., 2013; Mian and Sufi, 2014; Guren et al., 2018).[3]

Our focus on the consumption wealth channel complements research on the investment channel of the stock market that dates to Tobin (1969) and Hayashi (1982). Under the identifying assumptions we articulate below, our local labor market analysis absorbs the effects of changes in Tobin's Q or the cost of equity financing on investment into a time fixed effect, allowing us to isolate the consumption wealth channel.

Our theoretical framework builds upon the model in Mian and Sufi (2014) by incor-

[2]See also Bostic et al. (2009) and Paiella and Pistaferri (2017) for similar analyses of stock wealth effects in different contexts.

[3]See also Case et al.(2005; 2011), Campbell and Cocco (2007), Mian and Sufi (2011), Carroll et al. (2011), and Browning et al. (2013), among others. In terms of comparison of wealth effects from stock wealth versus housing wealth, Guren et al. (2018) estimate an MPC out of housing wealth of around 2.7 cents during 1978-2017, which is comparable in magnitude to our estimate of the stock wealth effect. This is substantially lower than the estimates in Mian et al. (2013) and Mian and Sufi (2014), which are in the range of 7 cents. See Guren et al. (2018) for a discussion of the possible drivers of these differences.

porating several features important for a structural interpretation of the results, including endogenous changes in wealth, monetary policy, partial wage adjustment, and imperfectly substitutable tradable goods. Our framework also shares features with models of small open economies with nominal rigidities (e.g. Gali and Monacelli, 2005) adapted to the analysis of monetary unions by Nakamura and Steinsson (2014) and Farhi and Werning (2016), but differs from these papers by including a fully nontradable sector. This feature facilitates the structural interpretation and aggregation of the estimated local general equilibrium effects.

Our structural interpretation and aggregation results represent methodological contributions that apply beyond our particular model. First, and similar to the approach in Guren et al. (2018) and formalized in Guren et al. (2020), we illustrate how the estimated local general equilibrium effects can be combined with external estimates of the local income multiplier (e.g., estimates from local government spending shocks) to obtain the partial equilibrium spending effect.[4] Our decomposition differs from theirs in that it applies to the coefficient for the nontradable labor bill—a variable that is easily observable at the regional level—and therefore includes an adjustment for the labor share of income. Second, we show how, under standard assumptions, the response of the local labor bill in the *nontradable* sector provides a direct and transparent bound for the response of the aggregate effect across *all* sectors when monetary policy does not react.

2 Data

In this section we explain how we measure the key objects in our empirical analysis: the ratio of geographic stock market wealth to labor income, the stock market return, employment, and payroll. Our geographical unit is a U.S. county. This level of aggregation leaves ample variation in stock market wealth while being large enough to encompass a substantial share of spending by local residents. The U.S. contains 3,142 counties using current delineations.

2.1 Stock Market Wealth

We denote our main regressor as $S_{a,t-1}R_{a,t-1,t}$, where $S_{a,t-1}$ is stock market wealth in county a in period $t-1$ normalized by the period $t-1$ labor bill and $R_{a,t-1,t}$ is the portfolio return between $t-1$ and t. In Section 5, we show that regressions of log changes in local labor market outcomes on this variable yield coefficients tightly related to the key parameters of our model.

[4]In contemporaneous work, Wolf (2019) formally establishes (in a closed economy setting) conditions under which the multiplier effects from private spending are exactly the same as the multiplier effects from public spending.

We construct local stock market wealth by capitalizing taxable dividend income and then adjusting for stock wealth held in non-taxable accounts. We summarize our methodology here and provide additional detail of the data, sample construction, and adjustments in Appendix A.1. Our capitalization method involves multiplying observed taxable dividend income by a price-dividend ratio to arrive at stock wealth held in taxable accounts.[5] We start with IRS Statistics of Income (SOI) data containing county aggregates of annual dividend income reported on individual tax returns, over the period 1989-2015. Dividend income as reported on form 1040 includes any distribution from a C-corporation. It excludes distributions from partnerships, S-corporations, or trusts, except in rare circumstances where S-corporations that converted from C-corporations distribute earnings from before their conversion. While we cannot separate distributions from publicly-traded and privately-held C-corporations, we show in Appendix A.1.4 that equity in privately-held C-corporations is too small (less than 7% of total equity of C-corporations) to meaningfully affect our results.

We construct a county-specific capitalization factor as the product of the price-dividend ratio on the value-weighted CRSP portfolio and a time-varying county-specific adjustment. The CRSP portfolio contains all primary listings on the NYSE, NYSE MKT, NASDAQ, and Arca exchanges and, therefore, covers essentially the entire U.S. equity market. The county-specific adjustment recognizes that older individuals both have higher average wealth and hold higher dividend-yield stocks, as first conjectured in Miller and Modigliani (1961) and documented in Graham and Kumar (2006). We believe we are the first to apply such an adjustment in capitalizing equity wealth. To do so, we follow Graham and Kumar (2006) and use the Barber and Odean (2000) data set of individual account-level stock holdings from a large discount broker over the period 1991-1996.[6] Specifically, as we describe in more detail in Appendix A.1.2, we merge the Barber and Odean (2000) data set with CRSP stock

[5]The literature has proposed other income measures and capitalization factors. Mian et al. (2013) and Mian and Sufi (2014) group dividends, interest, and other non-wage income together and use the ratio of total household financial wealth in the Financial Accounts of the United States (FAUS) to the national aggregate of this combined income measure as a single capitalization factor for all financial wealth. Saez and Zucman (2016) and Smith et al. (In progress) use both dividends and capital gains to allocate directly held corporate equities in the FAUS, with Smith et al. arguing forcefully for a low weight on the capital gains component because realized capital gains include many transactions other than sales of corporate equity. Relative to these alternatives, capitalizing dividends using a price-dividend ratio isolates the income stream most closely related to corporate equity wealth and facilitates the adjustment for heterogeneous dividend yields described below.

[6]The data are a random sample of accounts at the brokerage and have been used extensively to study individual trading behavior (Barber and Odean, 2000, 2001; Graham and Kumar, 2006; Barber and Odean, 2007; Mitton and Vorkink, 2007; Kumar, 2009; Seasholes and Zhu, 2010; Kent et al., 2019). Graham and Kumar (2006) compare the data with the 1992 and 1995 waves of the SCF and show that the stock holdings of investors in the brokerage data are fairly representative of the overall population of retail investors. We consider taxable accounts with at least one dividend-paying stock to mimic the dividends observed in the IRS data.

and mutual fund data and compute average dividend yields for five age groups, separately for each Census Region. The dividend yield slopes upward with age, with individuals 65 and over holding stocks with a dividend yield about 10% (*not* p.p.) higher than the market average and individuals 35 and younger holding stocks with a dividend yield about 10% lower than the market average. Importantly, variation by age accounts for essentially all of the variation in dividend yields across the wealth distribution, as shown in Figure A.1 and Table A.1. We combine the age-specific dividend yields with county-level demographic information and wealth by age group from the Survey of Consumer Finances (SCF). We then adjust the CRSP dividend yield in each county-year by the age-wealth-weighted average of the age-specific dividend yields.

We next adjust county taxable stock market wealth to account for wealth held in non-taxable accounts, primarily in defined contribution pension plans.[7] We do not include wealth in defined benefit pension plans, since household claims on that wealth do not fluctuate directly with the value of the stock market. Roughly one-third of total household stock market wealth is held in non-taxable accounts (see Figure A.4). In Appendix A.1.3, we estimate the relationship at the household level between total stock market wealth, taxable stock market wealth, and household demographic characteristics, using the SCF. Total and taxable stock market wealth vary almost one-to-one, reflecting statutory limits on contributions to non-taxable accounts that make non-taxable wealth much more evenly distributed than taxable wealth. The variables also explain total wealth well, with an R^2 above 0.9. We combine the coefficients on taxable wealth and demographic characteristics from the SCF with our county-level measure of taxable stock wealth and county-level demographic characteristics to produce our final measure of total county stock market wealth. Finally, we divide this measure by SOI (annual) county labor income to arrive at our measure of local stock market wealth relative to labor income, $S_{a,t}$.

Figure 1a shows the variation in this measure across U.S. counties in 1990. Figure 1b and Figure 1c show the variation in 1990 and 2015, respectively, after removing state-specific means. The within-state differences are persistent over time, with a within-state correlation between $S_{a,1990}$ and $S_{a,2015}$ of 0.81. Table A.4 reports summary statistics for $S_{a,t}$ and other variables used in the analysis.

2.2 Stock Market Return

We write the stock market return in county a as $R^*_{a,t-1,t} = \alpha_a + R^f_{t-1,t} + b_{a,t} \times (R^m_{t-1,t} - R^f_{t-1,t}) + e_{a,t-1,t}$, where $R^f_{t-1,t}$ is the risk-free rate in period t, $R^m_{t-1,t}$ is the market return, $b_{a,t}$

[7]This adjustment is appropriate if the marginal propensities to consume out of taxable and non-taxable stock wealth are the same. We revisit this assumption at the end of our analysis (see Footnote 44.)

Figure 1: Stock Market Wealth Relative to Labor Income Across U.S. Counties.

(a) 1990

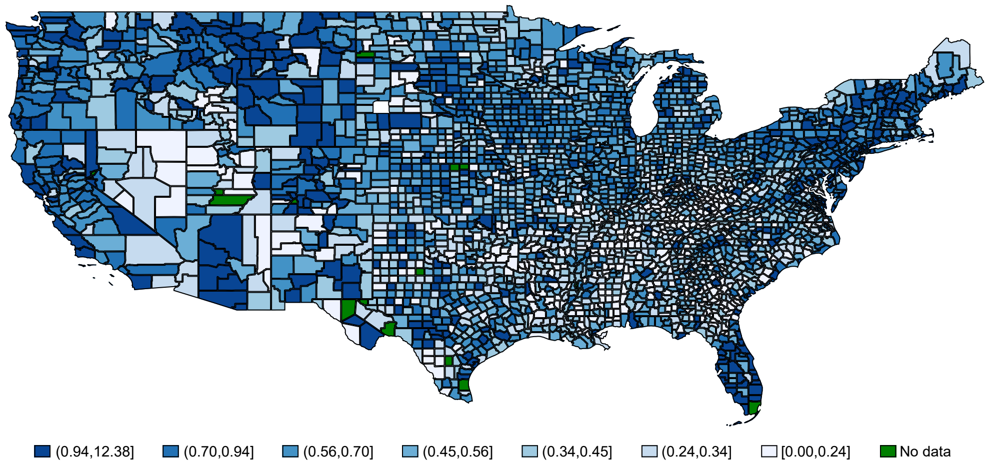

(b) 1990, within state

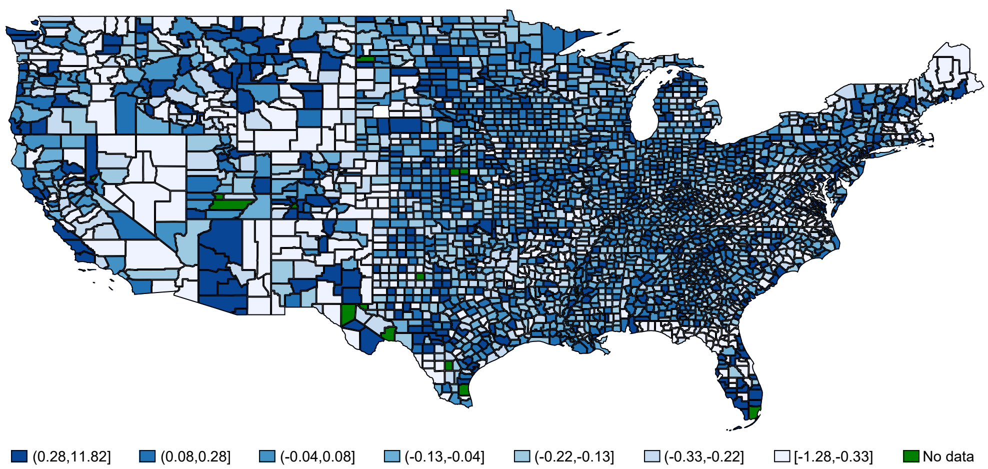

(c) 2015, within state

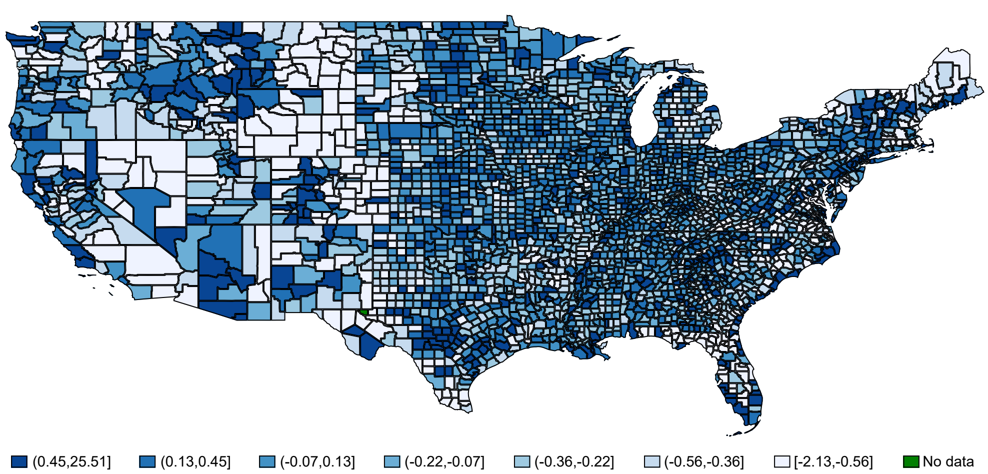

is a county-specific portfolio beta, and $e_{a,t-1,t}$ is an idiosyncratic component of the return. We do not observe $R^*_{a,t-1,t}$. Instead, we define the variable $R_{a,t-1,t}$ that enters into our main regressor as $R_{a,t-1,t} = R^f_{t-1,t} + b_{a,t} \times (R^m_{t-1,t} - R^f_{t-1,t})$. To operationalize $R_{a,t-1,t}$, we equate the risk-free rate $R^f_{t-1,t}$ with the interest rate on a 3-month Treasury bill, the market return $R^m_{t-1,t}$ with the total return on the value-weighted CRSP portfolio, and construct the county-specific portfolio beta $b_{a,t}$ using the relationship between market beta and age in the Barber and Odean (2000) data set and our measure of the county age-wealth distribution. This adjustment incorporates the tendency for older, wealthier households to hold stocks with lower betas, a pattern we document in Figure A.6 of the online appendix. Ignoring it would result in systematic over-counting of changes in wealth in high wealth areas when the stock market changes, leading to an under-estimate of the consumption wealth effect, although this effect turns out to be small in practice as the $b_{a,t}$ all lie between 0.97 and 1.03.

We now discuss the differences between the true county return $R^*_{a,t-1,t}$ and the measured return $R_{a,t-1,t}$ and why these differences do not affect the validity of our empirical analysis. Three possible differences exist. First, the true county return includes a county-specific α_a, reflecting differences in portfolio characteristics and the possibility that high wealth areas have systematically better portfolios, as suggested by Fagereng et al. (2016). Our empirical specification will include county fixed effects to absorb this type of heterogeneity. Second, high wealth areas could have systematically riskier or less risky stock portfolios beyond the correlation due to age, in which case we would systematically mis-measure $b_{a,t}$. While previous work has documented that wealthy households have portfolios tilted toward riskier asset classes than the general population (Carroll, 2000; Calvet and Sodini, 2014), here what matters is risk-taking *within* stock portfolios. Figure A.6 shows this correlation using the Barber and Odean (2000) data set. Except for the bottom wealth decile, who typically hold only one or two securities and have very low beta portfolios, there is a nearly flat relationship between beta and wealth decile within age bins. Therefore, this source of heterogeneity does not appear important in practice. Third, the true return $R^*_{a,t-1,t}$ contains an idiosyncratic component $e_{a,t-1,t}$, reflecting differences in portfolio allocation arising, for example, from home bias as documented in Coval and Moskowitz (1999) or from differences in market beta uncorrelated with wealth. This component has no impact on our empirical results because it gives rise to idiosyncratic changes in wealth that are uncorrelated with our main regressor. This statement remains true even if the idiosyncratic part of the return correlates with local economic activity, as might occur due to home bias in portfolio allocation.[8]

[8] Formally, assume the true structural model is $y_a = \beta\left(S_a R^*_a\right) + \epsilon_a$ and $R^*_a = R_a + e_a$, where y_a is an outcome, e_a is a mean-zero component of the return independent of wealth S_a or the measured part of the return R_a, and the structural residual ϵ_a is independent of the measured change in wealth $S_a R_a$. (We have dropped time subscripts and ignored the component α_a to simplify notation and without loss of generality).

Figure 2: Attributes of Quarterly Stock Returns

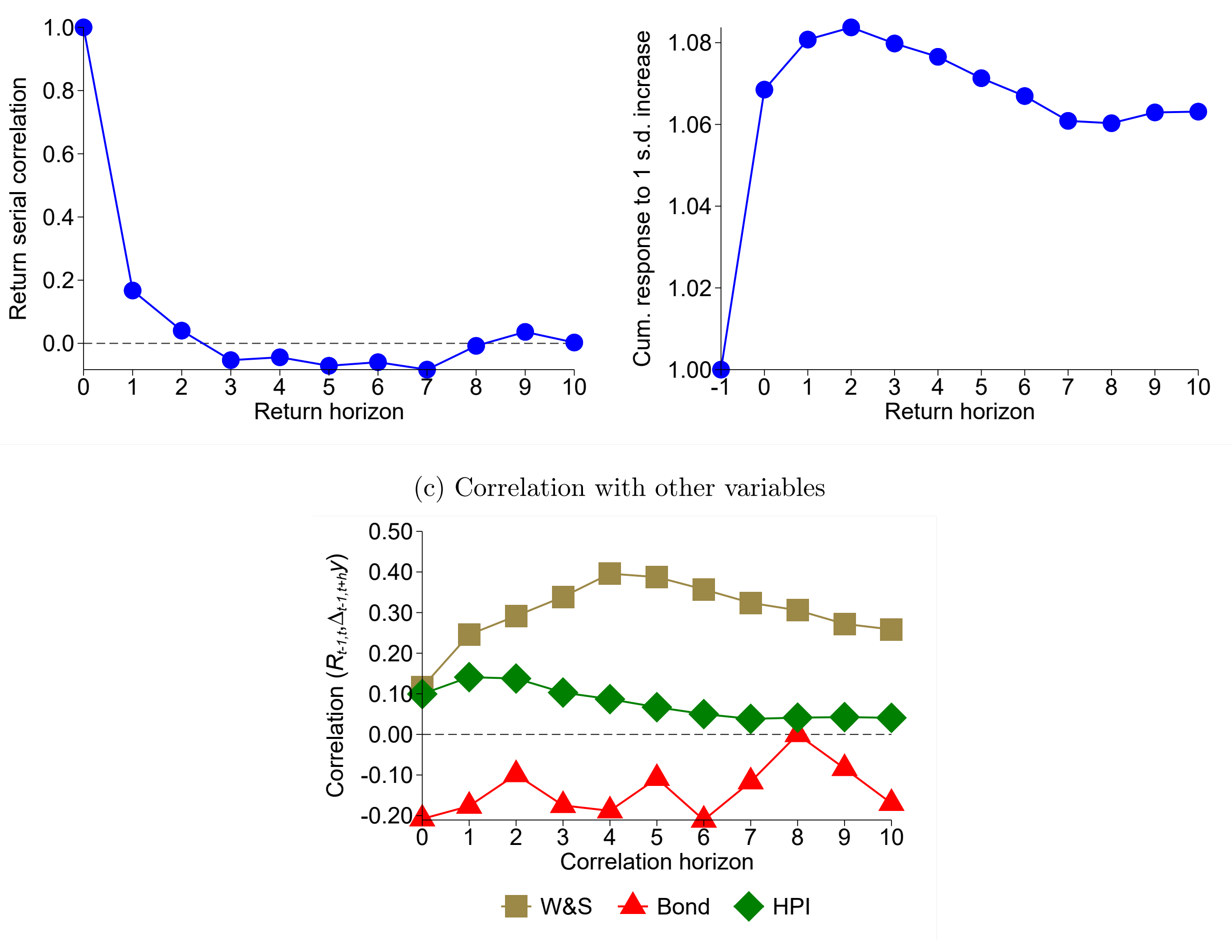

(a) Serial correlation of returns

(b) Cumulative return response

(c) Correlation with other variables

Notes: Panel (a) reports the coefficients β_h from estimating the regression $R_{t+h-1,t+h} = \alpha_h + \beta_h R_{t-1,t} + e_h$ at each quarterly horizon h shown on the lower axis, where $R_{t+h-1,t+h}$ is the total return on the value-weighted CRSP portfolio between quarters $t+h-1$ and $t+h$. Panel (b) reports the transformation $\Pi_{h=0}^{j}(1+\beta_h\sigma_R)$ at each quarterly horizon j shown on the lower axis, where σ_R is the standard deviation of the CRSP return. Panel (c) reports the correlation coefficients of $R_{t-1,t}$ and $y_{t-1,t+h}$ at each quarterly horizon h shown on the lower axis, where $y_{t-1,t+h}$ is the log change in aggregate labor compensation, the holding return on the 5 year Treasury, or the change in aggregate house prices between $t-1$ and $t+h$.

Substituting, we have $y_a = \beta(S_a R_a) + u_a$, where $u_a = \beta S_a e_a + \epsilon_a$ is a composite residual. Therefore, the coefficient $\hat{\beta}$ from regressing y_a on $S_a R_a$ asymptotes to β, since $Cov(S_a e_a, S_a R_a) = Cov(\epsilon_a, S_a R_a) = 0$ by the independence assumptions on e_a and ϵ_a. Alternatively, one can think of $S_{a,t-1}R_{t-1,t}$ as the excluded instrument and $S_{a,t-1}R^*_{a,t-1,t}$ as the endogenous variable in an instrumental variables design. Under the assumption of purely idiosyncratic heterogeneity, the first stage regression of $S_{a,t-1}R^*_{a,t-1,t}$ on $S_{a,t-1}R_{t-1,t}$ would yield a coefficient of 1, in which case the IV coefficient coincides with the reduced form coefficient that we estimate. Importantly, this argument extends straightforwardly to mis-measurement of $S_{a,t}$ due to heterogeneity in the price-dividend ratio uncorrelated with true wealth. Finally, the argument makes no assumption on the correlation between the idiosyncratic component of the return e_a and the structural residual ϵ_a, as might occur in the context of home bias in portfolio allocation. Hyslop and Imbens (2001) provide a more general discussion of measurement error that does not lead to biased estimation.

Figure 2a shows the serial correlation in the quarterly return on the CRSP portfolio and Figure 2b the cumulative return following a one standard deviation increase in the stock market during our sample period. As is well known, stock returns are nearly i.i.d., a result confirmed by the almost complete absence of serial correlation in Figure 2a. This pattern facilitates interpretation of our empirical results since it implies that a stock return in period t has a roughly permanent effect on wealth, and we mostly ignore the small momentum and subsequent reversal shown in Figure 2b in what follows. Figure 2c shows the correlation of the period t stock return with the changes in other macroeconomic aggregate variables over the horizon $t-1$ to $t+h$. In our sample, the stock market return is positively correlated with aggregate labor income and house prices, and negatively correlated with fixed income returns. However, the correlation coefficients are all well below one, reflecting the substantial movement in stock prices independent of these other factors (Shiller, 1981; Cochrane, 2011; Campbell, 2014).

2.3 Outcome Variables

Our main outcome variables are log employment and payroll from the Bureau of Labor Statistics Quarterly Census of Wages and Employment (QCEW). The source data for the QCEW are quarterly reports filed with state employment security agencies by all employers covered by unemployment insurance (UI) laws. The QCEW covers roughly 95% of total employment and payroll, making the data set a near universe of administrative employment records. We use the NAICS-based version of the data, which start in 1990, and seasonally adjust the published county-level data by sequentially applying Henderson filters using the algorithm contained in the Census Bureau's X-11 procedure.[9]

An important element of our analysis is to distinguish between responses in sectors affected by local demand shocks, which we refer to as "nontradable" sectors, and "tradable" sectors unlikely to be affected by local demand shocks. We follow Mian and Sufi (2014) and label NAICS codes 44-45 (retail trade) and 72 (accommodation and food services) as nontradable and NAICS codes 11 (agriculture, forestry, fishing and hunting), 21 (mining, quarrying, and oil and gas extraction), and 31-33 (manufacturing) as tradable.[10] The re-

[9]The NAICS version of the QCEW contains a number of transcription errors prior to 2001. We follow Chodorow-Reich and Wieland (Forthcoming, Appendix F) and hand-correct these errors before applying the seasonal adjustment procedure.

[10]Mian and Sufi (2014) exclude NAICS 721 (accommodation) from their definition of nontradable industries. We leave this industry in our measure to avoid complications arising from the much higher frequency of suppressed data in NAICS 3 than NAICS 2 digit industries in the QCEW data. The national share of nontradable employment and payroll in NAICS 721 are both less than 8% and we have verified using counties with non-suppressed data that including this sector does not affect the nontradable responses reported below.

tail trade sector includes a wide variety of establishments that cover essential (e.g. grocery stores, drug stores) and luxury (e.g. specialty food stores, jewelry stores) expenditure and everything in between (e.g. auto dealers, furniture stores, clothing stores). Nonetheless, this classification is conservative in the sense that it leaves a large amount of employment unclassified. This is in line with our model calibration, which depends only on having a subset of industries that produce truly nontradable goods. On the other hand, even most manufacturing shipments occur within the same *zip code* (Hillberry and Hummels, 2008), which suggests local consumption demand could impact our measure of tradables. We report robustness to using a classification scheme based on the geographic concentration of employment in an industry.

3 Econometric Methodology

This section provides a formal discussion of causal identification, presents our baseline specification, and discusses the main threats to identification.

3.1 Framework

Motivated by the model in Section 5, we assume a true data generating process of the form:

$$\Delta_{a,t-1,t+h}y = \beta_h[S_{a,t-1}R_{a,t-1,t}] + \Gamma_h' X_{a,t-1} + \epsilon_{a,t-1,t+h}, \tag{1}$$

where $\Delta_{a,t-1,t+h}y = y_{a,t+h} - y_{a,t-1}$ is the change in variable y in area a between $t-1$ and $t+h$, $S_{a,t-1}$ is stock market wealth in area a in period $t-1$ relative to labor market income in the area, $R_{a,t-1,t} = b_{a,t}R^m_{t-1,t} + (1-b_{a,t})R^f_{t-1,t}$ is the measured return on the stock portfolio, $X_{a,t-1}$ collects included covariates determined (from the perspective of a local area) as of time $t-1$, β_h and Γ_h are coefficients (with the latter possibly vector-valued), and $\epsilon_{a,t-1,t+h}$ contains un-modeled determinants of the outcome variable.

Let $\hat{\beta}_h$ and $\hat{\Gamma}_h$ denote the coefficients from treating $\epsilon_{a,t-1,t+h}$ as unobserved and Eq. (1) as a Jordà (2005) local projection to be estimated by OLS. Because the local portfolio betas $\{b_{a,t}\}$ all lie close to 1 and $R^f_{t-1,t}$ is much less volatile than $R^m_{t-1,t}$, we can use the approximation $S_{a,t-1}R_{a,t-1,t} \approx S_{a,t-1}b_{a,t}R^m_{t-1,t}$ in Eq. (1).[11] In that case, Eq. (1) has an approximate shift-share structure with a single national shifter given by the market return

[11] That is, for any (de-meaned) variable $v_{a,t}$, $E[S_{a,t-1}R_{a,t-1,t}v_{a,t}] = E[S_{a,t-1}b_{a,t}R^m_{t-1,t}v_{a,t}] + E[S_{a,t-1}(1-b_{a,t})R^f_{t-1,t}v_{a,t}] \approx E[S_{a,t-1}b_{a,t}R^m_{t-1,t}v_{a,t}]$, where the term $E[S_{a,t-1}(1-b_{a,t})R^f_{t-1,t}v_{a,t}]$ is negligible because $1-b_{a,t} \approx 0$ and $Var(R^f_{t-1,t}) << Var(R^m_{t-1,t})$. In fact, our results below change imperceptibly whether or not we include the term $S_{a,t-1}(1-b_{a,t})R^f_{t-1,t}$.

$R^m_{t-1,t}$, and the identifying assumption for $plim\hat{\beta}_h = \beta_h$ takes the form:

$$E\left[R^m_{t-1,t}\mu_t\right] = 0, \tag{2}$$

where $\mu_t \equiv E\left[S_{a,t-1}b_{a,t}\epsilon_{a,t-1,t+h}\right]$ is a time t cross-area average of the product of the beta-adjusted stock wealth–to–income $b_{a,t}S_{a,t-1}$ and the unobserved component $\epsilon_{a,t-1,t+h}$.[12] Intuitively, this condition will not hold if the outcome variable (e.g., employment or payroll) grows faster for unmodeled reasons ($\epsilon_{a,t-1,t+h} > 0$) in high wealth areas ($\Rightarrow \mu_t > 0$) in periods when the stock return is positive, and vice versa when the stock return is negative.

The econometrics of shift-share designs have recently received renewed attention in Goldsmith-Pinkham et al. (2018) and Borusyak et al. (2018). Condition (2) coincides with the exogeneity condition in Borusyak et al. (2018) in the case of a single national observed shock and multiple (asymptotically infinite) areas and time periods. As in their framework, the condition recasts the identifying assumption from a panel regression into a single time series moment by defining the cross-area average μ_t. Borusyak et al. (2018) defend the validity of shift-share instruments when the shifter is exogenous, a seemingly natural assumption in our setting given that stock market returns are nearly i.i.d. Nonetheless, since stock market returns are equilibrium outcomes (as most shifters are), identification of β_h also requires that other aggregate variables correlated with $R^m_{t-1,t}$ and not controlled for in X impact areas with high and low stock market wealth uniformly. Importantly, we do *not* require that stock market wealth be distributed randomly, and show in Table A.5 that $S_{a,t}$ correlates with the share of a county's population with a college education and the median age, among other variables. Instead, as illustrated by Eq. (2), we require that high and low wealth areas not be heterogeneously affected by other aggregate variables that co-move with stock returns. This insight motivates our baseline specification and the robustness analysis below.

3.2 Baseline Specification

Our baseline specification implements Eq. (1) at the county level and at quarterly frequency, with outcome y either log employment or log quarterly payroll. We include the following controls in $X_{a,t-1}$: a county fixed effect, a state × quarter fixed effect, and eight lags of the

[12] To derive this condition, let $\boldsymbol{Y}$ denote the $AT \times 1$ vector of $\Delta_{a,t-1,t+h}y$ stacked over A areas and T time periods, $\boldsymbol{S}$ the $AT \times T$ matrix containing the vector $\begin{pmatrix} b_{1,t}S_{1,t-1} & \dots & b_{A,t}S_{A,t-1}\end{pmatrix}'$ in rows $A(t-1)+1$ to At of column t and zeros elsewhere, $\boldsymbol{R}$ the $T \times 1$ vector of stock market returns, $\boldsymbol{X}$ the $AT \times K$ matrix of K covariates stacked over areas and time periods, and $\boldsymbol{\epsilon}$ the $AT \times 1$ stacked vector of $\epsilon_{a,t-1,t+h}$. Then we can rewrite Eq. (1) in matrix form as $\boldsymbol{Y} = \beta_h\boldsymbol{SR} + \boldsymbol{X}\Gamma_h + \boldsymbol{\epsilon}$. It follows that $plim\hat{\beta}_h = \beta_h$ if $0 = \lim_{A,T\to\infty}(\boldsymbol{SR})'\,\boldsymbol{\epsilon} = \lim_{A,T\to\infty}\boldsymbol{R}'\boldsymbol{S}'\boldsymbol{\epsilon} = \lim_{A,T\to\infty}\sum_t R_{t-1,t}\sum_a b_{a,t}S_{a,t-1}\epsilon_{a,t-1,t+h} = E\left[R_{t-1,t}\mu_t\right]$. Note that this identification condition presumes a homogenous treatment effect. We explore treatment heterogeneity explicitly in Section 4.4.

"shock" variable $\{S_{a,t-j-1}R_{a,t-j-1,t-j}\}_{j=1}^{8}$. We also include interactions of $S_{a,t-1}$ with changes in other forms of aggregate wealth: the holding return on a 5 year Treasury bond, the log growth of national house prices between $t-1$ and t, and the log change in national labor income and non-corporate business income from $t-1$ to the cumulative total over the next 12 quarters (to capture human capital and private business wealth).[13] Finally, we also include a Bartik (1991) shift-share measure of predicted employment growth at horizon h based only on industry composition, $\Delta_{a,t-1,t+h}e^{B}$.[14] We weight regressions by 2010 population and report standard errors two-way clustered by time and county. Clustering by county accounts for any residual serial correlation in stock market returns and has a small effect on the standard errors in practice. Clustering by time allows for areas with high or low stock market wealth to experience other common shocks and accords with the recommendation of Adão et al. (2019) in the special case of a single national shifter. Finally, we exclude from our baseline sample counties in the top 5% of the share of employees working at large (500+) firms, as these firms can have direct exposure to the stock market.[15]

3.3 Threats to Identification and Motivation for Covariates

Our identifying assumption is that following a positive stock return, areas with high stock market wealth relative to labor income do not experience unusually rapid employment or payroll growth—relative to their own mean and to other counties in the same state, and conditional on the included covariates—for reasons other than the wealth effect on local consumption expenditure. As emphasized by Goldsmith-Pinkham et al. (2018), this require-

[13] Specifically, we interact $S_{a,t-1}$ with (i) the holding return on a 5 year zero coupon Treasury bond using the updated Gürkaynak et al. (2006) data set, (ii) the log change in the Case-Shiller national house price series, and (iii,iv) $\ln\left(\sum_{j=0}^{11}\mathcal{R}^{-j}A_{t+j}\right)-\ln A_{t-1}$ for A_t=aggregate labor compensation (NIPA code A4002C) or aggregate non-corporate business income (nonfarm sole proprietor income and partnership income, NIPA code A041RC) and a quarterly discount factor $\mathcal{R}=1.03^{1/4}$. To see the rationale for the last two controls, let $H_t^{\infty}=\sum_{j=0}^{T}\mathcal{R}^{-j}A_{t+j}$ denote the discounted stream of labor (or private business) income A_t. The revision to human capital (or private business) wealth in period t is $\frac{E_t[H_t^{\infty}]-E_{t-1}[H_t^{\infty}]}{E_{t-1}[H_t^{\infty}]}\simeq \ln E_t\left[H_t^{\infty}\right]-\ln E_{t-1}\left[H_t^{\infty}\right]$. Relative to this definition, our control (i) truncates the horizon at $T=11$ (truncating at longer horizons gives similar results); (ii) replaces $E_t\left[H_t^{11}\right]$ with its perfect-forecast counterpart H_t^{11} (under rational expectations, this provides an unbiased measure of expected wealth); and (iii) replaces $E_{t-1}\left[H_t^{11}\right]$ with income in the last period, A_{t-1}. Under the efficient market hypothesis, this last step does not matter because both $E_{t-1}\left[H_t^{\infty}\right]$ and A_{t-1} are determined in period $t-1$ and therefore should be orthogonal to the stock return $R_{t-1,t}^{m}$.

[14] The Bartik (1991) industry shift-share predicted employment growth between $t-1$ and $t+h$ is defined as $\Delta_{a,t-1,t+h}e^{B}=\sum_{i\in\text{NAICS 3}}\left(\frac{E_{a,i,t-1}}{E_{a,t-1}}\right)\left(\frac{E_{i,t+h}-E_{i,t-1}}{E_{i,t-1}}\right)$, where $E_{a,i,t}$ denotes the (seasonally unadjusted) level of employment in NAICS 3-digit industry i in county a and period t, $E_{a,t}$ is total employment in county a, and $E_{i,t}$ is seasonally-adjusted total national employment in industry i.

[15] Data on payroll by firm size come from the Census Bureau's Quarterly Work Force Indicators. Because this data set has less historical coverage than our baseline sample, we use the time series mean share for each county. This step contains little loss of information because the large payroll share is extremely persistent at the county level, with an R^2 of 0.85 from a regression of the quarterly share on county fixed effects.

ment mirrors the parallel trends assumption in a continuous difference-in-difference design with multiple treatments. Two main threats to identification exist.

The first threat occurs because stock prices are forward-looking, so fluctuations in the stock market may reflect news about deeper economic forces such as productivity growth that independently affect consumption and investment. This "leading indicator" channel confounds interpretation of the relationship between consumption and the stock market in aggregate time series data. Our cross-sectional research design requires only the weaker condition that areas with high and low stock wealth to labor income ratios not load differently on other aggregate variables that co-move with the stock market. Conceptually, such differential loading could occur if stock wealth correlates with other forms of wealth and the return on the stock market correlates with the returns on other forms of wealth. Inclusion in $X_{a,t-1}$ of interactions of $S_{a,t-1}$ with other aggregate variables directly addresses the possible heterogeneity in exposure to changes in four other types of wealth: human capital wealth, non-corporate business wealth, fixed income wealth, and housing wealth.[16] For example, controlling for the interaction of $S_{a,t-1}$ and aggregate earnings addresses the possibility of high wealth areas having different exposure to aggregate earnings risk. Similarly, the Bartik variable controls for the possibility of high wealth counties concentrating in industries with higher stock market betas than those in low wealth counties or in industries that drive overall market returns, and the state-quarter fixed effects control non-parametrically for aggregate shocks that have heterogeneous impacts on different states. Finally, inclusion of the lags of $S_{a,t-1}R_{t-1,t}$ controls for the small serial correlation in stock returns shown in Figure 2a.

The second threat to identification concerns the separation of a consumption wealth effect from firm investment or hiring responding directly to the change in the cost of equity financing. Indeed, the response of total national employment to an increase in the stock

[16]For non-corporate business wealth, fixed income wealth, and housing wealth, we could alternatively try to control directly for changes in the local values of these variables. This alternative has two deficiencies. First, these variables may endogenously respond to local stock market wealth, making them an over-control. Second, measuring local business wealth and fixed income wealth poses a more formidable challenge than measuring local stock market wealth, because of the much larger variation in capitalization factors for the income streams generated by these variables and the particular sensitivity of fixed income wealth to the capitalization factor at interest rates near zero (Kopczuk, 2015; Smith et al., In progress). While this difficulty precludes estimation of the local labor market effects of changes in these other types of wealth, including interactions with the aggregate values of other wealth is still sufficient for identifying the stock market wealth effect. The reason is that heterogeneity in holdings of other wealth matters for our purpose only insofar as returns on such wealth correlate with our main regressor. Formally, denoting by $S^o_{a,t-1}R^o_{t-1,t}$ the change in some other type of wealth o, we can write $S^o_{a,t-1}R^o_{t-1,t} = \gamma S_{a,t-1}R^o_{t-1,t} + S^{o,\perp}_{a,t-1}R^o_{t-1,t}$, where $\gamma S_{a,t-1}$ is the fitted value from a regression of $S^o_{a,t-1}$ on $S_{a,t-1}$ and so by construction $S^{o,\perp}_{a,t-1}$ is orthogonal to $S_{a,t-1}$. Therefore, omitting the part $S^{o,\perp}_{a,t-1}R^o_{t-1,t}$ from the change in wealth of type o has no impact on the remaining variables in the regression (and note that we do not need to separately identify the parameter γ). As an example, interacting the Treasury return with stock wealth directly amounts to allowing for an arbitrary correlation between the levels of stock wealth and fixed income wealth across counties.

market cannot separately identify these two channels. Our local labor market analysis absorbs changes in the cost of issuing equity common across areas into the time fixed effect. Nonetheless, firms in high stock wealth areas may have a cost of capital more sensitive to the value of the stock market. Two aspects of our research design make such a correlation an unlikely driver of our results: (i) we find an employment response in nontradable but not in tradable industries, so differential access to capital markets would have to occur within areas and align with the tradable/nontradable sectoral distinction, and (ii) our baseline sample excludes counties in the top 5% of the share of employees working at large (500+) employee firms that might have greater access to public capital markets.

4 Results

4.1 Baseline Results

In this section we report our baseline results: (i) an increase in the stock market causes faster employment and payroll growth in counties with higher stock market wealth, (ii) the response is pronounced in industries that produce nontradable goods and in residential construction, and (iii) there is no increase in employment in industries that mostly produce tradable goods.

Figure 3 reports the time paths of responses of quarterly employment and payroll to an increase in stock market wealth; formally, the coefficients $\hat{\beta}_h$ from estimating Eq. (1). Table 1 reports the corresponding coefficients and standard errors for $h = 7$, where the stock market return occurs in period 0. Because the stock market is close to a random walk (Figure 2b), these time paths should be interpreted as the dynamic responses to a permanent change in stock market wealth. Panel A of Figure 3 shows no pre-trends in either total employment or payroll, consistent with the parallel trends assumption. Both series start increasing in period 1. Payroll responds more than employment, reflecting either rising hours per employee or rising compensation per hour. The point estimates indicate that a rise in stock market wealth in quarter t equivalent to 1% of labor income increases employment by 0.0077 log point (i.e. an approximately 0.69 basis point increase) and payroll by 0.0218 log point in quarter $t + 7$. The increases appear persistent.

Panels B and C examine the responses in industries classified as producing nontradable or tradable output, respectively. Employment and payroll in nontradable industries rise by more than the total effect. In contrast, the responses in tradable industries are flat following a positive stock market return. The horizon 7 differences between the tradable and nontradable employment and payroll coefficients are both significant at the 1% level.

Table 1: Baseline Results

	All		Non-traded		Traded	
	Emp.	W&S	Emp.	W&S	Emp.	W&S
	(1)	(2)	(3)	(4)	(5)	(6)
Right hand side variables:						
$S_{a,t-1}R_{a,t-1,t}$	0.77*	2.18**	2.02*	3.24**	−0.11	0.71
	(0.36)	(0.63)	(0.80)	(1.01)	(0.64)	(0.74)
Horizon h	Q7	Q7	Q7	Q7	Q7	Q7
Pop. weighted	Yes	Yes	Yes	Yes	Yes	Yes
County FE	Yes	Yes	Yes	Yes	Yes	Yes
State $\times$ time FE	Yes	Yes	Yes	Yes	Yes	Yes
Shock lags	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.66	0.64	0.39	0.48	0.35	0.36
Counties	2,901	2,901	2,896	2,896	2,877	2,877
Periods	92	92	92	92	92	92
Observations	265,837	265,837	263,210	263,210	252,928	252,928

Notes: The table reports coefficients and standard errors from estimating Eq. (1) for $h = 7$. Columns (1) and (2) include all covered employment and payroll; columns (3) and (4) include employment and payroll in NAICS 44-45 (retail trade) and 72 (accommodation and food services); columns (5) and (6) include employment and payroll in NAICS 11 (agriculture, forestry, fishing and hunting), NAICS 21 (mining, quarrying, and oil and gas extraction), and NAICS 31-33 (manufacturing). The shock occurs in period 0 and is an increase in stock market wealth equivalent to 1% of annual labor income. All columns also include eight lags $\{S_{a,t-j-1}R_{a,t-j-1,t-j}\}_{j=1}^{8}$, interactions of $S_{a,t-1}$ with the log change in national labor income and with non-corporate business income from $t-1$ to the cumulative total over the next 12 quarters, the interaction of $S_{a,t-1}$ and the holding return on a 5 year Treasury bond, the interaction of $S_{a,t-1}$ and the log growth of national house prices between $t-1$ and t, and a Bartik (1991) shift-share measure of predicted employment growth. For readability, the table reports coefficients in basis points. Standard errors in parentheses and double-clustered by county and quarter. * denotes significance at the 5% level, and ** denotes significance at the 1% level.

These patterns accord with the predictions of the theoretical model presented in the next section. They also militate against a leading indicator or cost-of-capital explanation since such confounding forces would have to apply only to the nontradable sector.

Figure 4 shows a large response of employment and payroll in the residential building construction sector (NAICS 2361). We show this sector separately because, while it also produces output consumed locally, the magnitude does not easily translate into our theoretical model since the sector produces a capital good (housing) that provides a service flow over many years. Thus, a desire by local residents to increase their consumption of housing services following a positive wealth shock will result in a front-loaded response of employ-

Figure 3: Baseline Results

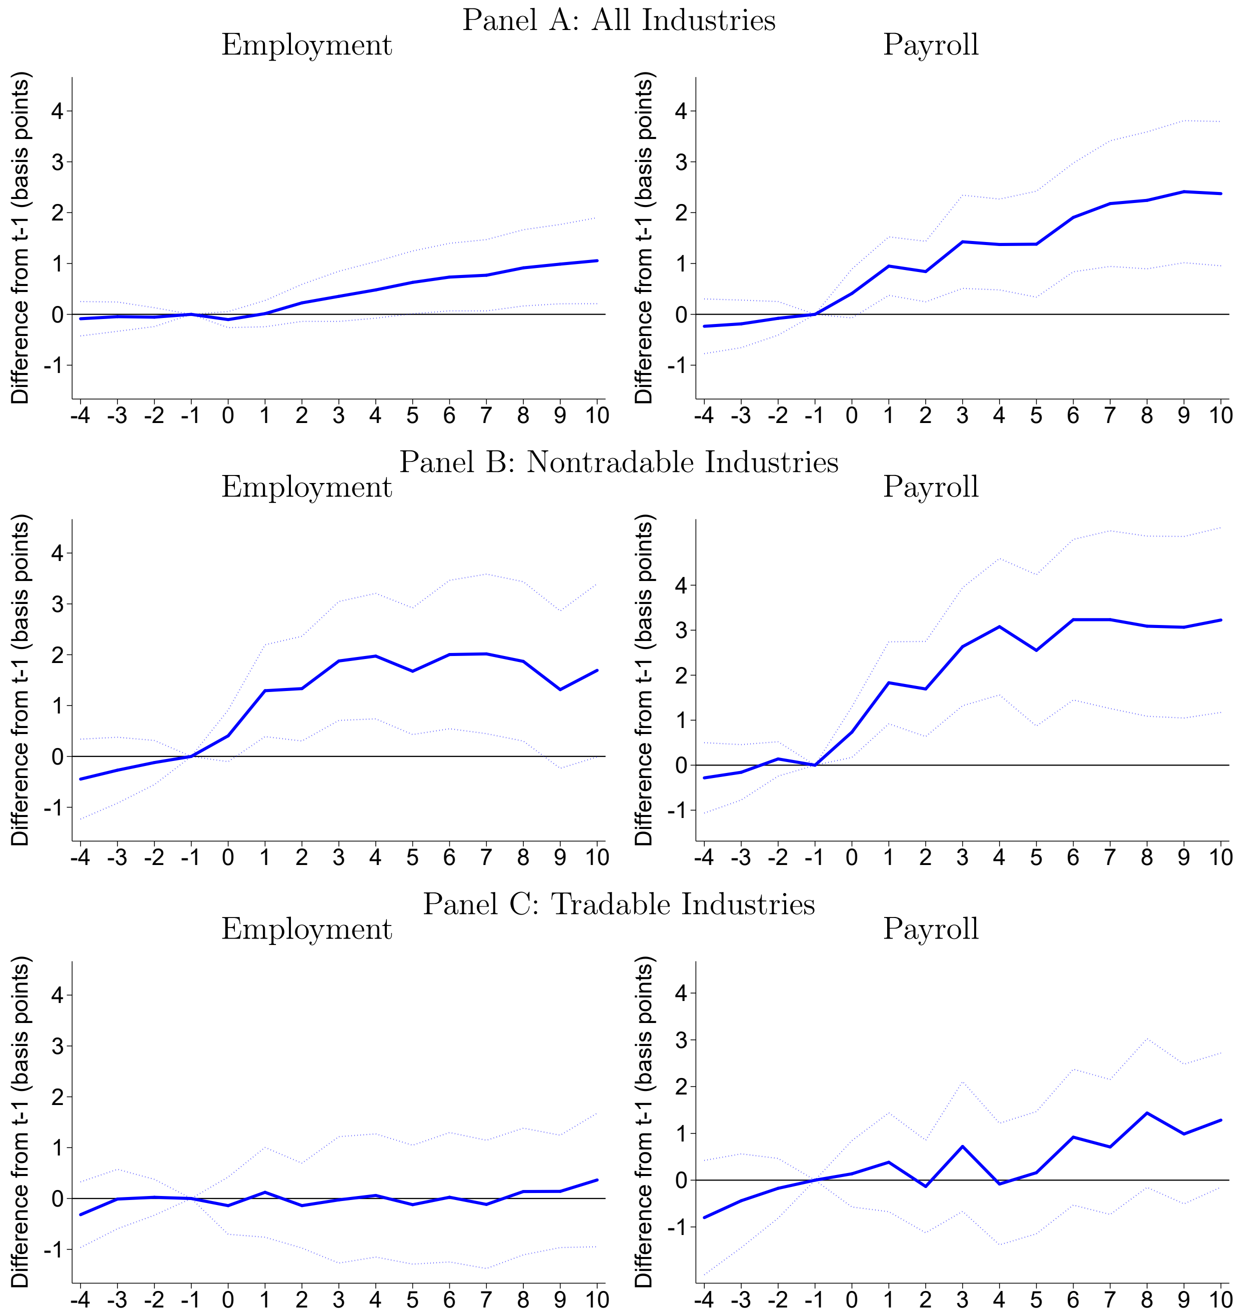

Notes: The figure reports the coefficients β_h from estimating Eq. (1) for quarterly employment (left panel) and wages (right panel) at each quarterly horizon h shown on the lower axis. Panel A includes all covered employment and payroll; Panel B includes employment and payroll in NAICS 44-45 (retail trade) and 72 (accommodation and food services); Panel C includes employment and payroll in NAICS 11 (agriculture, forestry, fishing and hunting), NAICS 21 (mining, quarrying, and oil and gas extraction), and NAICS 31-33 (manufacturing). The shock occurs in period 0 and is an increase in stock market wealth equivalent to 1% of annual labor income. The dashed lines show the 95% confidence bands based on standard errors two-way clustered by county and quarter.

Figure 4: Response of Residential Construction

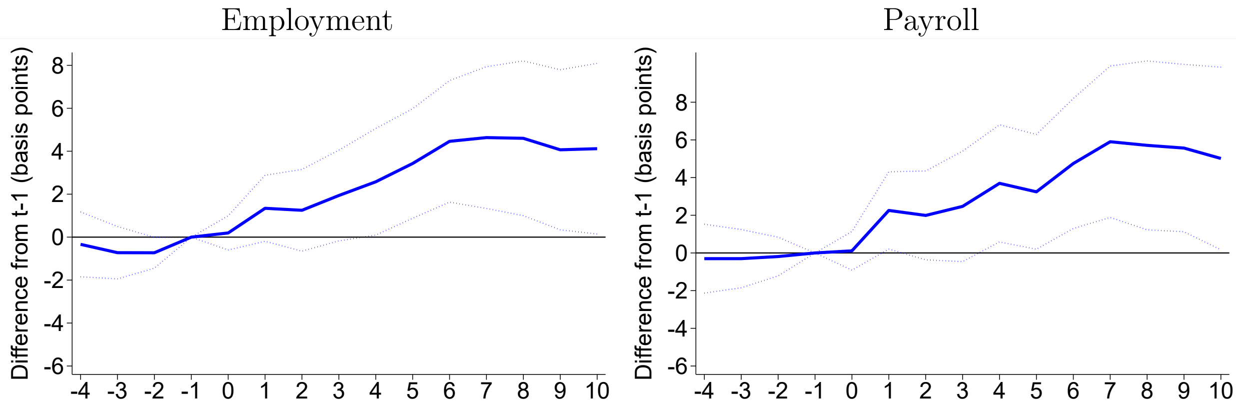

Notes: The figure reports the coefficients β_h from estimating Eq. (1) for residential building construction (NAICS 2361) employment and payroll at each quarterly horizon h shown on the lower axis. The shock occurs in period 0 and is an increase in stock market wealth equivalent to 1% of annual labor income. The dashed lines show the 95% confidence interval bands.

ment in the construction sector. Nonetheless, the large response of residential construction provides additional evidence of a local demand channel at work. We find no corresponding response in construction sectors unrelated to residential building.[17]

Figure 5 reports the response of population. The magnitude lies well below the response of total employment and the data cannot reject no population response at the horizon we examine.[18]

4.2 Robustness

Tables 2 and 3 report results from a number of robustness exercises for the horizon $h = 7$ overall, nontradable, and tradable responses of employment and payroll. The first row of each table reproduces the baseline specification.

Table 2 shows robustness to subtracting or adding covariates to the baseline specification. Rows 2 expands the variation used to identify the response by removing the interactions of $S_{a,t-1}$ with changes in aggregate labor income, non-corporate income, bond wealth, and house prices, and the Bartik control. The results are similar to the baseline specification. The insensitivity reflects a combination of two forces: (i) the loadings on the other aggregate variables do not vary too much with stock wealth, and (ii) as illustrated in Figure 2c, while

[17]In unreported results, we find smaller but statistically significant positive responses in specialty trade contractors (NAICS 238), a category that includes a number of sectors (electrical contractors, plumbers, etc.) involved in the construction of residential buildings. In sharp contrast, there is a flat or slightly negative response in heavy and civil engineering construction (NAICS 237). We also find a large and statistically significant response of new building permits using the Census Bureau residential building permits survey.

[18]The population data by county come from the Census Bureau. The Census reports population as of July 1 of each year. We linearly interpolate these data to obtain a quarterly series.

Figure 5: Response of Population

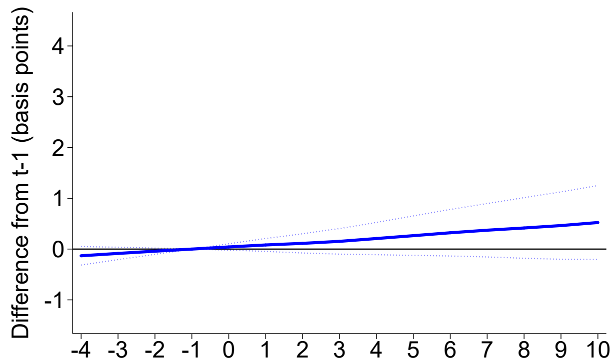

Notes: The figure reports the coefficients β_h from estimating Eq. (1) for total county population at each quarterly horizon h shown on the lower axis. The shock occurs in period 0 and is an increase in stock market wealth equivalent to 1% of annual labor income. The dashed lines show the 95% confidence interval bands.

stock prices are not strictly exogenous, much of the volatility in the stock market and hence the variation in our main regressor occurs for reasons unrelated to these other aggregate factors.

The remaining rows add additional control variables to the baseline specification to address particular concerns. While our baseline specification already includes a linear interaction of stock wealth/income and aggregate labor earnings, previous work has found especially high sensitivity among very high earners (Guvenen et al., 2014). To address this concern, row 3 includes an indicator for being in the top 5% of counties by share of returns with greater than $200,000 in adjusted gross income, interacted with time fixed effects. This row illustrates that controlling flexibly for cyclical patterns of counties with a large share of high earners has a small impact on the coefficients. Motivated by theories of news-driven business cycles (Beaudry and Portier, 2006), row 4 adds an interaction of $S_{a,t-1}$ with the Fernald (2012) measure of TFP growth between $t-1$ and $t+7$, again with little effect. Row 5 adds contemporaneous and 12 lags of local house prices. While our baseline specification controls for the sensitivity of wealthier areas to the aggregate housing cycle, adding the local controls allows this sensitivity to vary with the performance of the stock market.[19] Row 6 controls for the share of payroll in a county at establishments belonging to large (500+ employee) firms interacted with the stock market return. Large firms are more likely to have publicly traded equity and thus experience a direct reduction in their cost of capital when the stock market rises; the stability of coefficients indicates that our results do not reflect an

[19] We use the Federal Housing Finance Agency (FHFA) annual county-level repeat sales house price index and interpolate to obtain a quarterly series. In unreported results, we also find the response of residential construction remains quantitatively robust to controlling for contemporaneous and lags of house price growth so that the construction response does not merely reflect a run-up in local house prices in high wealth areas before the stock market rises.

Table 2: Robustness to Covariates

Dependent variable:	Total		Nontradable		Tradable	
	Emp.	Payroll	Emp.	Payroll	Emp.	Payroll
Specification:						
1. Baseline	0.77*	2.18**	2.02*	3.23**	−0.12	0.71
	(0.36)	(0.63)	(0.80)	(1.01)	(0.64)	(0.74)
2. Only county & stateXquarter FE	1.04*	2.82**	1.21	2.92**	0.02	1.55
	(0.40)	(0.74)	(1.06)	(1.06)	(0.87)	(1.04)
3. Control high earners	0.58	1.64**	1.90*	2.83**	−0.06	0.30
	(0.35)	(0.58)	(0.80)	(0.94)	(0.65)	(0.70)
4. Aggregate TFP sensitivity	0.66*	2.06**	1.85*	3.03**	−0.09	0.75
	(0.31)	(0.62)	(0.76)	(0.95)	(0.59)	(0.71)
5. Control local house prices	0.70+	2.15**	1.50*	2.82**	−0.45	0.35
	(0.37)	(0.65)	(0.64)	(0.92)	(0.62)	(0.70)
6. Control large firm share	0.70*	2.05**	1.93*	3.09**	−0.18	0.67
	(0.34)	(0.59)	(0.76)	(0.95)	(0.61)	(0.70)
7. Control lagged outcomes	0.75*	2.17**	2.10*	3.19**	−0.19	0.75
	(0.36)	(0.63)	(0.80)	(0.99)	(0.70)	(0.73)
8. CzoneXtime FE	1.09**	2.24**	2.41*	3.68**	−0.37	0.24
	(0.38)	(0.67)	(0.96)	(1.11)	(0.64)	(0.86)

Notes: The table reports alternative specifications to the baseline for $h = 7$. The shock occurs in period 0. Each cell reports the coefficient and standard error from a separate regression with the dependent variable indicated in the table header and the specification described in the left-most column. For readability, the table reports coefficients in basis points. Standard errors in parentheses and double-clustered by county and quarter. + denotes significance at the 10% level, * denotes significance at the 5% level, and ** denotes significance at the 1% level.

investment response by such firms. Row 7 includes lagged outcomes to control directly for any pre-trends.[20] Row 8 replaces the state-by-quarter fixed effects with commuting zone-by-quarter fixed effects. In this specification, identification comes from comparing the responses of high and low wealth counties within the same commuting zone. Adding these controls has a minor effect on the point estimates.

Table 3 collects other robustness exercises. Rows 2 and 3 show that the quarters with the most extreme stock returns and the counties with the largest and smallest values of $S_{a,t}$ do not drive the results, although excluding these quarters and counties increases the standard errors. Rows 4 excludes counties in which at least one S&P 500 constituent firm

[20]We include both a county fixed effect and lags of the dependent variable because of the large time dimension (roughly 100 quarters) of the data (Alvarez and Arellano, 2003).

has its headquarters, while row 4 excludes counties headquartering a firm on the Forbes list of the largest private companies. The coefficients remain qualitatively similar, although the payroll responses drop somewhat when excluding S&P 500 headquarter counties. We suggest caution in interpreting these results, however, because these 130 counties account for more than half of total stock wealth and payroll, so that excluding them substantially alters the characteristics of the sample. Rows 5 and 6 show robustness to not weighting the regressions and to trimming at the 1st and 99th percentile of county population.

The next three rows alter the shock variable. Row 8 uses only the price component of the S&P 500 return with similar results. Row 9 instruments $S_{a,t-1}R_{a,t-1,t}$ with $S_{a,t-8}R_{a,t-1,t}$ and row 10 uses the within-county mean ratio of dividend income to labor income interacted with the time-varying price-dividend ratio and return as an instrument. Because the dividend-labor income ratio changes little over time, instrumenting with the lagged wealth variable or fixing this ratio has a small effect on the results.

Row 11 uses an alternative classification of industries into tradable and nontradable based on their geographic concentration. Intuitively, if preferences for output of different industries are similar across locations, then industries with concentrated production must sell to buyers in other regions. This idea traces back at least to Krugman (1991, p. 55) and has been pursued in Ellison and Glaeser (1997), Jensen and Kletzer (2005), and Mian and Sufi (2014), among others. We follow these authors and define a tradability index for industry i as $G_i = \sum_a (s_{a,i} - x_a)^2$, where $s_{a,i}$ denotes the share of employment in industry i located in county a and x_a denotes the share of total employment located in county a, and classify industries in the bottom quartile of this index as nontradable and industries in the top quartile as tradable. We obtain responses very similar to those using our baseline categorization.[21]

The last row returns to the baseline specification but expands the geographic unit to a Core Based Statistical Area (CBSA).[22] The point estimates change little except in the

[21] We construct the index at the NAICS 3 digit level and group industries such that the share of total employment in each quartile is the same. The classification has substantial overlap with our baseline categorization: 7 of the 12 least-concentrated industries are in NAICS 44-45 or 72, and 27 of the 45 most-concentrated industries are in NAICS 11, 21, or 31-33 (the concentrated industries are smaller on average). Even at the 3 digit level, disclosure limitations affect the number of industries reporting employment and payroll in each period. We restrict to county-quarters with the same number of industries reporting non-missing employment and wages in periods $t-1$ and $t+7$, resulting in a final sample about one-half as large as our baseline and explaining why we prefer the simpler 2 digit-based classification for our baseline.

[22] The Office of Management and Budget (OMB) defines CBSAs as areas "containing a large population nucleus and adjacent communities that have a high degree of integration with that nucleus" and has designated 917 CBSAs of which 381 (covering 1,166 counties) are Metropolitan Statistical Areas (MSAs) and the remainder (covering 641 counties) are Micropolitan Statistical Areas (MiSAs). An MSA is a CBSA with an urban core of at least 50,000 people. The remaining counties not affiliated with a CBSA are rural and excluded from the estimation. Because CBSA's may contain counties from multiple states (e.g. the

Table 3: Other Robustness

Dependent variable:	Total		Nontradable		Tradable	
	Emp.	Payroll	Emp.	Payroll	Emp.	Payroll
Specification:						
1. Baseline	0.77*	2.18**	2.02*	3.23**	−0.12	0.71
	(0.36)	(0.63)	(0.80)	(1.01)	(0.64)	(0.74)
2. Keep if $R_{t-1,t} \in [P5, P95]$	1.14*	2.98**	3.52**	5.12**	0.27	1.84
	(0.46)	(0.93)	(1.01)	(1.25)	(0.97)	(1.23)
3. Trim top/bottom 1% of $S_{a,t}$	1.02*	2.93**	2.65*	4.56**	0.52	1.33
	(0.51)	(0.91)	(1.15)	(1.40)	(1.06)	(1.11)
4. Drop S&P 500 HQs	0.30	0.69+	1.68*	1.98**	0.02	0.72
	(0.21)	(0.39)	(0.66)	(0.68)	(0.82)	(0.88)
5. Drop Forbes Top Private HQs	0.40	0.89*	1.88*	2.66**	0.29	0.75
	(0.25)	(0.42)	(0.75)	(0.84)	(0.78)	(0.85)
6. Unweighted	0.47	0.84*	2.98*	2.86**	0.11	0.56
	(0.31)	(0.42)	(1.19)	(1.03)	(0.84)	(1.05)
7. Trim by population	0.82**	1.84**	2.15**	3.04**	0.51	1.52
	(0.31)	(0.56)	(0.78)	(0.91)	(0.74)	(0.94)
8. Price component only	0.74*	2.11**	1.93*	3.11**	−0.15	0.65
	(0.36)	(0.62)	(0.78)	(0.99)	(0.62)	(0.72)
9. IV with lagged wealth	0.76*	1.91**	1.60*	2.61**	−0.23	0.42
	(0.37)	(0.60)	(0.76)	(0.96)	(0.58)	(0.72)
10. IV with fixed dividends/income	0.89**	2.61**	1.63**	2.89**	0.21	1.20*
	(0.12)	(0.18)	(0.21)	(0.25)	(0.42)	(0.51)
11. Concentration-based T/NT	0.77*	2.18**	2.13**	4.26**	−0.72	0.62
	(0.36)	(0.63)	(0.62)	(1.00)	(0.98)	(1.61)
12. Across CBSAs	0.44	1.80+	2.56+	3.00+	0.68	1.95
	(0.47)	(1.03)	(1.53)	(1.63)	(1.61)	(1.73)

Notes: The table reports alternative specifications to the baseline for $h = 7$. The shock occurs in period 0. Each cell reports the coefficient and standard error from a separate regression with the dependent variable indicated in the table header and the specification described in the left-most column. For readability, the table reports coefficients in basis points. Standard errors in parentheses and double-clustered by county and quarter. + denotes significance at the 10% level, * denotes significance at the 5% level, and ** denotes significance at the 1% level.

tradable sectors where they rise slightly, while the standard errors increase substantially. The larger standard errors reflect the decrease in wealth variation after averaging across counties within a CBSA and the smaller sample size. The larger coefficients in the tradable sector could reflect spending on tradable goods produced outside of a resident's county but within the CBSA; however, the data do not reject equality of the coefficients in the county and CBSA specifications.

4.3 Decomposing Variation

In this section we provide evidence on whether certain areas "drive" the results in the sense of Andrews et al. (2017). Consider the specification reported in row 2 of Table 2 in which $X_{a,t}$ includes only a county fixed effect and state-by-quarter fixed effect. In this case, letting $\tilde{z}_{a,t}$ denote $S_{a,t-1}R_{t-1,t}$ demeaned by county and state-by-quarter, $\Delta_{a,t}\tilde{y}$ the outcome after demeaning with respect to county and state-by-quarter (where for notational simplicity we have suppressed the dependence of Δ on the horizon h), π_a the 2010 population in county a, and s index states, we can decompose the OLS coefficient as follows:

$$\beta = \sum_s w_s \beta_s$$

where

$$\beta_s \equiv \left(\sum_{a\in s}\sum_t \pi_a \tilde{z}^2_{a,t}\right)^{-1} \sum_{a\in s}\sum_t \pi_a \tilde{z}_{a,t}\Delta_{a,t}\tilde{y},$$

$$w_s \equiv \left(\sum_{a'}\sum_t \pi_{a'} \tilde{z}^2_{a',t}\right)^{-1} \left(\sum_{a\in s}\sum_t \pi_a \tilde{z}^2_{a,t}\right).$$

Here, β_s is the regression coefficient obtained by using only observations from state s and the weight w_s is the contribution to the total (residual) variation in the regressor from state s.[23] The weights $\{w_s\}$ are all positive and sum to one.

Table 4 reports the ten states with the largest weight in the regression. Not surprisingly, since the regression weights by population, California, Texas, and Florida rank among the states with the highest weights. More surprisingly, Florida, with 6% of the 2010 population, has a weight in the regression above 30%. This high share reflects the large variation across Florida counties in stock market wealth. On the other hand, Florida does *not* drive the

Boston-Cambridge-Newton MSA contains five counties in MA and two counties in NH), the specification in this row replaces the state×quarter fixed effects with quarter fixed effects.

[23] We could have done this decomposition for the baseline specification after partialing out the interactions of $S_{a,t-1}$ with other aggregate variables and the Bartik employment variable. In that case, the coefficient β_s would no longer equate to the coefficient from estimating the regression in state s only because the coefficient on these additional controls would differ across states. The alternative of re-estimating the baseline specification while dropping one state at a time yields conclusions similar to those obtained from Table 4.

Table 4: Ten States with Largest Weight

State	Population share	Weight	β_s, nontradable wage bill
Florida	0.061	0.313	0.30
California	0.121	0.081	5.01
Virginia	0.026	0.050	2.36
Texas	0.081	0.039	1.98
Ohio	0.037	0.034	1.90
North Carolina	0.031	0.032	3.14
Missouri	0.019	0.031	3.23
Illinois	0.042	0.027	7.88
Washington	0.022	0.027	8.46
Maryland	0.019	0.026	5.46

finding of a positive regression coefficient, as the Florida-only nontradable labor bill coefficient is *smaller* than the overall coefficient. Hence excluding Florida from the sample would *raise* the estimated coefficient. Virginia also receives a larger weight in the regression than its population share, reflecting the contrast in the state between wealthier northern suburbs of D.C. and poorer southern counties. Notably, all 10 of the states with the largest weight have $\beta_s > 0$. Thus, no one or two states drive the overall result.

4.4 Heterogeneity

This section reports heterogeneity of the response along the dimensions of whether the stock return is positive or negative, the sample period, and wealth level. For each dimension, we augment Eq. (1) by replacing $\beta_h[S_{a,t-1}R_{a,t-1,t}]$ with $\sum_{m=1}^{M} \beta_h^m \times \mathbb{I}\{o_{a,t} \in m\} \times [S_{a,t-1}R_{a,t-1,t}]$, where $\mathbb{I}\{o_{a,t} \in m\}$ is an indicator for observation $o_{a,t}$ belonging to set m. Figure 6 reports results for the coefficient on nontradable payroll, the variable most directly used in our theoretical analysis.

The left bars show a similar response of nontradable payroll to a negative or positive stock return. Nearly 75% of quarters in our sample contain a positive return, explaining the higher precision around the coefficient on positive returns. The middle bars show the response split before and after the end of the NASDAQ bust. The response is slightly larger in the more recent period, but not statistically significantly different.[24]

[24]Not shown, this pattern holds across other outcomes except total employment, which responds much more strongly in the latter period. Our theory can rationalize a larger response of employment if the more recent period featured greater wage rigidity.

Figure 6: Heterogeneity

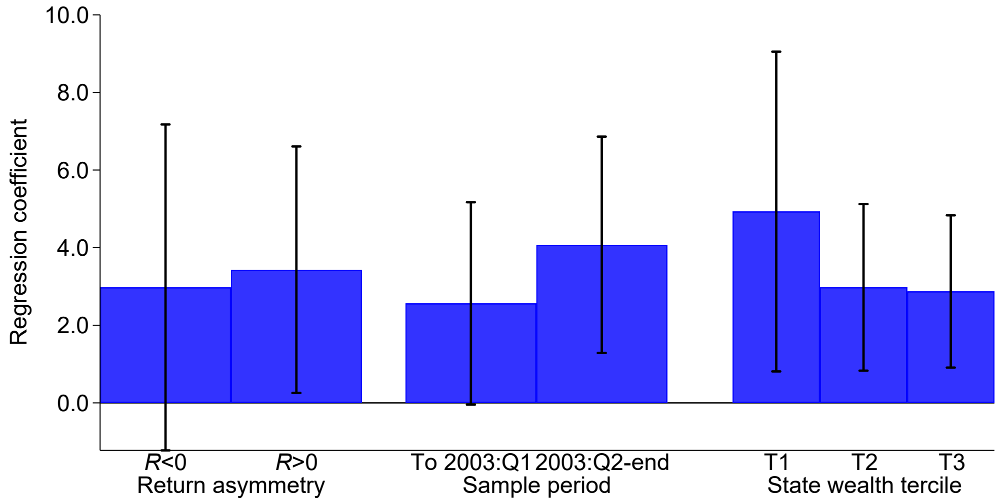

Notes: The figure reports the coefficients β_h^m from estimating Eq. (1) for the nontradable wage bill at horizon $h = 7$, where m indexes positive versus negative stock return (left bars), before or after 2003:Q2 (middle bars), or tercile of the state's per capita wealth distribution (right bars). The whiskers show the 95% confidence intervals.

Finally, many theories of consumption predict higher MPCs for less wealthy households. In the context of stock market wealth, Di Maggio et al. (forthcoming) find a higher MPC in Sweden among households in the lower half of the wealth distribution. In our regional context, such heterogeneity could also arise from local general equilibrium amplication declining in wealth (since, all else equal, a smaller MPC also leads to a smaller multiplier effect). The right bars show that the coefficient indeed declines in tercile of state wealth, although the differences are not statistically significant.[25] This insensitivity may partly reflect that in practice stock wealth heterogeneity is substantially greater within than across counties, and our cross-county analysis already reflects the wealth-weighted MPC within a typical county.

4.5 Labor Income versus Consumption Expenditure

Our analysis so far has focused on the impact on labor market variables. Shortly, we will use economic theory to relate the response of payroll in the nontradable sector to the marginal propensity to consume out of stock market wealth. Before turning to that analysis, we

[25] We split states by tercile of their time-averaged real (deflated by the price index for personal consumption expenditure) dividends per capita. Splitting by state wealth level maintains the identification of each coefficient as coming only from within-state variation. The terciles are: Alabama, Arkansas, Delaware, Georgia, Idaho, Indiana, Kentucky, Louisiana, Mississippi, New Mexico, North Carolina, North Dakota, Oklahoma, Tennessee, Texas, Utah, West Virginia (tercile 1); Alaska, Arizona, California, Hawaii, Iowa, Maine, Michigan, Montana, Nebraska, New York, Ohio, Oregon, Rhode Island, South Carolina, South Dakota, Wisconsin, Wyoming (tercile 2); Colorado, Connecticut, Florida, Illinois, Kansas, Maryland, Massachusetts, Minnesota, Missouri, Nevada, New Hampshire, New Jersey, Pennsylvania, Vermont, Virginia, Washington (tercile 3).

Figure 7: Nontradable Payroll, Consumption Expenditure, and Total Payroll

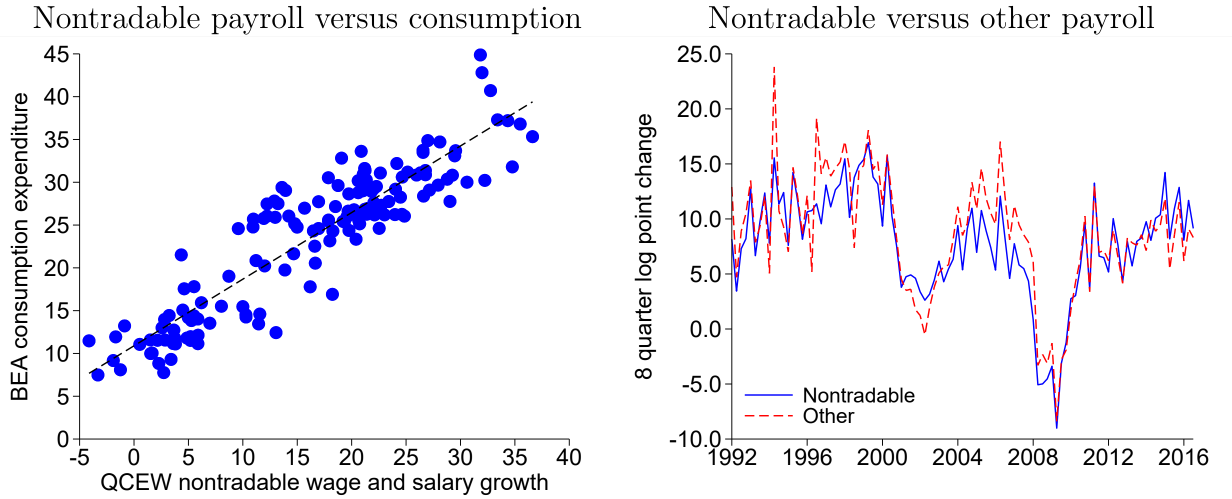

Notes: The left panel presents a scatter plot of five-year log changes in state-level QCEW nontradable wages and salaries and state-level BEA personal consumption expenditure, for each five-year period corresponding to processed quinquennial Economics Censuses (1997-2002, 2002-2007, 2007-2012). The right panel shows 8 quarter log changes in national QCEW wages and salaries in the nontradable sector (NAICS 44-45 and 72) and all other sectors.

present empirical evidence of the relationship between nontradable payroll and consumption expenditure, nontradable sectors and the rest of the economy, and direct evidence of a response of consumption expenditure.

We first show a tight relationship between nontradable payroll and consumption expenditure growth at the state level. The left panel of Figure 7 presents a scatter plot of five-year log changes in state-level QCEW nontradable wages and salaries and state-level BEA personal consumption expenditure (the lowest level of aggregation at which BEA reports consumption expenditure), for each five-year period corresponding to processed quinquennial Economics Censuses (1997-2002, 2002-2007, 2007-2012). We restrict attention to these five-year intervals in which consumption expenditure derives essentially entirely from actual sales data (Awuku-Budu et al., 2016). The two series exhibit a strong positive relationship.

Next, our theoretical analysis in Section 7 will require an assumption of homotheticity across nontradable and other sectors. The right panel of Figure 7 shows evidence of this relationship by plotting 8 quarter log changes of national nontradable payroll and total payroll in all other sectors in the QCEW. At the local level, these two series exhibit different responses to increases in stock wealth, with nontradable payroll rising more sharply. At the national level these series co-move uniformly over time, with a regression coefficient of 0.96 (Newey-West standard error 0.077) and R^2 of 0.79. The similarities in the mean growth rates and high frequency movements of these two series signify homotheticity across locally-nontradable spending and other categories. Intuitively, if the national economy is nearly closed, then all sectors are nontradable nationally and will co-move if homotheticity holds.

Appendix A.5 provides further evidence of preference homotheticity across nontradable and other sectors by extending the Dynan and Maki (2001) analysis of securities-owning households in the Consumer Expenditure Survey. We estimate the effect of the stock market on these households' consumption, separately for their retail expenditure and other expenditure. Consistent with homotheticity, we find similar (cumulative) effects across the two types of expenditure.

Finally, we provide direct evidence of the response of consumption expenditure to stock wealth in Table 5, using the BEA state-level data. These data start in 1997 and have an annual frequency, resulting in a very large reduction in both the cross-section (roughly 3000 counties to 50 states) and time (93 quarters to 18 years) dimensions relative to our baseline, county-quarter specification. Guided by the theoretical model in the next section, we also modify Eq. (1) by replacing $S_{a,t-1}$ with $S^C_{a,t-1}$, defined as the ratio of stock wealth to consumption expenditure in state a and year $t-1$.

We estimate a cross-state coefficient of 4.8. As we will see, this magnitude accords extremely well with the coefficient on nontradable payroll of 3.2 estimated in our baseline specification, providing additional support for the homotheticity assumption and the theoretical mapping of our baseline specification into the MPC out of stock wealth in the next section. From an econometric identification standpoint, this coincidence is remarkable, as our baseline specification uses only within-state variation while Table 5 uses only cross-state variation for identification. However, the coefficient is estimated less precisely than in our baseline, reflecting the large reduction in sample size. Moreover, since we have few clusters in the time dimension (18 years), the conventional clustered standard errors reported in column (2) might be biased. We address this issue by reporting in column (3) the standard error using the "LZ2" bias-reduction adjustment recommended by Imbens and Kolesár (2016) for samples with relatively few clusters and in column (4) the Imbens and Kolesár (2016) suggested degrees of freedom for the t-distribution implied by columns (1) and (3).

5 Theoretical Model

This section develops a stylized theoretical model to interpret the empirical analysis. We present the main equations and results in the main text and relegate additional details to Appendix B. We use the model to illustrate the cross-sectional effects of changes in aggregate stock prices and to validate our empirical specification. In subsequent sections, we calibrate the model and structurally interpret our empirical findings.

We start with a brief overview of the model. There is a continuum of areas denoted by

Table 5: Cross-state Expenditure Results

Coefficient	Conventional two-way clustered standard error	LZ2 two-way clustered standard error	BM degrees of freedom
4.82	1.97	2.85	4.50

The table reports results from estimating $\Delta_{a,t-1,t+h} y = \beta_h [S^C_{a,t-1} R_{a,t-1,t}] + \Gamma'_h X_{a,t-1} + \epsilon_{a,t-1,t+h}$, where y is total consumption expenditure in state a, $h = 2$ years, $S^C_{a,t-1}$ is the ratio of stock wealth to consumption expenditure in state a in period $t-1$, and the remaining variables are analogous to our baseline specification. The first column reports the regression coefficient. The second column reports the standard error clustered by state and year using the conventional degrees of freedom adjustment. Column (3) reports the standard error using the "LZ2" adjustment recommended by Imbens and Kolesár (2016) for samples with relatively few clusters. Column (4) reports the Imbens and Kolesár (2016) suggested degrees of freedom for the t-distribution implied by columns (1) and (3).

subscript a and two time periods denoted by subscripts 0 and 1. We interpret period 1 as the long-run, in which prices adjust and macroeconomic outcomes are determined solely by productivity. In contrast, period 0 is the short-run in which aggregate demand can matter. Hence, a period in the model may correspond to several years. There are two factors of production, labor and capital. Labor is specific to the area in period 0, which ensures that wages and employment in the short run are influenced by local demand. Capital is mobile across areas (in either period), which simplifies the analysis by ensuring that capital has a single price. The price of capital in period 0 is endogenous and can change due to fluctuations in its expected productivity in period 1. Importantly, capital ownership is heterogeneous across areas. We analyze how changes in the price of capital affect local labor market outcomes. We also separately model nontradable and tradable goods, which yields additional predictions and will play an important role in the calibration.

5.1 Environment and Equilibrium

In each period $t \in \{0,1\}$ and area a, a representative household divides its consumption $C_{a,t}$ between a tradable good that can be transported costlessly across areas, $C^T_{a,t}$, and a nontradable good that must be consumed in the area where it is produced, $C^N_{a,t}$, according to the preferences:

$$C_{a,t} = \left(C^N_{a,t}/\eta\right)^{\eta} \left(C^T_{a,t}/\left(1-\eta\right)\right)^{1-\eta}.$$

Competitive firms produce the nontradable good using labor $L^N_{a,t}$ and capital $K^N_{a,t}$ and the Cobb-Douglas technology:

$$Y^N_{a,t} = \left(K^N_{a,t}/\alpha^N\right)^{\alpha^N} \left(L^N_{a,t}/\left(1-\alpha^N\right)\right)^{1-\alpha^N}.$$

Here, $1-\alpha^N$ denotes the share of labor in the nontradable sector.

There are two technologies for producing the tradable consumption good. The first technology uses tradable inputs produced in each area using local labor $L^T_{a,t}$ and capital $K^T_{a,t}$ and the Cobb-Douglas technology:

$$\begin{aligned} Y^T_t &= \left(\int_a \left(Y^T_{a,t}\right)^{\frac{\varepsilon-1}{\varepsilon}} da\right)^{\frac{\varepsilon}{\varepsilon-1}} \\ \text{where } Y^T_{a,t} &= \left(K^T_{a,t}/\alpha^T\right)^{\alpha^T}\left(L^T_{a,t}/\left(1-\alpha^T\right)\right)^{1-\alpha^T}. \end{aligned}$$

The elasticity of substitution $\varepsilon > 0$ governs the effect of unit costs in an area on the exports from that area. The term $1-\alpha^T$ captures the share of labor in the tradable sector.

The second technology uses only capital $\tilde{K}^T_t$:

$$\tilde{Y}^T_t = D_t^{1-\alpha^T}\tilde{K}_t.$$

The productivity parameter D_t determines the rental rate of capital. This technology does not play an important role beyond the asset pricing side of the model. Specifically, we will obtain changes in stock prices in period 0 by varying the future productivity of this technology, D_1. The normalizing power $1-\alpha^T$ simplifies the expressions.

Areas are identical except for their initial capital wealth. The representative household in area a enters period 0 owning $1+x_{a,0}$ units of capital, where $\int_a x_{a,0}da = 0$. We let Q_0 denote the (cum-dividend) price of capital at the beginning of period 0 and normalize the aggregate capital supply to one. Therefore, $(1+x_{a,0})\,Q_0$ denotes the value of capital and, hence, the stock market wealth held by households in area a at the start of period 0. Consequently, the distribution of capital ownership, $\{x_{a,0}\}_a$, determines the cross sectional heterogeneity of stock wealth.

The representative household in each area separates its consumption and labor choices as follows. At the beginning of period 0, the household splits into a consumer and a continuum of workers.[26] The consumer makes a consumption-savings decision to maximize a time-separable log utility function subject to an intertemporal budget constraint:

$$\max_{C_{a,0},C_{a,1}} \log C_{a,0} + \delta \log C_{a,1} \tag{3}$$

[26]We choose to model consumption and labor decisions separately for two reasons. First, assuming workers choose labor according to Greenwood et al. (1988) (GHH) preferences allows us to ignore the wealth effects of labor supply. Second, we can endow consumers with standard time-separable preferences. In addition to simplifying the subsequent expressions, this setup accords with the fact that workers hold relatively little stock market wealth. At the same time, we sidestep some consequences of GHH preferences, such as leading to unplausibly large fiscal and monetary multipliers (Auclert and Rognlie, 2017).

$$\text{s.t.} \qquad P_{a,0}C_{a,0} + \frac{P_{a,1}C_{a,1}}{R^f} = W_{a,0}L_{a,0} + (1 + x_{a,0})\, Q_0 + \frac{W_{a,1}L_{a,1}}{R^f}. \tag{4}$$

Here, $P_{a,t}$ denotes the price level in period t in area a, $W_{a,t}$ the wage level, $L_{a,t}$ labor supply, and R^f the risk-free rate. The elasticity of intertemporal substitution (EIS) of one simplifies the analysis and is empirically plausible (see Appendix B.9 for a discussion of how a more general EIS affects our analysis).

In period 1 (the long run) labor is exogenous, $L_{a,1} = \overline{L}_1$, for all a, and the nominal wage is constant, $W_{a,1} = \overline{W}$. We model period 0 labor supply to incorporate both some degree of wage stickiness and disutility of labor. Specifically, a worker of type ν supplies specialized labor services $L_{a,0}(\nu)$ subject to a constant elasticity labor demand curve—determined by the aggregate demand for labor in the area as well as the elasticity of substitution between specialized labor types.[27] A fraction of the labor types (the sticky workers) supply labor at the preset wage $\overline{W}$ (the same wage as in the long-run). The remainder (the flexible workers) set a wage $W_{a,0}(\nu)$ to maximize:

$$\log\left(C_{a,0} - \frac{\chi}{1+\varphi}\int_0^1 L_{a,0}(\nu)^{1+\varphi}\, d\nu\right),$$

where φ denotes the inverse of the Frisch elasticity of labor supply. Thus, the worker chooses labor according to Greenwood et al. (1988) preferences, which omit a wealth effect on labor supply.

In Online Appendix B.1, we derive the optimal wage set by flexible workers and combine it with the wage of the sticky workers to obtain a labor supply curve (c.f. Eq. (B.17)). We linearize the resulting equation around a benchmark in which all areas have common wealth to derive the log-linear labor supply curve (c.f. Eq. (B.57)):

$$\log\frac{W_{a,0}}{\overline{W}} = \lambda\left(\log\frac{P_{a,0}}{P_0} + \varphi\log\frac{L_{a,0}}{\overline{L}_0}\right). \tag{5}$$

Here, P_0 and $\overline{L}_0$ denote the price level and labor that would obtain if all areas had the same wealth, and $\lambda \in [0,1]$ is a meta-parameter that is an inverse meausure of wage stickiness. When $\lambda = 0$, wages are fully sticky. When $\lambda = 1$, wages are fully flexible and the equation reduces to a neoclassical labor supply relationship between labor and the real wage.[28]

[27]Formally, the worker faces the labor demand curve $L_{a,0}(\nu) = \left(\frac{W_{a,0}(\nu)}{W_{a,0}}\right)^{-\varepsilon_w} L_{a,0}$, where $W_{a,0} = \left(\int_0^1 W_{a,0}(\nu)^{1-\varepsilon_w}\, d\nu\right)^{1/(1-\varepsilon_w)}$ and $L_{a,0} = \left(\int_0^1 L_{a,0}(\nu)^{\frac{\varepsilon_w - 1}{\varepsilon_w}}\, d\nu\right)^{\frac{\varepsilon_w}{\varepsilon_w - 1}}$. Here, $L_{a,0}$ denotes the aggregate demand for labor in area a that obtains in equilibrium.

[28]Letting λ_w denote the fraction of flexible workers that reset wages in period 0, $\lambda = \frac{\lambda_w}{1+(1-\lambda_w)\varphi\varepsilon_w}$.

Finally, at the end of period 0 the household recombines and makes a portfolio decision to allocate savings between capital (stock wealth) and a risk-free asset. The risk-free asset is in zero net supply and generates a gross nominal return in period 1 denoted by R^f. The monetary policy sets R^f to keep labor supply equal to its frictionless level on average.[29] Specifically, it ensures $\int_a L_{a,0} da = \overline{L}_0$, where $\overline{L}_0$ denotes the labor supply that would obtain if all areas had the same stock wealth and there were no wage rigidities. Appendix B.1 completes the description of the setup and defines the equilibrium.

5.2 Consumption Wealth Effect

In Appendix B.2, we characterize the equilibrium and establish the key mechanism behind our results: the consumption wealth effect. Specifically, in view of the preferences in (3), the time-zero consumption expenditure in area a satisfies:

$$P_{a,0} C_{a,0} = \frac{1}{1+\delta} \left(H_{a,0} + \left(1 + x_{a,0}\right) Q_0 \right). \tag{6}$$

Here, $H_{a,0}$ denotes human capital wealth, the present discounted value of labor income. Hence, we have the standard result with log utility that consumption expenditure is a fraction of lifetime wealth.

We now solve for the endogenous variables, first in a benchmark case in which areas have common wealth and then by linearizing the equilibrium equations around that benchmark. We use the common wealth benchmark to illustrate the source of stock price fluctuations, and we use the log-linearized equilibrium to describe the empirical predictions regarding the cross-sectional effects of these fluctuations.

5.3 Common Wealth Benchmark and Stock Price Fluctuations

First suppose all areas have the same stock wealth, $x_{a,0} = 0$ for each a. In this case, the equilibrium allocations and prices are the same across areas, so we drop the subscript a. We solve for the equilibrium in Appendix B.3. We make a parametric assumption on D_0 to ensure that firms are indifferent to using the capital-only technology in period 0 (but they do use it in period 1).[30] In this case, the equilibrium is particularly simple. To state the result,

[29] In practice, monetary policy affects the nominal interest rate by changing the money supply, in an environment where money provides liquidity services and the interest rate reflects the "price" of liquidity. To simplify the exposition, we do not explicitly model money or its liquidity services. These features can be added to the model without changing anything substantive (see Woodford (1998) for further discussion).

[30] For simplicity, we assume the capital-only technology can be used to produce tradables but not nontradables. This provides a potential source of nonhomotheticity across sectors. The assumption on D_0 ensures that production remains homothetic in period 0, which is important for some of our results.

we define the weighted average capital share across the nontradable and tradable sectors,

$$\overline{\alpha} = \eta\alpha^N + (1-\eta)\,\alpha^T. \tag{7}$$

The equilibrium with common wealth is then given by:

$$\begin{aligned}
W_0 &= \overline{W}, \quad L_0 = \overline{L}_0 \text{ where } \overline{L}_0 \text{ solves } (B.38)\,, \\
L_0^N/\overline{L}_0 &= \frac{1-\alpha^N}{1-\overline{\alpha}}\eta \text{ and } L_0^T/\overline{L}_0 = \frac{1-\alpha^T}{1-\overline{\alpha}}(1-\eta)\,. \\
R^f &= R^{f,*} = \frac{1}{\delta}\frac{\overline{L}_1 + D_1}{\overline{L}_0 + D_0}, \\
Q_0/\overline{W} &= D_0 + \frac{D_1}{R^f} = D_0 + \delta\left(\overline{L}_0 + D_0\right)\frac{D_1}{\overline{L}_1 + D_1}, \\
H_0/\overline{W} &= \overline{L}_0 + \frac{\overline{L}_1}{R^f} = \overline{L}_0 + \delta\left(\overline{L}_0 + D_0\right)\frac{\overline{L}_1}{\overline{L}_1 + D_1}.
\end{aligned} \tag{8}$$

The first line shows that the nominal wage is equal to its long-run level and labor supply is given by its frictionless level (see the appendix for a characterization). The second line shows that the share of labor employed in each sector is determined by the sectoral shares in household spending, adjusted by the differences in labor shares across sectors. The third line characterizes the interest rate that brings about this outcome ("rstar").

The last two lines characterize the prices of physical and human capital. An increase in the *future* productivity of capital, D_1, increases the equilibrium price of capital Q_0. Monetary policy responds to this change by raising R^f; however, the equilibrium price of capital increases even after incorporating the monetary policy response.[31]

We focus on the comparative statics of a change in the future productivity of capital from some D_1^{old} to D_1^{new}. By Eq. (8), the price of capital changes from Q_0^{old} to some Q_0^{new}, while leaving the aggregate labor market outcomes unchanged, $L_0 = \overline{L}_0, W_0 = \overline{W}$. We next investigate how this change affects *local* labor market outcomes when stock wealth is heterogeneously distributed across areas. In Appendix B.8, we generalize the model to incorporate uncertainty over D_1 and show that our analysis is robust to other sources of fluctuations in Q_0, such as changes in the level of uncertainty or changes in risk aversion.[32]

[31]Specifically, we have $\frac{dQ_0}{dD_1} = \frac{\delta\overline{W}\left(\overline{L}_0+D_0\right)D_1}{\left(\overline{L}_1+D_1\right)^2} > 0$ and $\frac{dR^f}{dD_1} = \frac{1}{\delta\left(\overline{L}_0+D_0\right)} > 0$. We also have $\frac{d(Q_0+H_0)}{dD_1} = 0$: that is, the interest rate response stabilizes the *total* wealth, $Q_0 + H_0$. This ensures that aggregate spending and thus aggregate employment remains unchanged.

[32]Specifically, we show that a reduction in households' perceived uncertainty about D_1 increases Q_0 and $R^{f,*}$. After extending the analysis to more general Epstein-Zin preferences, we also establish that a decrease in households' relative risk aversion parameter increases Q_0 and $R^{f,*}$ (see Proposition 3). Finally, we show that, conditional on generating the same increase in Q_0, the decline in risk or risk aversion has the same

5.4 Heterogeneous Wealth and Cross-Sectional Predictions

We now derive cross-sectional predictions for the empirically-relevant case of a heterogeneous distribution of stock wealth. We also highlight the properties of the coefficients that will inform our calibration exercise.

We first log-linearize the equations that characterize the equilibrium around the common wealth benchmark for a given D_1. Specifically, we let $w_{a,0} = \log\left(W_{a,0}/\overline{W}\right), p_{a,0} = \log\left(P_{a,0}/P_0\right)$ and $l_{a,0} = \log\left(L_{a,0}/\overline{L}_0\right)$ denote the log-deviations of nominal wages, nominal prices, and total labor for each area. We define $l^N_{a,0}$ and $l^T_{a,0}$ similarly for the nontradable and tradable sectors. In Appendix B.4 we present closed-form solutions for $p_{a,0}, w_{a,0}, l_{a,0}, l^N_{a,0}, l^T_{a,0}$ for a given level of D_1.

In particular, we express local prices in terms of local wages,

$$p_{a,0} = \eta\left(1-\alpha^N\right) w_{a,0}. \tag{9}$$

Combining this with Eq. (5), we obtain a reduced-form labor supply equation:

$$w_{a,0} = \kappa l_{a,0}, \text{ where } \kappa = \frac{\lambda\varphi}{1-\lambda\eta\left(1-\alpha^N\right)}. \tag{10}$$

Here, κ is a composite wage adjustment parameter that combines the effect of inverse wage stickiness, λ, and the inverse labor supply elasticity, φ. The parameter also depends on the share of nontradables, η, and the share of labor in nontradables, $1-\alpha^N$, because these parameters determine the extent to which a change in local nominal wages affects local prices and therefore local real wages.

Our key predictions correspond to the comparative statics as D_1^{old} changes to D_1^{new}. Since the benchmark we log-linearize around does not change, the first-order effect on local labor market outcomes is characterized by changes in log-deviations. We solve for these changes as follows (see Appendix B.5):

$$\Delta\left(w_{a,0}+l_{a,0}\right) = \frac{1+\kappa}{1+\kappa\zeta}\mathcal{M}\left(1-\alpha^N\right)\eta\frac{1}{1+\delta}\frac{x_{a,0}\Delta Q_0}{\overline{WL}_0}, \tag{11}$$

$$\Delta l_{a,0} = \frac{1}{1+\kappa}\Delta\left(w_{a,0}+l_{a,0}\right), \tag{12}$$

$$\Delta\left(w_{a,0}+l^N_{a,0}\right) = \mathcal{M}\frac{1}{1+\delta}\left[\left(1-\overline{\alpha}\right)\frac{x_{a,0}\Delta Q_0}{\overline{WL}_0} + \left(1-\alpha^T\right)\left(1-\eta\right)\Delta\left(w_{a,0}+l^T_{a,0}\right)\right], \tag{13}$$

$$\Delta\left(w_{a,0}+l^T_{a,0}\right) = -\left(\varepsilon-1\right)\left(1-\alpha^T\right)\Delta w_{a,0}, \tag{14}$$

quantitative effects on local labor market outcomes as in our baseline model.

$$\text{where } \mathcal{M} = \frac{1}{1 - (1-\alpha^N)\,\eta/(1+\delta)}$$

$$\text{and } \zeta = 1 + (\varepsilon - 1)\frac{\left(1-\alpha^T\right)^2}{1-\overline{\alpha}}(1-\eta)\,\mathcal{M}.$$

Here, $\Delta y \equiv y^{new} - y^{old}$ denotes the change in equilibrium variable y. In particular, $\Delta Q_0 = Q_0^{new} - Q_0^{old}$ denotes the dollar change in the aggregate stock wealth. Thus, $x_{a,0}\Delta Q_0$ denotes the change in stock wealth in area a relative to other areas. The equations describe how the (relative) stock wealth change normalized by the labor bill, $\frac{x_0 \Delta Q_0}{WL_0}$, affects the (relative) local labor market outcomes in the area.

These equations are intuitive. Eq. (11) shows that an increase in stock wealth in an area increases the total labor bill. To understand the coefficient, note that one more dollar of stock wealth in an area leads to $1/(1+\delta)$ dollars of additional total spending (cf. Eq. (6)), of which $\eta/(1+\delta)$ is spent on nontradable goods produced locally. The increase in spending, in turn, increases the local labor bill by $\left(1-\alpha^N\right)\eta/(1+\delta)$ dollars. This direct effect gets amplified by the local Keynesian income multiplier, denoted by $\mathcal{M}$. The remaining term, $\frac{1+\kappa}{1+\kappa\zeta}$, reflects potential adjustments to the labor bill due to changes in exports to other areas. Specifically, an increase in local wages makes the areas's goods more expensive, which reduces (resp. increases) the tradable labor bill (and thus the total labor bill) when tradable inputs are gross substitutes, $\varepsilon > 1$ (resp. gross complements, $\varepsilon < 1$).

Eq. (12) is a rearrangement of the reduced-form labor supply equation in (10), which relates changes in labor to changes in the labor bill according to the wage adjustment parameter, κ. In particular, how much employment responds relative to the total labor bill (given a change in stock wealth) will discipline κ in our calibration exercise.

Eqs. (13) and (14) characterize the effects on the labor bill separately for the nontradable and tradable sectors. These equations are particularly simple when tradable inputs have unit elasticity, $\varepsilon = 1$. In this case, the effect on the tradable labor bill is zero, $\Delta\left(w_{a,0} + l_{a,0}^T\right) = 0$. The coefficient multiplying the wealth change for the nontradable labor bill can be decomposed into three terms: the partial equilibrium MPC out of stock market wealth $1/(1+\delta)$, the weighted average labor share of income $1-\overline{\alpha}$, and the local multiplier $\mathcal{M}$. In Section 6 we use this decomposition to recover the partial equilibrium MPC given externally calibrated $1-\overline{\alpha}$ and $\mathcal{M}$. Notably, the expression does not require information on the share of nontradables in spending, η, or the share of labor in the nontradable sector, $1-\alpha^N$ (see Section 6 for the intuition).

When $\varepsilon \neq 1$, the decomposition for the nontradable sector does not hold exactly. In this case, as illustrated by Eq. (14), the stock wealth shock can affect the tradable labor bill if it has an effect on wages. As illustrated by Eq. (13), this affects local households' income and, therefore, creates knock-on effects in the nontradable sector (captured by the additional term in brackets). However, if wages do not adjust much, then the tradable adjustment has a small impact on the analysis even when ε is somewhat different from 1.

5.5 Summary and Mapping into the Empirical Analysis

According to Eqs. (11) to (14), an increase in *national* stock prices driven by, e.g., changes in expected *future* productivity of capital or in risk aversion, increases the *current* total labor bill and nontradable labor bill by more in areas with greater stock market wealth. The effect on the tradable labor bill is ambiguous and depends on whether tradable inputs are gross substitutes or complements. In Appendix B.4, we derive the additional predictions that nontradable employment, total employment, and wages weakly increase, and tradable employment weakly falls. All of these predictions accord with our empirical results.

The model also explains the functional form of our empirical regressions. In particular, define $S_{a,0} \equiv \frac{x_{a,0}Q_0}{\overline{WL}_0}$ as area a's (relative) stock wealth divided by its labor bill and $R_0 \equiv \frac{\Delta Q_0}{Q_0}$ as the stock return. Then, we have:

$$S_{a,0}R_0 = \frac{x_{a,0}\Delta Q_0}{\overline{WL}_0}. \tag{15}$$

This variable corresponds to our main regressor, the change in the stock wealth of the area normalized by the local labor bill. Eqs. (11) to (14) illustrate that the empirical coefficients using this regressor have a tight mapping into the key parameters of the model.[33] We next exploit this mapping and provide a structural interpretation of our empirical findings.[34]

6 Calibration and Structural Interpretation

In this section, we use our empirical results from Section 4 to calibrate two key parameters of the model: the strength of the direct stock wealth effect, $\frac{1}{1+\delta}$, and the degree of wage adjustment, κ. We only need two model equations to recover these parameters. Therefore,

[33] In the model, there is only one type of capital so all areas are associated with the same stock return, $R_{a,0} = R_0$ for each a. In the empirical exercise, we allow areas to have heterogenous risky portfolios and thus heterogeneous stock returns, $R_{a,0}$. Eqs. (11) to (14) would naturally generalize to a richer setting that features multiple risky assets and heterogeneous portfolios.

[34] As emphasized by Dynan and Maki (2001), such "dollar-dollar" specifications arise naturally in consumption-wealth models. An alternative approach would be to estimate an elasticity and to convert back into a dollar-dollar coefficient using the sample average ratio of stock market wealth to labor income (or consumption). This alternative has the drawback that the actual ratio varies substantially over time as the stock market booms and busts, a problem noted in the very different context of fiscal multipliers by Ramey and Zubairy (2018).

our calibration also applies in richer models as long as these equations hold. Throughout, we choose the coefficients reported in Table 1 as our calibration targets. As shown in Figure 3, the first few quarters of the impulse response feature sluggish adjustment for reasons outside the model, due e.g. to adjustment costs, consumer habit, or delayed recognition of the stock wealth changes, as found in Brunnermeier and Nagel (2008) and Alvarez et al. (2012). By quarter 7 adjustment is complete and the effect is relatively stable thereafter.

6.1 Direct Stock Wealth Effect

To determine the stock wealth effect parameter, we consider the nontradable labor bill in the special case with $\varepsilon = 1$. To facilitate interpretation, we rewrite Eq. (13) as:

$$\begin{aligned} \Delta\left(w_{a,0} + l^N_{a,0}\right) &= \mathcal{M}\left(1-\overline{\alpha}\right)\rho \times S_{a,0}R_0, \\ \text{where } \rho &= \frac{1}{\mathbb{T}}\frac{1}{1+\delta} \text{ and } S_{a,0} = \frac{x_{a,0}Q_0}{\overline{WL}_0/\mathbb{T}}, R_0 = \frac{\Delta Q_0}{Q_0}. \end{aligned} \tag{16}$$

Here, we have introduced the change of variables $\frac{1}{1+\delta} = \rho\mathbb{T}$, where we interpret ρ as the stock market wealth effect *per year* and $\mathbb{T}$ as the length of period 0 in years. Thus, the denominator of $S_{a,0}$, $\frac{\overline{WL}_0}{\mathbb{T}}$, equals the labor bill *per year* as in the empirical implementation, and the empirical coefficient maps into the stock wealth effect *per year*. In particular, the empirical coefficient can be decomposed into the product of three terms: ρ, the partial equilibrium MPC out of stock market wealth, the weighted-average labor share of income $1-\overline{\alpha}$, and the local Keynesian multiplier $\mathcal{M}$—equivalent to the multiplier on local government spending. We set the weighted-average labor share to a value standard in the literature, $1-\overline{\alpha} = 2/3$, and adjust other parameters to achieve a multiplier $\mathcal{M}$ =1.5, in line with empirical estimates (Nakamura and Steinsson, 2014; Chodorow-Reich, 2019).[35] We then calculate ρ by combining Eq. (16) with the empirical coefficient for the nontradable labor bill.

Specifically, using the coefficient from Table 1, we obtain:

$$\mathcal{M}\left(1-\overline{\alpha}\right)\rho = \frac{\Delta\left(w_{a,0} + l^N_{a,0}\right)}{S_{a,0}R_0} = 3.23\%. \tag{17}$$

[35] To see how we calibrate the multiplier, note that the change of variables in (16) creates one free parameter, $\mathbb{T}$. This parameter is not very meaningful since our model has stylized time periods (it has only two periods). The parameter affects the analysis mainly through its impact on the local multiplier, which is given by:

$$\mathcal{M} = \frac{1}{1-\left(1-\alpha^N\right)\eta/\left(1+\delta\right)} = \frac{1}{1-\left(1-\alpha^N\right)\eta\rho\mathbb{T}}.$$

Therefore, we use $\mathbb{T}$ to calibrate the local multiplier as $\mathcal{M} = 1.5$ given all other parameters. We avoid a literal interpretation of $\mathbb{T}$ and view it as a stand in for other features, such as borrowing constraints, which would affect $\mathcal{M}$ in richer models (see Appendix B.6 for intuition about why $\mathbb{T}$ affects $\mathcal{M}$ in our model).

Substituting $1-\overline{\alpha}=2/3$ and $\mathcal{M}$ =1.5, yields

$$\rho = 3.23\%.$$

Hence, our estimates suggest that a one dollar increase in stock wealth increases household spending by about 3.23 cents per year (at a horizon of two years). The implied magnitude is in line with the yearly discount rates typically assumed in the literature. It is also close to the estimates of the stock wealth effect on consumption for wealthy households in Sweden estimated in Di Maggio et al. (forthcoming).

We make four remarks on this approach. First, it does not depend on the labor supply block of the model. Second, we do not have to parameterize the spending share of nontradables, η, or the labor share in the nontradable sector, $1-\alpha^N$. To understand why, rewrite Eq. (16) as:

$$\frac{\Delta\left(W_{a,0}L_{a,0}^N/\mathbb{T}\right)}{\overline{WL}_0^N/\mathbb{T}}\frac{\overline{WL}_0^N}{\mathbb{T}} = \mathcal{M}\rho\left(1-\alpha^N\right)\eta\left(x_{a,0}\Delta Q_0\right) \text{ where } \frac{\overline{WL}_0^N}{\overline{WL}_0} = \eta\frac{1-\alpha^N}{1-\overline{\alpha}}. \qquad (18)$$

This expression illustrates that the effect of stock market wealth on the nontradable labor bill *in dollars*, $\Delta\left(W_{a,0}L_{a,0}^N/\mathbb{T}\right)$, does depend on both η and $1-\alpha^N$. However, with homothetic preferences and production across sectors, we have $\frac{\overline{WL}_0^N}{\overline{WL}_0} = \eta\frac{1-\alpha^N}{1-\overline{\alpha}}$: that is, the nontradable labor bill as a fraction of the total labor bill reflects the nontradable spending share as well as the sectoral differences in labor share. Therefore, since Eq. (16) normalizes the stock wealth change with the total labor bill, η and $1-\alpha^N$ drop out of the equation. Intuitively, with homothetic preferences these sectors' *average* share of the labor bill proxies for their *marginal* share of changes in the labor bill. As a consequence, the decomposition in (16) is robust to the nontradable spending share as well as the sectoral differences in labor share.[36] Moreover, since the decomposition does not depend on η, we can use it as long as we observe the response in a subset of nontradable sectors.

Third, when $\varepsilon \neq 1$, Eq. (16) applies up to an adjustment (see Eq. (13)). The adjustment reflects the possibility that the change in the tradable labor bill—due to the change in local wages—affects local households' income and creates knock-on effects on the nontradable labor bill. If wages are sufficiently rigid, then the tradable adjustment does not change the

[36] Eq. (18) suggests the decomposition is also robust to (certain types of) cross-county heterogeneity in labor shares. For instance, suppose that areas with high stock wealth ($x_{a,0}>0$) feature greater labor share in nontradables ($1-\alpha_a^N > 1-\alpha^N$)—perhaps because they spend more on high-quality goods that are more labor intensive as recently shown by Jaimovich et al. (2019). Then, the average labor bill of nontradables in these areas is also greater than average ($\overline{WL}_{a,0}^N > \overline{WL}_0^N$). As long as the average labor bill is proportional to the labor share, $\frac{\overline{WL}_{a,0}^N}{\overline{WL}_0^N} = \eta\frac{1-\alpha_a^N}{1-\alpha^N}$, Eq. (18) would still give the decomposition in (16).

analysis by much even if ε is somewhat different from 1. In practice, the value we obtain for κ (described next) implies that there is little loss of generality in ignoring this adjustment for empirically reasonable levels of ε, consistent with the small and statistically insignificant response of tradable payroll we estimate in the data. Therefore, we adopt $\varepsilon = 1$ as our baseline calibration in the main text and relegate the more general case to the appendix.[37]

Fourth, we can compare the ρ of 3.23 obtained from Eq. (17) to the ρ implied by the estimation using state-level consumption data. Following similar steps as in the derivation of Eq. (17), we obtain (see Eq. (B.67) in the appendix)

$$\Delta\left(p_{a,0} + c_{a,0}\right) = \mathcal{M}\rho \times S^C_{a,0} R_0. \tag{19}$$

Here, $p_{a,0} + c_{a,0}$ denotes log nominal consumption expenditure and $S^C_{a,0} = \frac{x_{a,0}Q_0}{P_0 C_0}$ denotes the ratio of area a's (relative) stock wealth to its consumption expenditure. Notably, the labor share does not enter into Eq. (19). Using $\mathcal{M} = 1.5$ and the coefficient from Table 5, we obtain a nearly identical ρ of $4.82/1.5 = 3.21$.

6.2 Wage Adjustment

We use Eq. (12) to determine the wage adjustment parameter κ,

$$\Delta l_{a,0} = \frac{1}{1+\kappa}\Delta\left(w_{a,0} + l_{a,0}\right) \ . \tag{20}$$

Recall that κ is a composite parameter that combines inverse wage stickiness and inverse labor supply elasticity [cf. Eq. (10)]. Therefore, it captures wage adjustment over the estimation horizon. One caveat is that, while the model makes predictions for total labor supply including changes in hours per worker, in the data we only observe employment. A long literature dating to Okun (1962) finds an elasticity of total hours to employment of 1.5. Applying this adjustment and using the coefficients for total employment and the total labor bill from Table 1 yields:

$$\begin{aligned} \frac{\Delta l_{a,0}}{S_{a,0}R_0} &= 1.5 \times 0.77\% \\ \frac{\Delta\left(w_{a,0} + l_{a,0}\right)}{S_{a,0}R_0} &= 2.18\%. \end{aligned}$$

[37] Specifically, in Appendix B.6.2 we consider alternative calibrations with $\varepsilon = 0.5$ and $\varepsilon = 1.5$. In these cases, since trade adjustment affects the analysis, the implied ρ also depends on the share of tradables, η. We allow this parameter to vary over a relatively large range, $\eta \in [0.5, 0.8]$, and show that the implied ρ remains within 5% of its baseline level. As expected, the greatest deviations from the baseline occur when η is low (that is, when the area is more open).

Combining these with Eq. (20), we obtain:

$$\kappa = 0.9. \tag{21}$$

Thus, a one percent change in labor is associated with a 0.9% change in wages at a horizon of two years.[38]

7 Aggregation when Monetary Policy is Passive

We next describe the effect of stock market changes on *aggregate* outcomes. In our model so far, these effects appear only in the interest rate ("rstar") because monetary policy adjusts to ensure aggregate employment remains at the frictionless level. We now consider an alternative scenario in which monetary policy is passive and leaves the interest rate unchanged in response to changes in stock prices. In this case, stock wealth changes affect aggregate labor market outcomes. These aggregate responses are of direct interest to monetary policymakers considering whether or not to accommodate a change in the stock market.

Our aggregation result for the labor bill is straightforward and relies on two observations. First, given homothetic preferences and production across sectors, a one dollar increase in stock market wealth has the same *proportional* effect on the *aggregate total* labor bill and the *local nontradable* labor bill, up to an adjustment for the difference in the aggregate and local spending multipliers. Second, since the aggregate spending multiplier is greater than the local multiplier, we can bound the aggregate effect from below. Therefore, our empirical estimate of the effect on the local nontradable labor bill is a lower bound for the effect on the aggregate total labor bill.

Our aggregation result for labor combines this finding with a third observation: since labor markets are local, the structural labor supply equation (5) remains unchanged as we switch from local to aggregate analysis (as emphasized by Beraja et al. (2016)). The reduced form labor supply equation in (10) changes slightly because shocks impact aggregate inflation and local inflation differently.

To establish these results formally, consider the model from Section 5, but assume that monetary policy keeps the nominal interest rate at a constant level, $R^f = \overline{R}^f$.[39] Appendix

[38] We can also estimate κ from the response of tradable employment [cf. Eq. (B.66)]. Intuitively, tradable employment declines only insofar as local wages and prices rise, so the response of l^T provides information about κ. Auclert et al. (2019) implement this approach in a different empirical setting. We prefer not to rely on this relationship because in practice (unlike in our model) even tradable goods may be influenced by local demand due to home bias, non-zero transportation costs, and supply chains. Nonetheless, the flat response of employment in the industries we classify as tradable in the data accords with a low value of κ.

[39] As before, monetary policy stabilizes the long-run wage level at the constant level, $\overline{W}$.

B.7 extends our theoretical analysis to this case. The aggregate equilibrium with a fixed interest rate is described by the tuple, (Q_0, L_0, W_0, P_0), that solves four equations provided in Appendix B.7. These equations illustrate that changes in the expected productivity of capital, D_1, affect not only the price of capital—as in the baseline model—but also aggregate income, employment, wages, and prices.

To characterize these effects further and to compare them with their local equilibrium counterparts, we log-linearize the equilibrium around the frictionless benchmark. Specifically, we let $\overline{D}_1$ denote the level of capital productivity such that $\overline{R}^f = R^{f,*}$ given $\overline{D}_1$. Considering the equilibrium variables as a function of D_1, and log-linearizing around $D_1 = \overline{D}_1$, we obtain the following equations for the aggregate labor bill and labor:

$$\Delta\left(w_0 + l_0\right) = \mathcal{M}^A \left(1 - \overline{\alpha}\right) \frac{1}{1+\delta} \frac{\Delta Q_0^A}{\overline{WL_0}}, \tag{22}$$

$$\Delta l_0 = \frac{1}{1+\kappa^A} \Delta\left(w_0 + l_0\right), \tag{23}$$

$$\text{where } \mathcal{M}^A \equiv \frac{1}{1 - 1/\left(1+\delta\right)} \frac{1+\kappa^A}{1 - \overline{\alpha} + \kappa^A}$$

$$\text{and } \kappa^A \equiv \frac{\lambda\varphi}{1-\lambda}.$$

Here, $l_0 = \log\left(L_0/\overline{L}_0\right)$ and $w_0 = \log\left(W_0/\overline{W}\right)$ denote log deviations of aggregate employment and wages from the frictionless benchmark. The variable Q_0^A is the log-linear approximation to the exogenous part of stock wealth, $\frac{\overline{W}D_1}{\overline{R}^f}$.[40] As before, $\Delta y \equiv y^{new} - y^{old}$ denotes the change in equilibrium variable y when expected future dividends change. Hence, Eqs. (22) and (23) describe the effect of a change in stock wealth on aggregate labor market outcomes. The parameter $\mathcal{M}^A$ captures the aggregate multiplier. The parameter κ^A captures the degree of aggregate wage adjustment.

Eq. (22) shows that the effect on the aggregate labor bill closely parallels its local counterpart (Eq. (12)), with three differences. First, the direct spending effect is greater in the aggregate than at the local level, $\frac{1-\overline{\alpha}}{1-\delta} > \frac{\left(1-\alpha^N\right)\eta}{1-\delta}$. Intuitively, some spending falls on goods that are tradable across local areas but nontradable in the aggregate. Second, the aggregate labor bill does not feature the export adjustment term $\frac{1+\kappa}{1+\kappa\zeta}$. Third, the aggregate multiplier is greater than the local multiplier, $\mathcal{M}^A > \mathcal{M}$, because spending on tradables (as

[40] The stock price satisfies $Q_0 = W_0 D_0 + \frac{\overline{W}D_1}{\overline{R}^f}$. In this setting, a one dollar increase in $\frac{\overline{W}D_1}{\overline{R}^f}$ increases the equilibrium stock price, Q_0, by more than one dollar. This is because the increase in aggregate demand and output in period 0 also increases the rental rate of capital, $W_0 D_0$. We focus on the comparative statics for a one dollar change in the exogenous component of the stock wealth (as opposed to actual stock wealth) as the appropriate counterfactual scenario for what would happen if monetary policy did not react to an observed stock price shock in an environment where it usually stabilizes the demand effects of these shocks.

well as the mobile factor, capital) diminish the local but not the aggregate multiplier.[41]

Likewise, Eq. (23) shows that the reduced-form labor supply equation closely parallels its local counterpart (cf. Eqs. (12) and (10)). In fact, since labor markets are local, the structural labor supply equation (5) that features prices and labor does not change as we switch from local to aggregate analysis. However, while the aggregate price level moves one-for-one with wages, $p_0 = w_0$, the price level for local consumption does not, since the prices of tradable goods and capital are determined nationally, $p_{a,0} = w_{a,0}\eta\left(1-\alpha^N\right)$ [cf. Eq. (9)]. Therefore, the real wage $w-p$ responds locally but not in the aggregate. The real wage response generates a neoclassical local labor supply response, with strength determined by the magnitude of the Frish elasticity $1/\varphi$, that does not extend to the aggregate level. Rewriting the expressions for κ and κ^A to eliminate the wage stickiness parameter, λ, we obtain:

$$\frac{1}{\kappa} = \frac{1}{\varphi}\left(1-\eta\left(1-\alpha^N\right)\right) + \frac{1}{\kappa^A}. \tag{24}$$

This expression illustrates that the local labor response, $\frac{1}{\kappa}$, combines a neoclassical response to higher real wages, $\frac{1}{\varphi}\left(1-\eta\left(1-\alpha^N\right)\right)$, that only occurs locally, and a term due to wage stickiness that extends to the aggregate, $\frac{1}{\kappa^A}$.

We now use our estimates for the local effects to quantify the aggregate effects on the labor market. We first use Eq. (22) to quantify the effect on the aggregate labor bill. Using the change of variables, $\frac{1}{1+\delta} = \rho\mathbb{T}$, we rewrite this equation as follows:

$$\begin{aligned} \Delta\left(w_0 + l_0\right) &= \mathcal{M}^A\left(1-\overline{\alpha}\right)\rho \times S_0^A R_0^A \\ \text{where } S_0^A &= \frac{Q_0^A}{\overline{WL}_0/\mathbb{T}} \text{ and } R_0^A = \frac{\Delta Q_0^A}{Q_0^A}. \end{aligned} \tag{25}$$

We define S^A as the ratio of aggregate stock wealth to the aggregate yearly labor bill, and R^A as the shock to stock valuations. Hence, $S_0^A R_0^A$ is the aggregate analog of $S_{a,0}R_0$ from the local analysis.

The coefficient in Eq. (25) is the same as its local counterpart in Eq. (16) for the *local nontradable* labor bill, up to an adjustment for the differences in the local and aggregate spending multipliers. Hence, we can combine our estimate for the local nontradable labor

[41] The aggregate spending multiplier is captured by the term $\tilde{\mathcal{M}}^A \equiv \frac{1}{1-1/(1+\delta)}$, which exceeds the local multiplier $\mathcal{M} = \frac{1}{1-(1-\alpha^N)\eta/(1+\delta)}$. In our setting, there is also a second multiplier effect in the aggregate, captured by the term $\mathcal{F}^A \equiv \frac{1+\kappa^A}{1-\overline{\alpha}+\kappa^A} > 1$. This effect emerges because demand-driven fluctuations in our model are absorbed by labor only. We refer to $\mathcal{F}^A$ as *the factor-share multiplier*. The composite multiplier, $\mathcal{M}^A = \mathcal{F}^A\tilde{\mathcal{M}}^A$, combines the standard spending multiplier with the factor-share multiplier. Our model is too stylized to provide an exact mapping between the local and aggregate multipliers. The inequality $\frac{\mathcal{M}^A}{\mathcal{M}} \geq 1$ is a robust feature of settings with constrained monetary policy (Chodorow-Reich, 2019).

bill (for quarter 7) with the inequality $\frac{\mathcal{M}^A}{\mathcal{M}} \geq 1$ to bound the coefficient from below:

$$\mathcal{M}^A (1-\overline{\alpha}) \rho = 3.23\% \frac{\mathcal{M}^A}{\mathcal{M}} \geq 3.23\%.$$

Therefore, if not countered by monetary policy, a one dollar increase in stock valuations increases the aggregate labor bill per year by at least 3.23 cents. Why does the effect on the *local nontradable* labor bill provide information about the implied effect on the *aggregate total* labor bill? With homothetic preferences and production technologies (and ignoring trade effects, $\varepsilon = 1$), a given amount of spending generates the same proportional change on the labor bill in *all* sectors. In particular, the proportional change of the labor bill in the nontradable sectors—which we estimate with our local labor market approach—is the same as the proportional change of the labor bill in the tradable sectors, which we cannot estimate directly due to demand slippage to other regions. Importantly, while clearly convenient for aggregation, the homotheticity assumption also has empirical grounding, as we demonstrated in Section 4.5.

We next quantify the effect on aggregate labor. Using Eqs. (21) and (24) and setting the Frisch elasticity φ^{-1} to 0.5 (Chetty et al., 2012), the nontradable labor share $1-\alpha^N$ to 2/3 (a conservative value), and the nontradable share η to 0.5 (a conservative value), yields $\kappa^A = 1.3$.[42] Then, Eqs. (23) and (25) imply:

$$\Delta l_0 = \frac{1}{1+\kappa^A} \Delta (w_0 + l_0) = \frac{1}{1+\kappa^A} \mathcal{M}^A (1-\overline{\alpha}) \rho \times S_0^A R_0^A. \tag{26}$$

Substituting in the value of κ^A and the response of the labor bill, we obtain:

$$\frac{1}{1+\kappa^A} \mathcal{M}^A (1-\overline{\alpha}) \rho \geq \frac{3.23\%}{1+1.3} \frac{\mathcal{M}^A}{\mathcal{M}} \geq 1.4\%.$$

Therefore, a one dollar increase in stock valuations increases aggregate labor (total hours worked) by the equivalent of at least 1.4 cents (i.e. the labor bill for the additional hours worked is at least 1.4 cents) if monetary policy does not respond.

We can combine these estimates with the ratio of aggregate stock wealth to the aggregate yearly labor bill, S_0^A, to obtain the responses to a stock return, R_0^A. Using data from 2015 (weighting counties by their income), we obtain $S^A = 2.67$.[43] Substituting this value into

[42] As we have emphasized, the nontradable share of consumption expenditure η is a difficult parameter to calibrate given available regional data. Dupor et al. (2019) use the Commodity Flow Survey to estimate that two-thirds of shipments remain within a metropolitan area and 61% remain within a county. This estimate excludes the services component of consumption, which likely has a higher nontradable share. On the other hand, it may include some shipments within a local supply chain that eventually produces a tradable good.

[43] This value coincides almost exactly with the corresponding ratio of 2.63 obtained using C-corporation

Eqs. (25) and (26), we obtain:

$$\begin{aligned}\Delta\left(w_0+l_0\right) &= 3.23\%\frac{\mathcal{M}^A}{\mathcal{M}}\times 2.67\times R_0^A \geq 8.6\%\times R_0^A,\\ \Delta l_0 &\geq 1.4\%\frac{\mathcal{M}^A}{\mathcal{M}}\times 2.67\times R_0^A \geq 3.73\%\times R_0^A.\end{aligned}$$

Therefore, if not countered by monetary policy, a 20% stock return—approximately the yearly standard deviation of the return on the market portfolio—would increase the aggregate labor bill by at least 1.7%, and aggregate hours by at least 0.75%, at a horizon of two years.[44]

8 Conclusion

We estimate the effect of stock market wealth on labor market outcomes by exploiting regional heterogeneity in stock wealth across U.S. counties. An increase in stock wealth in a county increases local employment and the labor bill, especially in nontradable industries but also in total, but does not increase employment in tradable industries. We use a theoretical model to convert the estimated local general equilibrium effect into a household-level MPC out of stock market wealth of around 3.2 cents per year. We also calculate the aggregate general equilibrium effects of the stock wealth consumption channel on the labor market: a 20% change in stock valuations, unless countered by monetary policy, affects the aggregate labor bill by at least 1.7% and aggregate hours by at least 0.75% two years after the shock.

Our estimate for the household-level MPC out of stock market wealth is broadly in line with the quantitative predictions from frictionless models such as the permanent income hypothesis, but considerably smaller than the estimated MPCs out of liquid income found in the recent literature (Parker et al., 2013), even among higher income households (Kueng, 2018; Fagereng et al., 2019). One interpretation is that households that hold stock wealth are affected relatively less by borrowing constraints or by behavioral frictions that increase MPCs. Another possibility is that these households are subject to similar frictions as other households, but stock wealth is associated with more severe transaction costs (such as tax frictions or information frictions) that lead to lower MPCs than other types of liquid income. The latter view is consistent with recent evidence from Di Maggio, Kermani and Majlesi

equity wealth in the FAUS and total wages and salaries in NIPA. This value has increased to 2.89 in 2018.

[44] The magnitude of this calculation changes slightly if we instead assume consumption only responds to changes in taxable stock wealth. In that case, we would recover a larger marginal effect on payroll (intuitively, a larger consumption response would be required to rationalize the same cross-county changes in labor income given smaller wealth), but we would multiply that response by a smaller change in wealth given a 20% change in the stock market. Combining these changes, we would find that a 20% stock return increases the aggregate labor bill by at least 1.3%.

(forthcoming), who argue that Swedish households respond to capital gains significantly less than they respond to dividend payouts.

Our *regional* analysis complements household-level studies of the stock wealth effect by providing direct evidence of stock market affecting labor market outcomes—a key concern for monetary policymakers. Our findings support "the Fed put"—the central banks' tendency to cut interest rates after stock market declines unrelated to productivity (see e.g., Rigobon and Sack (2003); Bjørnland and Leitemo (2009); Cieslak and Vissing-Jorgensen (2017)). Our estimates and aggregation results can be used to calibrate the appropriate interest rate response. If the interest rate is constrained, e.g., due to the zero lower bound or fixed exchange rates, then our analysis implies that stock price declines would induce a sizeable reduction in aggregate labor bill and employment (see Caballero and Simsek (forthcoming) for a related dynamic setup that illustrates the downturn would be further amplified by feedbacks between output and asset prices).

An important question for policymakers concerns the speed at which stock wealth changes affect the economy. We find evidence of sluggish adjustment, with the effect on labor markets starting after 1 to 2 quarters and stabilizing between quarters 4 and 8. This pattern suggests that large stock price declines that quickly reverse course—such as the stock market crash of 1987 or the Flash crash of 2010—are unlikely to impact labor markets, whereas more persistent price changes—such as the NASDAQ boom in the late 1990s or the stock market boom of recent years—have more sizeable effects.

On the other hand, our focus on the consumption channel and our empirical design omit factors that could further increase the effect of stock market wealth changes on aggregate labor markets. First, as discussed by Chodorow-Reich (2019), the Keynesian multiplier effects are likely greater at the aggregate level (when monetary policy is passive) than at the local level. Second, other channels, such as the response of investment, also create a positive relationship between stock prices and aggregate demand (see Caballero and Simsek, forthcoming). Relatedly, while our industry-level analysis mostly focuses on sectors that produce nondurable goods and services, we also find that stock price changes have a large effect on the construction sector. The construction response provides further qualitative evidence that stock wealth affects the economy by changing local demand and inducing an accelerator-type effect on housing investment (see Rognlie et al., 2018; Howard, 2017). We leave a quantitative assessment of these additional factors for future work.

References

Adão, R., Kolesár, M., Morales, E., 08 2019. Shift-Share Designs: Theory and Inference. The Quarterly Journal of Economics 134 (4), 1949–2010.

Alvarez, F., Guiso, L., Lippi, F., 2012. Durable consumption and asset management with transaction and observation costs. American Economic Review 102 (5), 2272–2300.

Alvarez, J., Arellano, M., 2003. The time series and cross-section asymptotics of dynamic panel data estimators. Econometrica 71 (4), 1121–59.

Andrews, I., Gentzkow, M., Shapiro, J. M., 06 2017. Measuring the Sensitivity of Parameter Estimates to Estimation Moments*. The Quarterly Journal of Economics 132 (4), 1553–1592.

Auclert, A., Dobbie, W., Goldsmith-Pinkham, P., 2019. Macroeconomic effects of debt relief: Consumer bankruptcy protections in the great recession.

Auclert, A., Rognlie, M., 2017. A note on multipliers in nk models with ghh preferences. Stanford University working paper.

Auclert, A., Rognlie, M., Straub, L., 2018. The intertemporal keynesian cross. Tech. rep., January 2017 working paper.

Awuku-Budu, C., Gholizadeh, M., Guci, L., Ingber, J. S., Lucas, C. A., Weible, L. J., Nov. 2016. Personal consumption expenditures by state. Survey of Current Business, 1–11.

Barber, B. M., Odean, T., 2000. Trading is hazardous to your wealth: The common stock investment performance of individual investors. The Journal of Finance 55 (2), 773–806.

Barber, B. M., Odean, T., 2001. Boys will be boys: Gender, overconfidence, and common stock investment. The quarterly journal of economics 116 (1), 261–292.

Barber, B. M., Odean, T., 2007. All that glitters: The effect of attention and news on the buying behavior of individual and institutional investors. The review of financial studies 21 (2), 785–818.

Bartik, T. J., 1991. Who Benefits from State and Local Development Policies? W.E. Upjohn Institute for Employment Research.

Beaudry, P., Portier, F., September 2006. Stock prices, news, and economic fluctuations. American Economic Review 96 (4), 1293–1307.

Beraja, M., Hurst, E., Ospina, J., 2016. The aggregate implications of regional business cycles. Tech. rep., National Bureau of Economic Research.

Bjørnland, H. C., Leitemo, K., 2009. Identifying the interdependence between us monetary policy and the stock market. Journal of Monetary Economics 56 (2), 275–282.

Borusyak, K., Hull, P., Jaravel, X., September 2018. Quasi-experimental shift-share research designs. Working Paper 24997, National Bureau of Economic Research.

Bostic, R., Gabriel, S., Painter, G., 2009. Housing wealth, financial wealth, and consumption: New evidence from micro data. Regional Science and Urban Economics 39 (1), 79–89.

Browning, M., Gørtz, M., Leth-Petersen, S., 2013. Housing wealth and consumption: a micro panel study. The Economic Journal 123 (568), 401–428.

Brunnermeier, M. K., Nagel, S., 2008. Do wealth fluctuations generate time-varying risk aversion? micro-evidence on individuals. American Economic Review 98 (3), 713–36.

Caballero, R. J., Simsek, A., forthcoming. A risk-centric model of demand recessions and speculation. The Quarterly Journal of Economics.

Calvet, L., Sodini, P., 2014. Twin picks: Disentangling the determinants of risk-taking in household portfolios. The Journal of Finance 69 (2), 867–906.

Campbell, J. Y., 2014. Empirical asset pricing: Eugene fama, lars peter hansen, and robert shiller. The Scandinavian Journal of Economics 116 (3), 593–634.

Campbell, J. Y., Cocco, J. F., 2007. How do house prices affect consumption? evidence from micro data. Journal of monetary Economics 54 (3), 591–621.

Carroll, C., 2000. Household Portfolios: Theory and Evidence. Ch. Portfolios of the Rich.

Carroll, C. D., Otsuka, M., Slacalek, J., 2011. How large are housing and financial wealth effects? a new approach. Journal of Money, Credit and Banking 43 (1), 55–79.

Case, K. E., Quigley, J. M., Shiller, R. J., 2005. Comparing wealth effects: the stock market versus the housing market. Advances in macroeconomics 5 (1).

Case, K. E., Quigley, J. M., Shiller, R. J., 2011. Wealth effects revisited 1978-2009. Tech. rep., National Bureau of Economic Research.

Chetty, R., Guren, A., Manoli, D., Weber, A., 2012. Does indivisible labor explain the difference between micro and macro elasticities? a meta-analysis of extensive margin elasticities. NBER Macroeconomics Annual 27, 1–56.

Chodorow-Reich, G., May 2019. Geographic cross-sectional fiscal spending multipliers: What have we learned? American Economic Journal: Economic Policy 11 (2), 1–34.

Chodorow-Reich, G., Wieland, J., Forthcoming. Labor reallocation and business cycles. Journal of Political Economy.

Cieslak, A., Vissing-Jorgensen, A., 2017. The economics of the fed put.

Cochrane, J. H., 2011. Presidential address: Discount rates. The Journal of finance 66 (4), 1047–1108.

Cooper, D., Dynan, K., 2016. Wealth effects and macroeconomic dynamics. Journal of Economic Surveys 30 (1), 34–55.

Coval, J. D., Moskowitz, T. J., 1999. Home bias at home: Local equity preference in domestic portfolios. The Journal of Finance 54 (6), 2045–2073.

Davis, M. A., Palumbo, M. G., 2001. A primer on the economics and time series econometrics of wealth effects 9.

Di Maggio, M., Kermani, A., Majlesi, K., forthcoming. Stock market returns and consumption. Journal of Finance.

Dupor, B., Karabarbounis, M., Kudlyak, M., Mehkari, M., 2019. Regional consumption responses and the aggregate fiscal multiplier.

Dynan, K., 2010. Wealth effects and the changing economy. Tech. rep.

Dynan, K., Maki, D., 2001. Does stock market wealth matter for consumption?

Ellison, G., Glaeser, E., 1997. Geographic concentration in u.s. manufacturing industries: A dartboard approach. Journal of Political Economy 105 (5), 889–927.

Fagereng, A., Blomhoff, M., Natvik, G., 2019. Mpc heterogeneity and household balance sheets.

Fagereng, A., Guiso, L., Malacrino, D., Pistaferri, L., November 2016. Heterogeneity and persistence in returns to wealth. Working Paper 22822, National Bureau of Economic Research.

Farhi, E., Werning, I., 2016. Fiscal multipliers: Liquidity traps and currency unions. In: Handbook of Macroeconomics. Harvard University mimeo.

Fernald, J., 2012. A quarterly, utilization-adjusted series on total factor productivity. Federal Reserve Bank of San Francisco Working Paper No. 2012-19.

Galí, J., 2011. The return of the wage phillips curve. Journal of the European Economic Association 9 (3), 436–461.

Gali, J., Monacelli, T., 2005. Monetary policy and exchange rate volatility in a small open economy. The Review of Economic Studies 72 (3), 707–734.

Goldsmith-Pinkham, P., Sorkin, I., Swift, H., March 2018. Bartik instruments: What, when, why, and how. Working Paper 24408, National Bureau of Economic Research.

Graham, J. R., Kumar, A., 2006. Do dividend clienteles exist? evidence on dividend preferences of retail investors. The Journal of Finance 61 (3), 1305–1336.

Greenwood, J., Hercowitz, Z., Huffman, G. W., 1988. Investment, capacity utilization, and the real business cycle. The American Economic Review 78 (3), 402–417.

Guren, A., McKay, A., Nakamura, E., Steinsson, J., 2020. What do we learn from cross-sectional empirical estimates in macroeconomics?

Guren, A. M., McKay, A., Nakamura, E., Steinsson, J., 2018. Housing wealth effects: The long view. Tech. rep., National Bureau of Economic Research.

Gürkaynak, R., Sack, B., Wright, J., 2006. The u.s. treasury curve: 1961 to the present. Finance and Economics Discussion Series 2006-28, The Federal Reserve Board.

Guvenen, F., Ozkan, S., Song, J., 2014. The nature of countercyclical income risk. Journal of Political Economy 122 (3), 621–660.

Hayashi, F., 1982. Tobin's marginal q and average q: A neoclassical interpretation. Econometrica 50 (1), 213–224.

Hillberry, R., Hummels, D., 2008. Trade responses to geographic frictions: A decomposition using micro-data. European Economic Review 52 (3), 527 – 550.

Howard, G., 2017. The migration accelerator: Labor mobility, housing, and aggregate demand. MIT working paper.

Hyslop, R., Imbens, G. W., 2001. Bias from classical and other forms of measurement error. Journal of Business & Economic Statistics 19 (4), 475–481.

Imbens, G. W., Kolesár, M., 2016. Robust standard errors in small samples: Some practical advice. The Review of Economics and Statistics 98 (4), 701–712.

Jaimovich, N., Rebelo, S., Wong, A., 2019. Trading down and the business cycle. Journal of Monetary Economics 102, 96–121, carnegie-Rochester-NYU ConferenceSeries on Public Policy.

Jensen, J., Kletzer, L., 2005. Tradable services: Understanding the scope and impact of services offshoring. Brookings Trade Forum, 75–133.

Jordà, Ò., 2005. Estimation and inference of impulse responses by local projections. American economic review 95 (1), 161–182.

Kent, D., Garlappi, L., Xiao, K., 2019. Monetary policy and reaching for income.

Kopczuk, W., February 2015. What do we know about the evolution of top wealth shares in the united states? Journal of Economic Perspectives 29 (1), 47–66.

Krugman, P., 1991. Geography and Trade. The MIT Press.

Kueng, L., 06 2018. Excess sensitivity of high-income consumers. The Quarterly Journal of Economics 133 (4), 1693–1751.

Kumar, A., 2009. Who gambles in the stock market? The Journal of Finance 64 (4), 1889–1933.

Lettau, M., Ludvigson, S., Steindel, C., 2002. Monetary policy transmission through the consumption-wealth channel. FRBNY Economic Policy Review 5, 117–133.

Lettau, M., Ludvigson, S. C., 2004. Understanding trend and cycle in asset values: Reevaluating the wealth effect on consumption. american economic review 94 (1), 276–299.

Mankiw, N., Zeldes, S. P., 1991. The consumption of stockholders and nonstockholders. Journal of Financial Economics 29 (1), 97 – 112.

Mian, A., Rao, K., Sufi, A., 2013. Household balance sheets, consumption, and the economic slump. The Quarterly Journal of Economics 128 (4), 1687–1726.

Mian, A., Sufi, A., 2011. House prices, home equity-based borrowing, and the us household leverage crisis. American Economic Review 101 (5), 2132–56.

Mian, A., Sufi, A., 2014. What explains the 2007–2009 drop in employment? Econometrica 82 (6), 2197–2223.

Miller, M. H., Modigliani, F., 1961. Dividend policy, growth, and the valuation of shares. the Journal of Business 34 (4), 411–433.

Mitton, T., Vorkink, K., 2007. Equilibrium underdiversification and the preference for skewness. The Review of Financial Studies 20 (4), 1255–1288.

Nakamura, E., Steinsson, J., 2014. Fiscal stimulus in a monetary union: Evidence from us regions. American Economic Review 104 (3), 753–92.

Ogden, R., Thomas, D., Warusawitharana, M., 2016. Corporate equities by issuer in the financial accounts of the united states. Tech. rep., Board of Governors of the Federal Reserve System.

Okun, A., 1962. Potential gnp: Its measurement and significance. Proceedings of the Business and Economics Statistics Section, 98–103.

Paiella, M., Pistaferri, L., 2017. Decomposing the wealth effect on consumption. Review of Economics and Statistics 99 (4), 710–721.

Parker, J. A., Souleles, N. S., Johnson, D. S., McClelland, R., 2013. Consumer spending and the economic stimulus payments of 2008. American Economic Review 103 (6), 2530–53.

Poterba, J. M., 2000. Stock market wealth and consumption. Journal of economic perspectives 14 (2), 99–118.

Poterba, J. M., Samwick, A. A., 1995. Stock ownership patterns, stock market fluctuations, and consumption. Brookings papers on economic activity 1995 (2), 295–372.

Ramey, V. A., Zubairy, S., 2018. Government spending multipliers in good times and in bad: Evidence from us historical data. Journal of Political Economy 126 (2), 850–901.

Rigobon, R., Sack, B., 2003. Measuring the reaction of monetary policy to the stock market. The quarterly journal of Economics 118 (2), 639–669.

Rognlie, M., Shleifer, A., Simsek, A., 2018. Investment hangover and the great recession. American Economic Journal: Macroeconomics 10 (2), 113–53.

Saez, E., Zucman, G., 2016. Wealth inequality in the united states since 1913: Evidence from capitalized income tax data. The Quarterly Journal of Economics 131 (2), 519–578.

Seasholes, M. S., Zhu, N., 2010. Individual investors and local bias. The Journal of Finance 65 (5), 1987–2010.

Shiller, R. J., 1981. Do stock prices move too much to be justified by subsequent changes in dividends? The American Economic Review 71 (3), 421–436.

Smith, M., Zidar, O., Zwick, E., In progress. Top wealth in the united states: New estimates and implications for taxing the rich.

Tobin, J., 1969. A general equilibrium approach to monetary theory. Journal of Money, Credit and Banking 1 (1), 15–29.

Wolf, C., 2019. The missing intercept: A demand equivalence approach.

Woodford, M., 1998. Doing without money: controlling inflation in a post-monetary world. Review of Economic Dynamics 1 (1), 173–219.

Zhou, X., Carroll, C. D., 2012. Dynamics of wealth and consumption: new and improved measures for us states. The BE Journal of Macroeconomics 12 (2).