Review of Economic Dynamics 31 (2019) 347–362

Contents lists available at ScienceDirect

Review of Economic Dynamics

www.elsevier.com/locate/red

Technology shocks and hours revisited: Evidence from household data

Hikaru Saijo [1]

University of California, Santa Cruz, United States of America

ARTICLE INFO

Article history:
Received 1 July 2017
Received in revised form 21 May 2018
Available online 13 September 2018

JEL classification:
E24
E32

Keywords:
Limited stock market participation
Nominal rigidity
Technology shock

ABSTRACT

I exploit heterogeneous impulse responses at the household level due to limited stock market participation to provide novel evidence on the degree of nominal rigidities. A number of studies show that positive technology shocks reduce aggregate hours. The finding is often interpreted as evidence in favor of sticky prices. Using the Consumer Expenditure Survey, I show that, while non-stockholders reduce hours in response to a positive technology shock, stockholders increase them. Aggregate hours fall because most households are non-stockholders. This finding is inconsistent with models featuring a high degree of nominal rigidities.

© 2018 Elsevier Inc. All rights reserved.

1. Introduction

One of the fundamental questions in macroeconomics is to what extent nominal rigidities matter for the transmission of aggregate shocks. For example, the extent to which the central bank can influence the real interest rate and hence economic activity depends on how prices adjust to changes in the nominal rate. Similarly, government spending is likely to have greater effects in stimulating aggregate demand when prices and wages are rigid. In this paper, I exploit heterogeneous impulse responses at the household level due to limited stock market participation and provide novel evidence on the degree of nominal rigidities.

In an influential paper, Galí (1999) finds that positive technology shocks, identified from a long-run restriction in a structural vector auto-regression (VAR), reduce aggregate hours worked. Many researchers view the finding as important because it provides evidence against standard real business cycle (RBC) models in favor of New Keynesian models with weak monetary accommodation to technology shocks. However, it is also well known that if the income effect of labor supply is stronger than the substitution effect, then hours could fall in response to positive technology shocks even in RBC models.[2] That both New Keynesian and RBC models can explain the empirical finding by Galí (1999) suggests that macro data do not provide enough restrictions to differentiate among competing theories that differ in their degree of nominal rigidities.

E-mail address: hsaijo@ucsc.edu.

[1] I thank the Editor, the Associate Editor, and an anonymous referee for useful suggestions. I am also grateful to Susanto Basu, George Bulman, James Cloyne, Francesco Furlanetto, Chad Jones, Amir Kermani, Serguei Maliar, Thuy Lan Nguyen, B. Ravikumar, Carl Walsh as well as conference and seminar participants at End of Year Macroeconomic Conference (Canon Institute for Global Studies, Tokyo), Winter Meeting of the Econometric Society (Boston), and Santa Clara University for helpful comments and discussions. All shortcomings are mine.

[2] For example, Francis and Ramey (2002) construct an RBC model with consumption habit and investment adjustment cost that generates a negative employment response to positive technology shocks. See also Rotemberg (2003) and Lindé (2009).

https://doi.org/10.1016/j.red.2018.09.002
1094-2025/© 2018 Elsevier Inc. All rights reserved.

This leads me to exploit cross-sectional heterogeneities at the household level. Intertemporal substitution of labor supply through capital accumulation is central to modern business cycle theories. For example, in an RBC model, households increase their labor supply in response to a permanent increase in technology in order to reap the benefit of a higher return on investment. In principle, the strength of this intertemporal substitution effect critically depends on households' stock market participation status and hence, differences in conditional movements of hours worked to technology shocks may arise among households that participate in the stock market and those that do not. Because most U.S. households do not participate in the stock market,[3] conclusions based on aggregate hours could thus be misleading. Indeed, when limited stock market participation is taken into account, aggregate hours could fall in response to an improvement in technology in both flexible-price and sticky-price models. However, household-level data are useful in discriminating among alternative models even when aggregate data are not. In flexible-price models, stockholders increase their labor supply because of the standard intertemporal substitution effect, while non-stockholders reduce their labor supply due to the income effect. In contrast, both stockholders and non-stockholders reduce their labor supply in sticky-price models.

To empirically determine the impact of limited stock market participation, I use micro data from the Consumer Expenditure Survey (CEX) to estimate the responses of hours worked by stockholders and non-stockholders to a technology shock identified by a long-run restriction as in Galí (1999). I find that, in response to a positive technology shock, stockholders increase their hours but non-stockholders reduce them. Since most households are non-stockholders, hours worked decline in the aggregate. Thus, the micro data is consistent with the hypothesis that, due to limited stock market participation, the degree of intertemporal substitution of labor supply varies across households, and that the aggregate data masks that heterogeneity. The empirical finding is robust to an alternative specification of using the utilization-adjusted total-factor productivity (TFP) series constructed in Fernald (2014), which controls for heterogeneity across different types of inputs and variations in factor utilization such as labor effort and the workweek of capital.

Importantly, I show that this heterogeneity in impulse responses does not arise among households who hold other assets such as bonds or savings and those who do not. The holding statuses of other assets do not matter because interest rates on these assets do not move following a technology shock. In contrast, the return on stocks increases by 5 annual percentage points in response to a technology improvement. In addition, splits based on the amount of total wealth do not generate the heterogeneous impulse responses I find from the classification based on the households' stock-holding status. The evidence thus supports the idea that stockholders increase their hours in order to reap the benefit of the higher return on investment.

In the second half of the paper, I interpret my findings through the lens of a parsimonious, two-agent dynamic stochastic general equilibrium (DSGE) model with limited stock market participation based on Galí et al. (2007) and Bilbiie (2008). The framework is attractive because it nests a standard representative-agent DSGE model widely used in applied research as a special case. The model features nominal rigidities and thus also nests basic RBC and New Keynesian models as special cases. As in Altig et al. (2011), Dupor et al. (2009), and Liu and Phaneuf (2007), the structural parameters are estimated by matching the empirical impulse responses to technology shocks with the model counterpart. The key difference is that, in addition to standard macro variables, I also match the labor supply responses at the household level estimated from the CEX. I find that the limited stock market participation model is able to replicate the heterogeneous impulse responses and the estimated price and wage rigidities are much smaller than are typically found. In particular, the estimates imply that both firms and households adjust their prices and wages roughly every quarter. In contrast, when the model is estimated under the conventional assumption of full stock market participation using aggregate data only, the estimated frequencies of price and wage adjustments are 4 and 2 quarters, respectively. The exercise thus provides evidence in favor of the transmission mechanism of the limited stock market participation model and underscores its impact on the inference for competing business cycle theories that differ in their degrees of nominal rigidities.

The rest of the paper is organized as follows. After reviewing the relevant literature, in Section 2, which is the main part of this paper, I estimate the household-level labor supply responses to technology shocks using micro data from the CEX. In Section 3, I explore the quantitative implications of my empirical findings by estimating DSGE models with limited stock market participation. Finally, Section 4 concludes with some directions for future research.

1.1. Related literature

This paper is related to an emerging literature that exploits heterogeneous impulse responses at the micro level to understand the transmission of aggregate shocks. In particular, recent papers by Nakamura and Steinsson (2014) and Gorodnichenko and Weber (2016) are also interested in distinguishing competing theories that differ in the degree of nominal rigidities. Nakamura and Steinsson (2014) exploit regional variation in military buildups to estimate the government spending multiplier in a monetary and fiscal union and use it to compare alternative macro models. Gorodnichenko and Weber (2016) utilize differential reactions of conditional volatilities of stock returns for firms with different price adjustment frequencies to a nominal shock to quantify the cost of price adjustment. In my paper, I exploit the variation in labor supply responses to a technology shock across households with different stock market participation statuses in order to test the predictions from alternative theories. Another closely related works are Cloyne and Surico (2017) and Cloyne et al. (2015), who exploit household heterogeneity and study the role of liquidity constraints, proxied by the housing tenure status, on

[3] For example, in the Consumer Expenditure Survey dataset I use, about 14% of all households hold stocks.

the transmission of income tax shocks and monetary policy shocks, respectively. In contrast, I explore the impact of heterogeneity in the degree of intertemporal substitution, proxied by the stock market participation status, on the transmission of technology shocks. Additional studies in this literature include Coibion et al. (2017), who explore the effects of monetary policy shocks on consumption and income inequality, Anderson et al. (2016), who study the impacts of fiscal policy shocks on households with different characteristics such as income and age, and De Giorgi and Gambetti (2017), who characterize the cyclical fluctuations of the consumption data using a factor model.

The paper also belongs to a vast literature that started from Galí (1999) on the labor market effects of technology shocks. Several authors defend or question the identification restriction used in Galí (1999).[4] However, using different methodologies, Basu et al. (2006), Shea (1999), and Canova et al. (2010) reach the same conclusion as in Galí (1999), that positive technology shocks reduce aggregate hours.[5] Some authors emphasize the importance of firm-side heterogeneity to understand the transmission of technology shocks. For example, Chang and Hong (2006) find that technology's effect on hours varies greatly across manufacturing industries and Bocola et al. (2014) argue that existing estimates of neutral technology are biased in the presence of input heterogeneity. My paper, instead, focuses on heterogeneity on the household side.

Finally, this paper is also connected to a research agenda that explores implications of limited asset market participation. Vissing-Jørgensen (2002) and Attanasio et al. (2002) argue that accounting for limited asset market participation in the micro data is important for estimating intertemporal elasticity of substitution (IES) in consumption (see also Guvenen, 2006 for a macroeconomic perspective on this issue). My paper, instead, focuses on the labor supply. The theoretical framework of my paper is based on a parsimonious heterogeneous-agent setting used in, for example, Galí et al. (2007), who introduced non-asset holders into a New Keynesian model to explain the consumption response to a government spending shock.[6] In a related theoretical work, Furlanetto and Seneca (2012) study the labor supply response to a technology shock in a New Keynesian model with limited asset market participation. However, in their model all households are assumed to supply the same amount of hours. As I show below, this assumption is inconsistent with the CEX data. Finally, Bilbiie (2008) shows that limited asset market participation has interesting implications for monetary policy. Similar to my paper, his theoretical analysis emphasizes the importance of the heterogeneity in the intertemporal substitution effect of labor supply.

2. Empirical evidence

2.1. The Consumer Expenditure Survey

The Consumer Expenditure Survey (CEX), which started in 1980, consists of two surveys, the Quarterly Interview Survey and the Diary Survey. The CEX is collected by the Bureau of Labor Statistics (BLS) and provides information on household expenditures, income, employment, and characteristics. I use the Interview portion of the survey, available for download at the website of Inter-university Consortium for Political and Social Research, University of Michigan. In the survey, each household is interviewed 5 times over a 15 month period. The initial interview collects demographic and family information. The second through fifth interviews (total of 4 interviews) collect information for the 3 months prior to the interviews. Because all households are asked about their financial information in the fifth interview, the data allows me to identify their asset market participation status in a relatively precise manner.[7]

I construct a pseudo-panel by averaging total hours worked by households identified by their asset-holding status. Households are asked about their amounts in "Checking accounts, brokerage accounts and other similar accounts", "Savings accounts at banks, savings and loans, credit unions, etc.", "Stocks, bonds, mutual funds and other such securities", and "U.S. savings bonds". In standard business cycle models, intertemporal substitution made possible through capital accumulation is an important channel through which technology shocks affect labor supply. Thus, I start my empirical analysis by comparing the labor supply of households that hold stocks to those that do not. I will also compare hours worked for classifications based on their holding statuses of other assets, such as checking or savings accounts. I classify households with positive responses to "Stocks, bonds, mutual funds and other such securities" as stockholders and the rest as non-stockholders. As is evident from this categorization, not all households with the positive response hold stocks but most likely do.[8] The inability to perfectly identify stockholders should bias against finding difference in the labor supply behavior between the two groups.

[4] See Francis and Ramey (2002) and Fernald (2007) for papers that defend it and Christiano et al. (2003), Uhlig (2004), and Chari et al. (2008) for papers that question it. Erceg et al. (2005) conduct Monte Carlo simulations and conclude that the approach by Galí (1999) performs reasonably well with several caveats.

[5] See, however, Alexopoulos (2011) who finds that aggregate hours rise in response to technology improvements, which are measured based on books published in the field of technology.

[6] Recent works that use a similar framework include Broer et al. (2016), Walsh (2016), and Saijo (2018). Broer et al. (2016) focus on the transmission mechanism of a monetary policy shock while Walsh (2016) studies impulse responses to various shocks and the welfare implications. Both papers assume that agents who receive capital income do not supply labor. Saijo (2018) shows that limited capital market participation amplifies the macroeconomic impact of uncertainty shocks to fiscal policy.

[7] In contrast, in the Panel Study of Income Dynamics (PSID), financial information is collected only in certain years (1984, 1989, 1994, 1999, 2001, 2003, 2005, and 2007). An additional disadvantage of the PSID data is that full information about household labor supply is not available after 1993.

[8] For example, in 2008, according to Investment Company Institute (2008), 96% of households who own either stocks or bonds hold stocks.

Table 1
Sample sizes, CEX data, 1980–2011.

Group	Mean number of obs. per quarter	Education		Birth cohort		
		no college	college	young	middle	old
All households	4195	3164	1031	1721	1429	1044
Stockholders	594	312	282	191	244	159
Non-stockholders	3600	2851	749	1530	1185	885

Notes: "young", "middle", and "old" refer to households whose heads were born 1955–, 1935–1954, and –1934, respectively.

An important consideration in examining the causality of stock-holding status to the labor supply response is the issue of reverse causality. For example, households that increase their labor supply for whatever reason choose to invest in stocks when they did not before, rather than their stock market participation leading them to increase their labor supply in response to a technology shock. To address this issue, I define stockholders as households that were holding stocks *a year before* and non-stockholders as households that were not holding stocks *a year before* at each point in time. This categorization effectively allows me to estimate the responses of households based on their pre-shock stock market participation status as far as the initial responses to the shocks are concerned. To do this, I employ two additional variables in the CEX. The first variable asks whether a household's amount of stockholding remained the same, increased, or decreased from a year prior. The second variable asks the difference in the estimated market value of a household's stock holding from a year ago. I define a household as holding stocks a year before if the household (i) reports having the same amount of stock holding from a year before and reports a positive stock holding for the current year, (ii) reports having had an increase in the amount of stock holding but the amount of the increase is less than the amount of stock holding reported for the current year, and (iii) reports having had a decrease in the amount of stock holding. I define a household as not holding stocks a year before if the household (i) reports having the same amount of stock holding from a year before and reports zero stock holding for the current year and (ii) reports having had an increase in the amount of stock holding but the amount of the increase is equal to or more than the amount of stock holding reported for the current year.

Table 1 reports the sample sizes for different groups in the CEX data. About (3600/4195≈) 86% of households are non-stockholders.[9] The Table also shows the demographic properties of the stockholders and non-stockholders. Non-stockholders tend to be less educated and they tend to skew towards a younger birth cohort.[10] However, it is also important to note that a significant fraction of households whose heads have college degrees (749/1031≈73%) and households whose heads belong to the older birth cohort (885/1044≈85%) are non-stockholders.

2.2. Hours worked from the CEX

To measure annual total hours worked, I use two sets of information collected in the CEX: (i) hours worked per week and (ii) number of weeks worked full or part time in the last 12 months, both by head and by spouse. I define annual total hours worked as a product of (a) the cross-sectional average of hours worked per week, averaged over 4 quarters, and (b) the cross-sectional average of the number of weeks worked in the last 12 months, which I use only the information collected in the fourth quarter.[11]

Example. To understand the procedure, consider how stockholders' annual total hours worked in 1985 is constructed. First, I calculate the average hours worked per week, H_{1985}, using all surveyed households in 1985 identified as being stockholders a year prior. Second, I calculate the total number of weeks worked over the past 12 months, W_{1985}, using households surveyed in 1985:Q4 identified as being stockholders a year prior. The product $TH_{1985} = H_{1985} \times W_{1985}$ is the stockholders' hours worked in 1985. Similarly, I can compute stockholders' hours worked in 1986: $TH_{1986} = H_{1986} \times W_{1986}$. I repeat the procedure and obtain the time series of stockholders' annual hours worked.

Remark 1. Since the CEX does not ask the number of weeks worked in the last 3 months, it is not possible to measure quarterly total hours worked without additional assumptions. To see this, consider the problem of extracting quarterly numbers of weeks worked. For simplicity, suppose there is only one household. Suppose further that they are interviewed on 1985:Q4, 1986:Q1, 1986:Q2, and 1986:Q3. (Recall that each household is interviewed 4 times about their employment.) For example, in the interview conducted in 1985:Q4 the household reports the total number of weeks worked from 1985:Q1 to 1985:Q4 and in the interview conducted in 1986:Q1 the household reports the total number of weeks worked from

[9] As discussed in Vissing-Jørgensen (2002), the percentage of stockholders is smaller than in other sources. This may be because households who hold stocks through their pension plan do not report them. While this would lead them to be mis-categorized as non-stockholders, the mis-categorization should bias against finding differences between stockholders and non-stockholders.

[10] I classify households based on birth cohorts rather than age cohorts to be consistent with the microeconometric literature which argues that the classification for the grouping estimator should be constant or predictable over time.

[11] As recommended by the BLS, I use population weights.

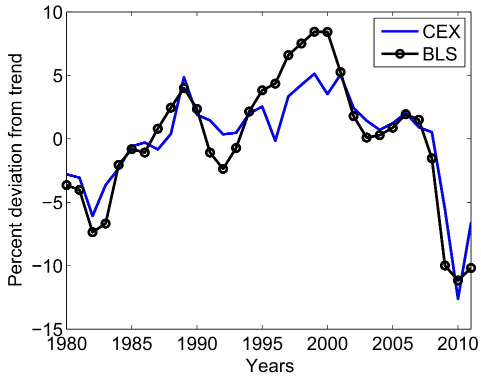

Fig. 1. Aggregate hours: CEX and BLS data. *Notes*: Time series of aggregate hours from the CEX and the BLS data. Both series are detrended using log deviations from linear trends and are multiplied by 100 so that they are expressed in percentage terms.

1985:Q2 to 1986:Q1 and so on. This means that there are three degrees of freedom: A researcher needs to infer quarterly numbers of weeks worked from 1985:Q1 to 1986:Q3 (7 quarters) from 4 interviews. One could potentially sidestep this issue by, for example, assuming that the household worked the same number of weeks from 1985:Q1 to 1985:Q4. In practice, however, I found that this method tends to produce quarterly numbers that are noisy and quite sensitive to the assumption being made.

Remark 2. A potential way to aggregate annual total hours worked would be to take the product of (a) and (b) for each household first and then take the cross-sectional average. An important caveat of this approach is that, while the entrance and exit of households to the survey happens every quarter, we can only use households who started reporting their weekly hours in the first quarter of the year. For example, if a household reported their hours during 1985:Q3, 1985:Q4, 1986:Q1, and 1986:Q2, then we cannot use the household's response to compute annual total hours worked because we do not know the household's weekly hours during 1985:Q1 and 1985:Q2 (if we want to compute annual hours for 1985) or during 1986:Q3 and 1986:Q4 (if we want to compute annual hours for 1986). In order to avoid the small sample problem, I thus take the cross-sectional averages of (a) and (b) separately across households and then take the product of the household averages of (a) and (b). A concern here is that the cross-sectional average of total hours,

$$TH = \frac{1}{N}\sum_{i=1}^{N}(H^i \times W^i),$$

where N is the total number of households i in the sample, is in principle different from the object being calculated here:

$$\widetilde{TH} = \frac{1}{N}\sum_{i=1}^{N}H^i \times \frac{1}{N}\sum_{i=1}^{N}W^i.$$

The magnitude of the bias arising from this procedure could be evaluated by comparing total hours series constructed from the CEX with that from the Bureau of Labor Statistics (BLS).

Remark 3. As shown in Fig. 1, annual hours for all households tracks the business sector hours series provided by the BLS quite well; the correlation between the two series (1980–2011), detrended using log deviations from linear trends, is 0.92. This increases the confidence that CEX is measuring household labor supply quite accurately.

2.3. Measuring the household-level responses to technology shocks

To measure the household labor supply responses to a technology shock, I first consider a long-run restriction using labor productivity as in Galí (1999). Specifically, assuming there are technology shocks ν_t^z and a vector of non-technology shocks ν_t^x, log labor productivity growth Δpr_t and a vector of the growth rates of the variables of interest Δx_t can be expressed in a MA form as

$$\begin{bmatrix} \Delta pr_t \\ \Delta x_t \end{bmatrix} = \begin{bmatrix} C^{11}(L) & C^{12}(L) \\ C^{21}(L) & C^{22}(L) \end{bmatrix} \begin{bmatrix} \nu_t^z \\ \nu_t^x \end{bmatrix},$$

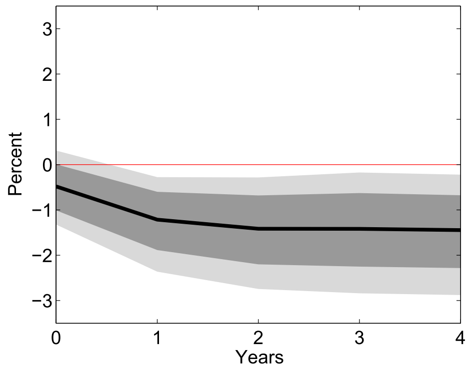

Fig. 2. Aggregate hours decline in response to a positive technology shock. *Notes*: Impulse responses of the levels of hours to a one-standard-deviation positive technology shock identified from labor productivity using a long-run restriction as in Galí (1999). The dark and light shaded areas are 68 and 90-percent confidence intervals, respectively. The sample period is 1981–2011.

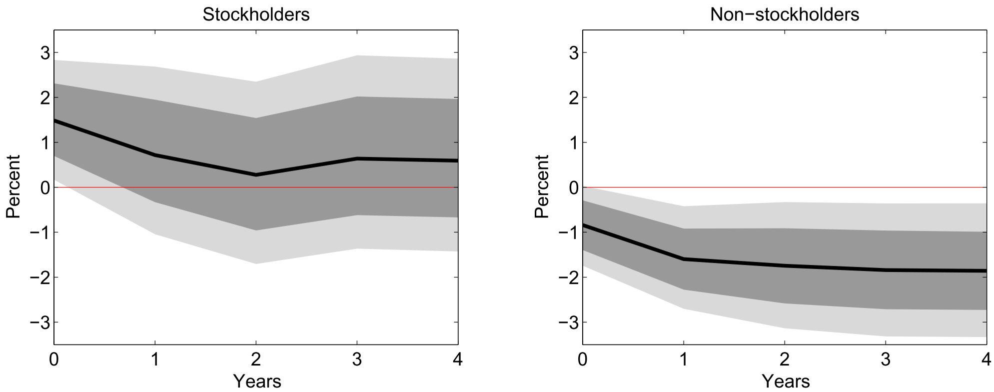

Fig. 3. In response to a positive technology shock, stockholders increase their hours but non-stockholders reduce them. *Notes*: Impulse responses of the levels of hours to a one-standard-deviation positive technology shock identified from labor productivity using a long-run restriction as in Galí (1999). The dark and light shaded areas are 68 and 90-percent confidence intervals, respectively. The sample period is 1981–2011.

where I impose $C^{12}(L) = 0$, i.e., only the technology shock has a permanent effect on labor productivity. I estimate the VAR on one lag using annual data from 1981–2011,[12] where x includes all households, stockholders, and non-stockholders: $\Delta x_t = [\Delta H_t^{all} \Delta H_t^{stock} \Delta H_t^{nonstock}]'$.

Fig. 2 reports impulse response to aggregate hours (hours worked by all households) to a one-standard deviation improvement in technology. On impact, hours decline by about 0.5 percent and stays persistently below the pre-shock level. The reduction of aggregate hours in response to a positive technology shock is consistent with the previous research such as Galí (1999) and Basu et al. (2006). Fig. 3 presents the central result of this paper. When technology improves, stockholders increase their hours worked on impact while non-stockholders reduce them. The initial responses are both statistically significant according to the bootstrap confidence intervals.

Another popular way of measuring the effect of a technology shock is to use the utilization-adjusted TFP series constructed in Fernald (2014). The series controls for heterogeneity across different types of inputs and variations in factor utilization such as labor effort and the workweek of capital. To estimate the conditional response of hours worked to technology shocks, I run a regression on current and lagged technology and lagged hours growth:

$$\Delta x_t = a + \sum_{i=1}^{I} b_i \Delta x_{t-i} + \sum_{j=1}^{J} c_j \Delta z_{t+1-j} + e_t, \quad (1)$$

where Δx is the growth rate of the variable of interest and Δz is the technology growth measured using the utilization-adjusted TFP by Fernald (2014). I set $I = J = 5$. Similar regression specifications have been adopted in, for example, Romer and Romer (2004). The advantage of including lagged dependent variables is that it sharpens estimates in short samples since it controls for dynamics of hours caused by other shocks. Finally, note that while the technology growth is a generated regressor, Coibion and Gorodnichenko (2012) have shown that explicitly adjusting for standard errors for the presence of

[12] I lose the first year of the CEX data because I use growth rates.

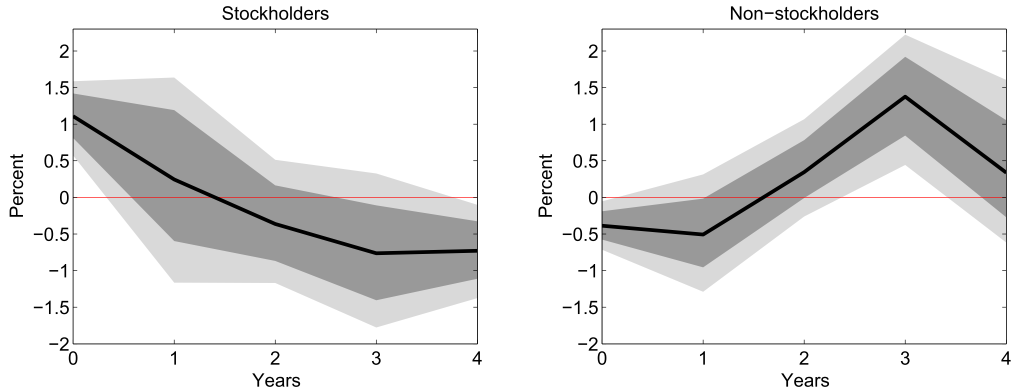

Fig. 4. Technology shocks identified using Fernald (2014) series. *Notes*: Impulse responses of the levels of hours to a one-standard-deviation positive technology shock, estimated from a regression on its own lags and current and lagged technology growth. Technology growth is measured using the utilization-adjusted TFP series constructed in Fernald (2014). The dark and light shaded areas are 68 and 90-percent confidence intervals, respectively. The sample period is 1981–2011.

generated regressors has a minor impact in this environment because the technology growth is a residual from the first stage and not the fitted value.

Fig. 4 shows that the main result is robust; in response to a positive technology shock, stockholders increase their labor supply and non-stockholders reduce them, where the initial responses are both statistically significant.[13]

For the rest of the empirical analysis and the structural estimation, I examine the responses of variables of interest to a technology shock using the utilization-adjusted TFP by Fernald (2014) based on the regression (1) instead of a long-run restriction as in Galí (1999). The utilization-adjusted TFP approach has two advantages. First, it facilitates comparison across different specifications since the technology process is estimated outside the regression. In contrast, in the long-run restriction the technology process is jointly estimated along with the responses of variables of interest. Second, the utilization-adjusted TFP approach allows me to simultaneously look at many variables; this property is useful for the structural estimation I conduct below. The long-run restriction approach requires me to estimate a large number of VAR coefficients as more variables are added. This could potentially make the inference imprecise especially in small samples.

CEX also asks households about variables other than labor supply, such as consumption and wages, which could potentially be informative. For consumption (defined as expenditures on non-durables and services), the responses of stockholders and non-stockholders are similar. They both increase their levels of consumption around 1 percent in the long run (4 years after the shock). Unfortunately, it is not possible to get an accurate measure of real wage from CEX. Because CEX does not ask directly for real wage per hour, I compute its proxy RW from a formula $RW = S/(P \times TH)$, where S is total annual salaries received, P is CPI, and TH is total annual hours worked. I find that correlation between growth rates of RW and its BLS counterpart (real hourly compensation in the business sector) is -0.04, and the standard errors of the impulse responses of RW for each group of households are quite large.

2.4. Splits based on holdings of other assets

I now estimate the responses to a technology shock when households are classified based on the holding statuses of assets other than stocks. As in the baseline analysis, I define households as holding a certain type of asset if they were holding that asset a year before at each point in time. I classify households with positive responses to "Checking accounts, brokerage accounts and other similar accounts", "Savings accounts at banks, savings and loans, credit unions, etc.", and "U.S. savings bonds" as holding checking accounts, savings accounts, and bonds, respectively. According to this classification, 73%, 56%, and 10% of all households hold checking accounts, saving accounts, and bonds, respectively. Fig. 5 shows that, as opposed to households who hold stocks, households who hold other assets (but may or may not hold stocks) do not increase their hours worked when technology improves.

To understand why households holding stocks increase their hours but households holding other assets do not, note that in standard RBC models households increase their hours in order to reap the benefit of the higher return on capital. Fig. 6 shows that the annual stock return, measured using S&P 500 index, indeed increases when technology improves.[14] In

---

[13] The confidence intervals are calculated based on 500 Monte Carlo draws. Specifically, I repeatedly draw coefficients from a multivariate-normal distribution centered around point-estimates of the regression coefficients with a covariance matrix based on the Newey and West (1987) estimator, and compute the implied responses to a one-standard-deviation improvement in technology.

[14] It is important to point out that the increase in the stock return in response to a positive technology shock does not necessarily imply that the market is inefficient. Indeed, there is an extensive literature documenting the predictability of stock returns. Fama and French (1988), for example, find that 25–40 percent of the variation of 3–5 year stock returns is predictable and note that the predictability can result from time-varying equilibrium expected returns in an efficient market. See also Balvers et al. (1990), who formally construct a general equilibrium model to show that some components of stock returns can be predictable to the extent that there is predictability in aggregate output due to persistence.

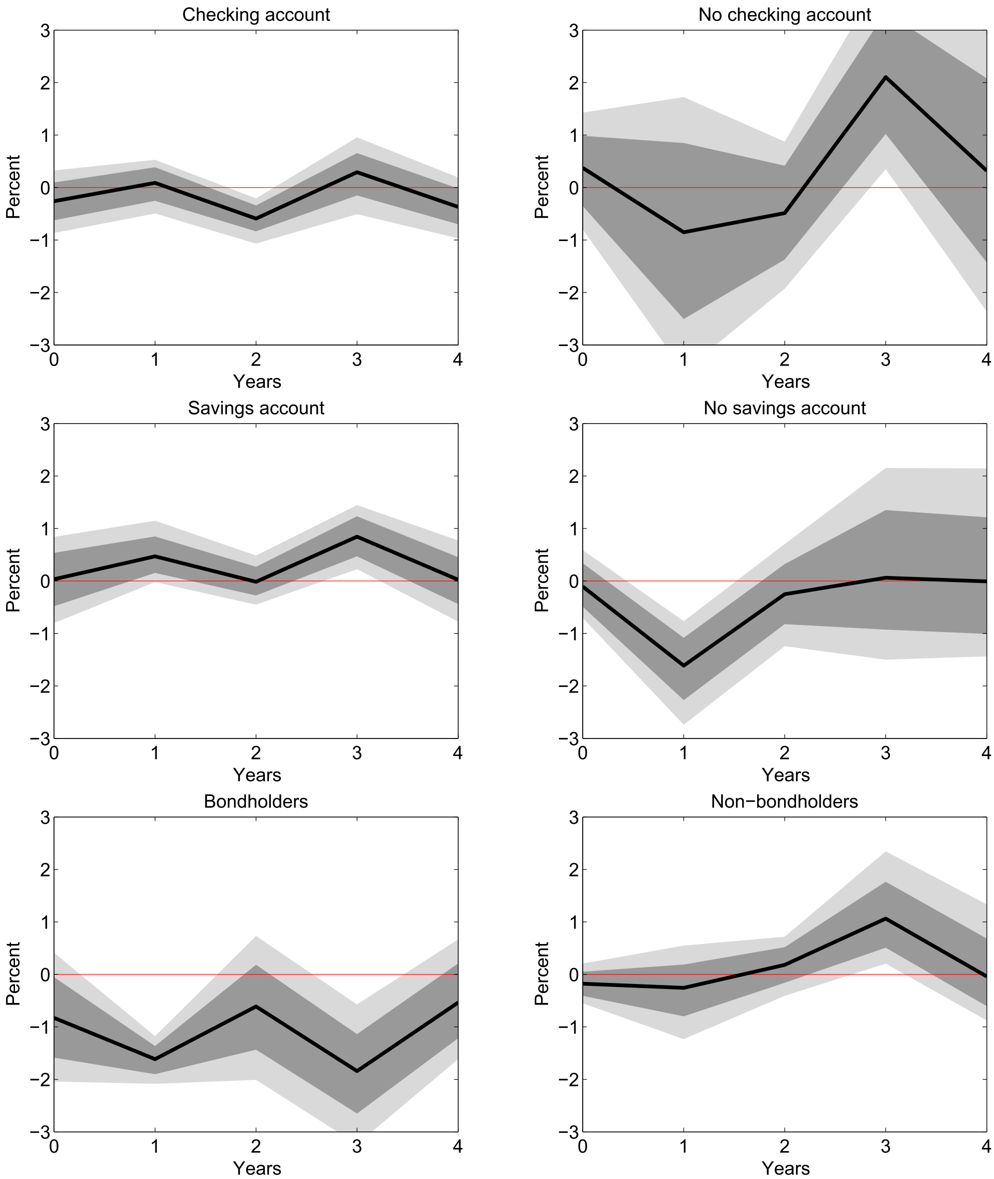

Fig. 5. Splits based on holdings of other assets. *Notes*: Impulse responses of the levels of hours to a one-standard-deviation positive technology shock, estimated from a regression on its own lags and current and lagged technology growth. Technology growth is measured using the utilization-adjusted TFP series constructed in Fernald (2014). The dark and light shaded areas are 68 and 90-percent confidence intervals, respectively. The sample period is 1981–2011.

contrast, returns on treasury bills or bonds do not increase after a positive technology shock (Fig. 7). The evidence supports the hypothesis that stockholders increase their hours in order to increase their investment in stocks, whose returns are higher after an improvement in technology. Households holding other assets do not increase hours because the returns on those assets do not increase when technology improves.[15]

One potential issue with my empirical finding is that not stock holdings per se, but the high amount of wealth may be driving the positive labor supply response to technology shocks. To address this concern, I split households based on the amount of all assets held. To ease comparison with the classification based on households' stock holding statuses, I set the cutoff so that the share of households that are classified to have a large amount of assets is equal to the share of stockholders (14% of the sample). Fig. 8 shows that households with high amount of wealth do not increase their hours in response to a positive technology shock. This result further provides support to the theory that holdings of stocks and not those of other assets are driving the positive labor supply response to a technology shock.

[15] The asset returns data is constructed by Aswath Damodaran (http://www.damodaran.com).

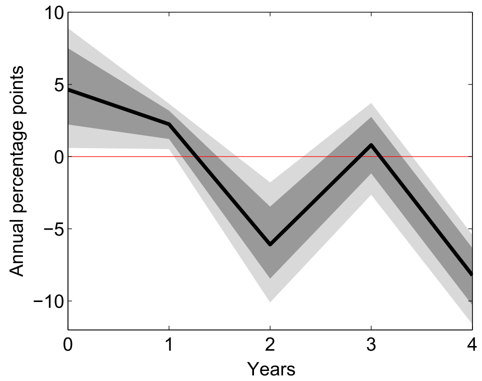

Fig. 6. Stock return after a positive technology shock. *Notes*: Impulse responses of the real annual return on S&P 500 index to a one-standard-deviation positive technology shock, estimated from a regression on its own lags and current and lagged technology growth. Technology growth is measured using the utilization-adjusted TFP series constructed in Fernald (2014). The dark and light shaded areas are 68 and 90-percent confidence intervals, respectively. The sample period is 1981–2011.

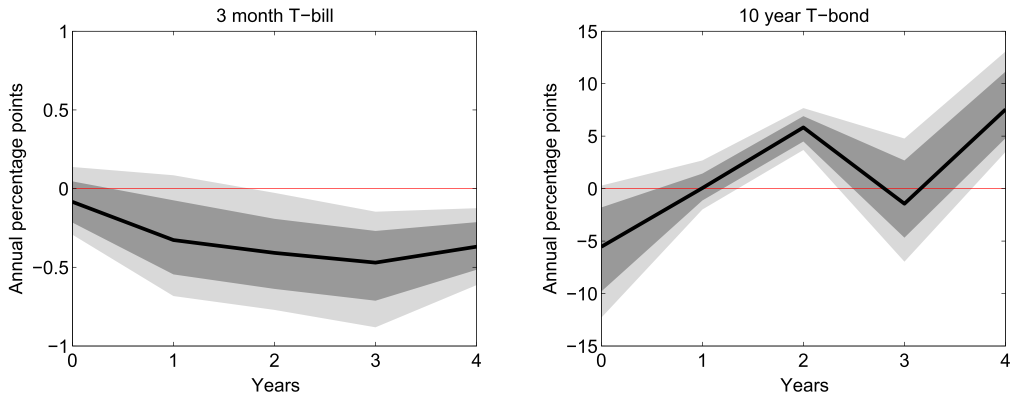

Fig. 7. Return on government bonds after a positive technology shock. *Notes*: Impulse responses of the real annual return on a 3 month treasury bill and a 10 year treasury bond to a one-standard-deviation positive technology shock, estimated from a regression on its own lags and current and lagged technology growth. Technology growth is measured using the utilization-adjusted TFP series constructed in Fernald (2014). The dark and light shaded areas are 68 and 90-percent confidence intervals, respectively. The sample period is 1981–2011.

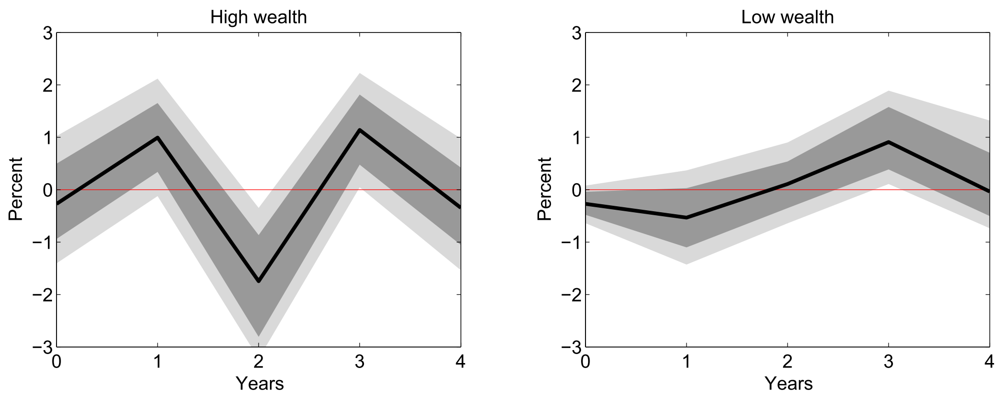

Fig. 8. Split based on the amount of wealth. *Notes*: Impulse responses of the levels of hours to a one-standard-deviation positive technology shock, estimated from a regression on its own lags and current and lagged technology growth. Technology growth is measured using the utilization-adjusted TFP series constructed in Fernald (2014). The cutoff, $12,050 in 1980 dollars (CPI for all urban consumers), is set so that the share of households that are classified to have a large amount of assets is equal to the share of stockholders (14% of the sample). The low wealth households are households whose total amount of assets is below the cutoff, including those with zero and negative wealth. The dark and light shaded areas are 68 and 90-percent confidence intervals, respectively. The sample period is 1981–2011.

To summarize, I find that, when technology improves, stockholders increase their hours but non-stockholders reduce them. Because most households are non-stockholders, aggregate hours fall. The heterogeneity does not arise among households who hold other assets and those who do not.[16] The evidence thus indicates that stockholders raise their hours in order to take advantage of the higher stock returns.

3. A structural analysis

3.1. The model

What are the theoretical and quantitative implications of the micro evidence I found in the previous section? To answer this question, I study a two-agent DSGE model with limited stock market participation based on Galí et al. (2007) and Bilbiie (2008). A fraction $1-\chi$ of households participate in the stock market and the remaining χ fraction of households do not. When $\chi = 0$, the model reduces to a standard representative agent model with full stock market participation. I assume that non-stockholders have access to an asset that bears a constant interest rate. This assumption is motivated from the empirical results from Section 2.4, where I show that the reason why stock ownership matters for the labor supply response while ownerships of other assets do not is because the stock return increases after a technology improvement while returns on other assets do not.

In each period, stockholders choose consumption C_t^s, share holdings S_t^s, and bond holdings B_t^s to maximize utility:

$$\max_{\{C_t^s, S_t^s, B_t^s\}} E_0 \sum_{t=0}^{\infty} \beta^t \left[\ln(C_t^s - bC_{t-1}^s) - \frac{(H_t^s)^{1+\eta}}{1+\eta} \right],$$

where H_t^s is hours worked, β is the discount factor, b controls the degree of habit persistence, and $1/\eta$ is the Frisch elasticity of labor supply when $b = 0$.

Unlike Guvenen (2006) and Guvenen (2009) who assumes heterogeneous intertemporal elasticity of substitution (IES), I assume log utility for both stockholders and non-stockholders. There are two reasons for this assumption. First, preference needs to be consistent with balanced growth. Second, while the Cobb–Douglas preference could potentially allow for both heterogeneous IES and balanced growth, it cannot separately specify the IES, the fraction of time allocated to work, and the Frisch elasticity of labor supply (see Guvenen, 2009 for details on this point). This poses difficulty in interpreting the results because it would confound the impact of limited stock market participation with the effect of heterogeneity in labor supply elasticity. Empirical evidence shows that the IES is likely higher for stockholders and non-stockholders (Vissing-Jørgensen, 2002 and Attanasio et al., 2002). Introducing IES heterogeneity would amplify the heterogeneity in labor supply because it would make the difference in the size of intertemporal substitution effect among stockholders and non-stockholders larger. Thus, assuming log utility for both stockholders and non-stockholders would give me a conservative estimate on the impact of limited stock market participation.

The stockholder's budget constraint is

$$P_t C_t^s + P_t S_t^s + B_t^s \leq W_t^s H_t^s + P_t D_t^s + P_t S_{t-1}^s + R_{t-1} B_{t-1}^s + \Pi_t^s + T^s,$$

where P_t is the price level, S_t^s is the share issued by the mutual funds (which I describe below), D_t^s is the dividends paid per share, W_t^s is the nominal wage, R_{t-1} is the gross nominal interest rate on risk-free bonds from period $t-1$ to t. Π_t^s is the combined profit of the intermediate-goods firms and labor unions distributed equally to each household and T^s is a transfer.

The mutual funds are perfectly competitive and use the funds from stockholders to form physical capital ($K_t = (1-\chi)S_t^s$), which is rented out to intermediate-goods firms at price r_t^k. I introduce variable capital utilization u_t so that the effective capital that is rented out is $u_t K_t$ and assume that higher utilization raises the depreciation rate $\delta(u_t)$.[17] The introduction of capital utilization into the model is motivated by the fact that the long-run identification by Galí (1999) or "purified" Solow residual approach by Basu et al. (2006) and Fernald (2014) are especially concerned about controlling for variable utilization. The mutual funds pay out dividends each period: $D_t^s = (r_t^k u_t - \delta(u_t))S_{t-1}^s$.

Non-stockholders choose consumption C_t^n and asset A_t^n, which earns a constant interest rate $\bar{r}$, to maximize utility:

$$\max_{\{C_t^n, A_t^n\}} E_0 \sum_{t=0}^{\infty} \beta^t \left[\ln(C_t^n - bC_{t-1}^n) - \frac{(H_t^n)^{1+\eta}}{1+\eta} \right],$$

where H_t^n is hours worked, subject to the budget constraint:

[16] In the previous version of the paper, I find that the main empirical finding is robust to a more sophisticated classification of households using the Survey of Consumer Finance (SCF), controlling for business cycle conditions, and several other checks. I also show that the heterogeneous impulse response cannot be explained by demographic factors such as education or birth cohorts. Further details are available from the author upon request.

[17] I assume the function form: $\delta(u_t) = \delta_0 + \delta_1(u_t - 1) + \frac{\delta_2}{2}(u_t - 1)^2$, where $\delta_0 > 0, \delta_1 > 0, \delta_2 > 0$.

$$P_t C_t^n + A_t^n \leq W_t^n H_t^n + \bar{r} A_{t-1}^n + \Pi_t^n + T_t^n - \frac{\kappa}{2}\left(\frac{A_t^n}{\chi Y_t}\right)^2,$$

where Π_t^n is the combined profit of the intermediate labor unions and the last term is the bond-holding cost. $\kappa > 0$ is a parameter that controls the size of the bond-holding cost. The cost eliminates non-stationarity otherwise built into models with exogenous interest rates and induces a unique steady state (Schmitt-Grohe and Uribe, 2003). I assume a segmented asset market: stockholders have access to mutual funds share S_t^s and risk-free bonds B_t^s but not the asset A_t^n. I lower the discount factor of the non-stockholders (while holding the stockholders' discount factor constant) so that $\bar{r} = R$ even with the bond-holding cost.[18] Finally, I also adjust non-stockholders' habit persistence parameter (while holding the non-stockholders' habit parameter constant) so that at the steady state, the marginal utilities of stockholders and non-stockholders are equal even with heterogeneous discount factors.[19]

In each period t, the final goods, Y_t, are produced by a perfectly competitive representative firm that combines a continuum of intermediate goods, indexed by $j \in [0, 1]$, with technology

$$Y_t = \left[\int_0^1 Y_{j,t}^{\frac{\theta_p - 1}{\theta_p}} dj\right]^{\frac{\theta_p}{\theta_p - 1}}.$$

$Y_{j,t}$ denotes the time t input of intermediate good j and θ_p controls the price elasticity of demand for each intermediate good. The intermediate-goods sector is monopolistically competitive. In period t, each firm j rents $K_{j,t}$ units of capital stock from the mutual fund sector and buys $H_{j,t}$ units of aggregate labor input from the employment sector to produce intermediate good j using technology

$$Y_{j,t} = z_t (u_{j,t} K_{j,t})^\alpha (\gamma^t H_{j,t})^{1-\alpha}.$$

γ is the rate of labor-augmenting deterministic technological growth and z_t is the level of technology that follows

$$\ln z_t = (1 - \rho_z)\bar{z} + \rho \ln z_{t-1} + \epsilon_t,$$

where ϵ_t is i.i.d. distributed from a normal distribution with mean zero and variance σ^2. Firms face a Calvo-type price-setting friction: In each period t, a firm can re-optimize its intermediate-goods price with probability $(1 - \xi_p)$. Firms that cannot re-optimize index their price according to the steady-state inflation rate, π.

In addition to sticky prices, I introduce sticky wages for two reasons. First, previous research such as Christiano et al. (2005) have shown that wage stickiness is an important ingredient for New Keynesian models. Second, New Keynesian models with limited asset market participation lead to equilibrium indeterminacy under monetary policy that satisfies the Taylor principle when the share of non-asset holders is sufficiently high (Galí et al., 2004). This is problematic for my analysis because the share of non-stockholders is quite high in the data. As shown in Colciago (2011), however, even a mild degree of wage stickiness restores equilibrium determinacy.[20]

Households supply their homogeneous labor to monopolistically competitive labor unions, who differentiate the household's labor service and set wages subject to a Calvo friction: In each period t, a union i can re-optimize its nominal wage $W_{i,t}$ with probability $(1 - \xi_w)$.[21] A union that cannot re-optimize index their wage according to the steady-state wage growth rate, $\gamma\pi$. I assume that stockholders and non-stockholders sell their labor service to their own unions (i.e., each labor union uses the labor service of only stockholders or only non-stockholders). This assumption allows me to generate heterogeneity in labor supply between stockholders and non-stockholders. A union then sells the differentiated labor to perfectly competitive labor packers. In turn, labor packers aggregate labor service H_t according to a technology

$$H_t = \left[\int_0^1 H_{i,t}^{\frac{\theta_w - 1}{\theta_w}} di\right]^{\frac{\theta_w}{\theta_w - 1}},$$

where θ_w controls the price elasticity of demand for each differentiated labor service, and sells the aggregated labor to the intermediate production firms for a nominal price of W_t per unit.

---

18 The non-stockholders' Euler equation is

$$\gamma \lambda_t^n = \beta E_t \lambda_{t+1}^n \frac{\bar{r}}{\pi_{t+1}} - \gamma \lambda_t^n \kappa \left(\frac{A_t^n}{\chi Y_t}\right) \frac{1}{\chi Y_t},$$

where the last term reflects the bond-holding cost.

19 I have also tried a version where stockholders and non-stockholders share the same discount factor and the habit parameter and obtained similar results.

20 In contrast to Colciago (2011), whose labor market structure implies that both asset holders and non-asset holders supply the same amount of labor, I allow for a heterogeneous labor supply between the two groups.

21 Smets and Wouters (2007) use a similar setting.

The central bank follows a Taylor rule with interest-rate smoothing:

$$\frac{R_t}{R} = \left(\frac{R_{t-1}}{R}\right)^{\rho_R} \left\{ \left(\frac{\pi_t}{\pi}\right)^{\phi_\pi} \left(\frac{Y_t}{\bar{Y}}\right)^{\phi_Y} \left(\frac{Y_t}{Y_{t-1}}\right)^{\phi_{\Delta Y}} \right\}^{1-\rho_R},$$

where R is the steady-state level of the nominal interest rate, ρ_R is the persistence of the rule, and ϕ_π, ϕ_Y, $\phi_{\Delta Y}$ are the sizes of the policy responses to the deviations of inflation, output, and output growth from their steady states, respectively.

The transfer payments balance period by period: $(1-\chi)T^s = \chi T^n$. I set the transfers T^s and T^n so that the levels of consumption and hours along the balanced growth path are the same for both stockholders and non-stockholders.[22] The assumption is motivated by my focus on the heterogeneity of impulse responses rather than the steady-state differences. I tried a version of the model where the steady-state levels of consumption and hours are different across stockholders and non-stockholders but the change had little effect on dynamics. Finally, the aggregate resource constraint is $C_t + I_t = Y_t$.

3.2. Impulse-response-matching estimation

The model is quarterly frequency and the simulated data is aggregated to annual frequency to match the empirical impulse responses in the estimation. I divide the structural parameters into two categories. The first set of parameters is fixed throughout the estimation. The discount factor β is 0.99 and the steady-state quarterly depreciation rate, δ_0, is 2.5%. The labor-augmenting technological growth factor γ is set to 1.004 so that the steady-state annual growth rate of output is 1.6%. θ_p and θ_w are both 11, which generate steady-state markups of 10%. I set $1-\chi = 0.14$. The value implies that the stock market participation rate is 14%, which is in line with the CEX evidence described in the previous section. The steady-state inflation rate π is set to be equal to the average inflation rate over the sample period.

The remaining set of parameters is estimated using a Bayesian version of an impulse-response-matching method. The method was developed by Christiano et al. (2010) and was also used in Christiano et al. (2013). The method first finds the "likelihood" of the data motivated by approximation based on asymptotic distribution theory. Let $\hat{\psi}$ denote the impulse response function computed using the data and let $\psi(\gamma)$ denote the impulse response function computed using the simulated data from the DSGE model, both based on the regression specification (1).[23] Suppose the DSGE model is correct and let γ_0 denote the true parameter vector; hence $\psi(\gamma_0)$ is the true impulse response function. Then we know that

$$\sqrt{T}(\hat{\psi} - \psi(\gamma_0)) \xrightarrow{d} N(0, W(\gamma_0)),$$

where T is the number of observations and $W(\gamma_0)$ is the asymptotic sampling variance, which depends on γ_0. The asymptotic distribution of $\hat{\psi}$ can be rewritten as

$$\hat{\psi} \xrightarrow{d} N(\psi(\gamma_0), V), \qquad V \equiv \frac{W(\gamma_0)}{T}.$$

When implementing the estimation algorithm, I use a consistent estimator of V. Specifically, V is a diagonal matrix with the sample variance of the $\hat{\psi}$ along the main diagonal. The non-diagonal terms are set to zero due to small sample considerations.

The method then uses the approximation of the likelihood

$$\mathcal{L}(\psi|\gamma) = (2\pi)^{-\frac{N}{2}} |V|^{-\frac{1}{2}} \exp\{-0.5[\hat{\psi} - \psi(\gamma)]' V^{-1} [\hat{\psi} - \psi(\gamma)]\},$$

where N is the number of impulse responses to be matched, in order to obtain the posterior distribution of γ, $p(\gamma|\psi)$, using the Bayes law:

$$p(\gamma|\psi) = \frac{p(\gamma)\mathcal{L}(\psi|\gamma)}{p(\psi)},$$

where $p(\gamma)$ is the prior and $p(\psi)$ is the marginal likelihood. I numerically characterize the posterior distribution using the random-walk Metropolis–Hastings algorithm.

Table 2 reports the prior distributions. Most priors are centered around standard values found in the literature. I include impulse responses of output, consumption, investment, aggregate hours, GDP deflator inflation, and Federal funds rate to the technology shock for the estimation.[24] These macro impulse responses are obtained by running separate regressions (1) for each variable on current and lagged technology growth, measured using the utilization-adjusted TFP series (Fernald, 2014), and on their own lag. The resulting impulse responses are displayed in Fig. 9. In response to a technology improvement, output, investment, consumption move little initially and increase over time. Nominal variables such as inflation and the

[22] Galí et al. (2007) also make this assumption.

[23] Thus, the data and the theoretical model are treated symmetrically in the estimation (Kehoe, 2006).

[24] Output is defined as real GDP, consumption is defined as (expenditures on nondurables)+(expenditures on services), and investment is defined as (expenditures on durables)+(private fixed investment). For aggregate hours, I use total annual hours worked for all households constructed from the CEX.

Table 2
Estimated parameters.

	Description	Prior			Posterior mode		
		Type	Mean	Std.	Baseline	χ estimated	Rep. agent
α	Capital share	B	0.35	0.03	0.35 (0.03)	0.35 (0.03)	0.34 (0.03)
δ_2	Convexity of depreciation	G	0.50	0.20	0.42 (0.24)	0.42 (0.20)	0.41 (0.21)
η	Inverse Frisch elasticity	G	0.60	0.30	0.15 (0.18)	0.29 (0.23)	0.42 (0.34)
b	Consumption habit	B	0.50	0.20	0.95 (0.07)	0.98 (0.21)	0.54 (0.20)
κ	Bond holding cost	G	1.00	0.30	1.09 (0.23)	0.98 (0.26)	–
χ	Non-stockholders share	B	0.50	0.10	0.86	0.59 (0.09)	0
ξ_p	Calvo price	B	0.50	0.20	**0.28** (0.10)	**0.31** (0.16)	**0.76** (0.15)
ξ_w	Calvo wage	B	0.50	0.20	**0.17** (0.10)	**0.32** (0.17)	**0.50** (0.18)
ρ_R	Interest smoothing	B	0.60	0.10	0.60 (0.09)	0.61 (0.09)	0.66 (0.09)
ϕ_π	Inflation response	N	1.70	0.30	1.73 (0.14)	1.69 (0.15)	1.40 (0.20)
ϕ_Y	Output response	N	0.15	0.05	0.13 (0.04)	0.16 (0.04)	0.17 (0.04)
$\phi_{\Delta Y}$	Output growth response	N	0.15	0.05	0.15 (0.05)	0.15 (0.05)	0.16 (0.05)
ρ_z	Technology shock	N	0.60	0.20	0.99 (0.02)	0.97 (0.18)	0.89 (0.10)
100σ	Technology shock	IG	1.00	1.00	0.28 (0.04)	0.25 (0.06)	0.28 (0.06)

Notes: *B* refers to the Beta distribution, *N* to the Normal distribution, *G* to the Gamma distribution, and *IG* to the Inverse-gamma distribution. Posterior standard deviations are in parentheses.

Federal funds rate fall. These patterns are broadly in line with previous studies such as Basu et al. (2006). In addition to those standard macro variables, when I estimate the baseline model with limited stock market participation, I also use stockholders and non-stockholders' hours responses from the CEX, again estimated using the regression (1).

3.3. Results

Table 2 gives the estimated parameter values for the baseline model (limited stock market participation) and the counterfactual model (full stock market participation). First, consider the posterior estimates of the counterfactual model (labeled "Rep. agent"). The estimated Calvo parameters for price and wage re-optimizations are 0.76 and 0.50, respectively. To put those values into perspective, the estimates imply that prices and wages are adjusted on average $(1/(1-0.76) \approx)$ 4 and $(1/(1-0.50) \approx)$ 2 quarters, respectively.[25] Due to strong nominal rigidities, the model can replicate the decline in aggregate hours and hence also the mild increases in quantities such as output and consumption. (For the interest of space, these impulse responses are not shown.)

Next, consider the posterior of the baseline model (labeled "Baseline" in Table 2). The estimates of Calvo price and wage parameters are 0.28 and 0.17, respectively and are significantly smaller than the estimates from the full stock market participation model. The values imply that firms and households adjust their prices and wages roughly every quarter.[26]

The blue solid lines with circles in Fig. 9 show that the model is able to reproduce the heterogeneous response to a technology shock. In the estimated model, a positive technology shock raises return on investment. Stockholders increase their labor supply due to the standard intertemporal substitution effect: They work more in order to reap the benefit of higher returns. Non-stockholders reduce their hours due to the income effect: Due to consumption habit, these households prefer to increase their consumption slowly to the new steady-state level. This can be achieved by less hours thanks to

[25] In the DSGE estimation literature, wage stickiness is typically estimated to be larger than price stickiness. For example, in Smets and Wouters (2007), the estimated frequencies of price and wage adjustments are 3 and 4 quarters, respectively. The reason why I find the opposite result is likely due to the estimation method used. For example, in the literature models featuring multiple structural shocks are estimated using full information methods while in this paper I use the impulse response to a technology shock to estimate parameters using a limited information procedure. In the context of my model, I find that increasing wage stickiness tends to lower the labor supply in the medium run (2–4 years after the shock). This, in turn, worsens the fit of the model to the data since the empirical response of aggregate hours, identified using Fernald (2014) series, returns to the pre-shock level 2 years after impact.

[26] The estimated Frisch elasticity, $1/\eta$, is higher than the usual estimates. To explore this further, I have re-estimated the model by fixing η to 0.5, which implies a Frisch elasticity of 2, and found the main results unaffected. Further details are available upon request.

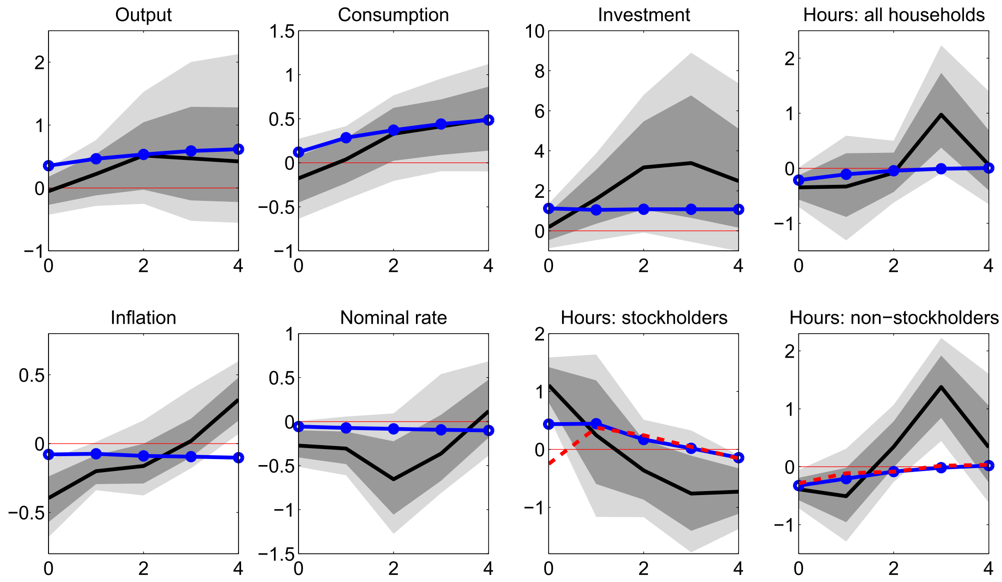

Fig. 9. Data and model impulse responses. *Notes*: Impulse responses from the data (black solid lines) and the baseline model with limited stock market participation (blue solid lines with circles), estimated from a regression on its own lags and current and lagged technology growth. The red dashed lines are the counterfactual impulse responses where I set the Calvo price and wage parameters to the estimated values from the full participation model while other parameters are fixed at the estimated values from the limited participation model. The units are in percentage deviations from the steady states except for the inflation and the nominal interest rate, which are in quarterly percentage-point deviations.

higher wages. The estimation prefers low nominal rigidities because the labor supply responses are heterogeneous in the data. Indeed, when I increase the Calvo price and wage parameters to the estimated values from the full participation model while holding other parameters fixed, I find that hours of both stockholders and non-stockholders decrease in response to a positive technology shock (red dashed lines in Fig. 9). In this case, stockholders reduce their hours because a positive technology shock reduces the return on investment due to countercyclical markups. The Figure also shows that, in addition to the household-level hours, the heterogeneous-agent model is able to replicate the responses of other standard macro variables.

A possible concern is to what extent my results regarding nominal rigidities are sensitive to the calibrated share of non-stockholders ($\chi = 0.86$), which is higher than what is usually adopted in the literature.[27] To explore this issue, I estimate a version of the baseline model where I also estimate the share of non-stockholders χ. The prior for χ is centered around 0.5 with a relatively tight standard deviation of 0.1. The resulting parameter estimates are found under the column "χ estimated" in Table 2. First, the estimated share of non-stockholders is lower at $\chi = 0.59$. Second, the estimated Calvo price and wage parameters are $\xi_p = 0.31$ and $\xi_w = 0.32$, which are slightly higher than the baseline estimation but are still substantially lower than the full stock market participation model. The estimated Calvo parameters are slightly higher than the baseline estimation because the estimation tries to reconcile the higher share of stockholders and the negative empirical response of aggregate hours by lowering the labor supply response of stockholders. Overall, the exercise indicates that the main conclusion — the estimated nominal rigidities are lower when household-level impulse responses are taken into account — is robust to a lower degree of non-participation rate.

4. Concluding remarks

In this paper, I exploited heterogeneous impulse responses at the household level due to limited stock market participation to provide novel evidence on the degree of nominal rigidities. Using micro data from the Consumer Expenditure Survey, I have shown that, in response to a positive technology shock, stockholders increase their hours but non-stockholders reduce them. The heterogeneity does not arise among households who hold other assets and those who do not. Instead, this heterogeneity arises because the strength of the intertemporal substitution effect on labor supply varies across the two groups. In the aggregate, hours fall because most households are non-stockholders. A parsimonious two-agent DSGE model with limited stock market participation is able to replicate the heterogeneity in the household-level impulse responses and generates an estimate regarding the degree of nominal rigidities that is smaller than is found under the representative-agent assumption.

I point out two directions for future research. First, the stark difference between the empirical impulse responses of stockholders and non-stockholders implies that technology shocks may have substantial distributional consequences, as

[27] For example, Galí et al. (2007) use $\chi = 0.5$.

found in Coibion et al. (2017) for monetary policy shocks. How technology shocks affect inequality and welfare is a crucial research question. Second, empirical and theoretical results in this paper suggest that limited stock market participation has important implications on transmissions of other aggregate shocks as well. A thorough investigation of this possibility is needed.

References

Alexopoulos, Michelle, 2011. Read all about it!! What happens following a technology shock? The American Economic Review 101, 1144–1179.
Altig, David, Christiano, Lawrence J., Eichenbaum, Martin, Lindé, Jesper, 2011. Firm-specific capital, nominal rigidities and the business cycle. Review of Economic Dynamics 14, 225–247.
Anderson, Emily, Inoue, Atsushi, Rossi, Barbara, 2016. Heterogeneous consumers and fiscal policy shocks. Journal of Money, Credit, and Banking 48, 1877–1888.
Attanasio, Orazio P., Banks, James, Tanner, Sarah, 2002. Asset holding and consumption volatility. Journal of Political Economy 110, 771–792.
Balvers, Ronald J., Cosimano, Thomas F., Mcdonald, Bill, 1990. Predicting stock returns in an efficient market. The Journal of Finance 45, 1109–1128.
Basu, Susanto, Fernald, John G., Kimball, Miles S., 2006. Are technology improvements contractionary? The American Economic Review 96, 1418–1448.
Bilbiie, Florian, 2008. Limited asset markets participation, monetary policy and (inverted) aggregate demand logic. Journal of Economic Theory 140, 162–196.
Bocola, Luigi, Hagedorn, Marcus, Manovskii, Iourii, 2014. Identifying neutral technology shocks. Manuscript.
Broer, Tobias, Harbo Hansen, Niels-Jakob, Krusell, Per, Oberg, Erik, 2016. The new Keynesian transmission channel: a heterogeneous-agent perspective. Manuscript.
Canova, Fabio, Lopez-Salido, David, Michelacci, Claudio, 2010. The effects of technology shocks on hours and output: a robustness analysis. Journal of Applied Econometrics 25, 755–773.
Chang, Yongsung, Hong, Jay H., 2006. Do technological improvements in the manufacturing sector raise or lower employment? The American Economic Review 96, 352–368.
Chari, V.V., Kehoe, Patrick J., McGrattan, Ellen R., 2008. Are structural VARs with long-run restrictions useful in developing business cycle theory? Journal of Monetary Economics 55, 1337–1352.
Christiano, Lawrence J., Eichenbaum, Martin, Vigfusson, Robert, 2003. What Happens After a Technology Shock? NBER Working Paper.
Christiano, Lawrence J., Eichenbaum, Martin, Evans, Charles L., 2005. Nominal rigidities and the dynamic effects of a shock to monetary policy. Journal of Political Economy 113, 1–45.
Christiano, Lawrence J., Trabandt, Mathias, Walentin, Karel, 2010. Involuntary Unemployment and the Business Cycle. NBER Working Paper.
Christiano, Lawrence J., Eichenbaum, Martin, Trabandt, Mathias, 2013. Unemployment and business cycles. Econometrica 84, 1523–1569.
Cloyne, James, Surico, Paolo, 2017. Household debt and the dynamic effects of income tax changes. The Review of Economic Studies 84, 45–81.
Cloyne, James, Ferreira, Clodomiro, Surico, Paolo, 2015. Monetary policy when households have debt: new evidence on the transmission mechanism. Manuscript.
Coibion, Olivier, Gorodnichenko, Yuriy, 2012. What can survey forecasts tell us about information rigidities? Journal of Political Economy 120, 116–159.
Coibion, Olivier, Gorodnichenko, Yuriy, Kueng, Lorenz, Silvia, John, 2017. Innocent bystanders? Monetary policy and inequality in the U.S. Journal of Monetary Economics 88, 70–89.
Colciago, Andrea, 2011. Rule-of-thumb consumers meet sticky wages. Journal of Money, Credit, and Banking 43, 325–353.
De Giorgi, Giacomo, Gambetti, Luca, 2017. Business cycle fluctuations and the distribution of consumption. Review of Economic Dynamics 23, 19–41.
Dupor, Bill, Han, Jing, Tsai, Yi-Chan, 2009. What do technology shocks tell us about the new Keynesian paradigm? Journal of Monetary Economics 56, 560–569.
Erceg, Christopher J., Guerrieri, Luca, Gust, Christopher, 2005. Can long-run restrictions identify technology shocks? Journal of the European Economic Association 3, 1237–1278.
Fama, Eugene F., French, Kenneth R., 1988. Permanent and temporary components of stock prices. Journal of Political Economy 96, 246–273.
Fernald, John G., 2007. Trend breaks, long-run restrictions, and contractionary technology improvements. Journal of Monetary Economics 54, 2467–2485.
Fernald, John G., 2014. A quarterly, utilization-adjusted series on total factor productivity. Manuscript.
Francis, Neville, Ramey, Valerie A., 2002. Is the technology-driven real business cycle hypothesis dead? Shocks and aggregate fluctuations revisited. Journal of Monetary Economics 110, 1379–1399.
Furlanetto, Francesco, Seneca, Martin, 2012. Rule-of-thumb consumers, productivity, and hours. Scandinavian Journal of Economics 114, 658–679.
Galí, Jordi, 1999. Technology, employment, and the business cycle: do technology shocks explain aggregate fluctuations? The American Economic Review 89, 249–271.
Galí, Jordi, López-Salido, David J., Vallés, Javier, 2004. Rule-of-thumb consumers and the design of interest rate rules. Journal of Money, Credit, and Banking 36, 739–763.
Galí, Jordi, López-Salido, David J., Vallés, Javier, 2007. Understanding the effects of government spending on consumption. Journal of the European Economic Association 5, 227–270.
Gorodnichenko, Yuriy, Weber, Michael, 2016. Are sticky prices costly? Evidence from the stock market. The American Economic Review 106, 165–199.
Guvenen, Fatih, 2006. Reconciling conflicting evidence on the elasticity of intertemporal substitution: a macroeconomic perspective. Journal of Monetary Economics 53, 1451–1472.
Guvenen, Fatih, 2009. A parsimonious macroeconomic model for asset pricing. Econometrica 77, 1711–1740.
Investment Company Institute, 2008. Equity and bond ownership in America, 2008. Available at http://www.ici.org/research/investors.
Kehoe, Patrick J., 2006. How to Advance Theory with Structural VARs: Use the Sims–Cogley–Nason Approach. NBER Working Paper.
Lindé, Jesper, 2009. The effects of permanent technology shocks on hours: can the RBC-model fit the VAR evidence? Journal of Economic Dynamics and Control 33, 597–613.
Liu, Zheng, Phaneuf, Louis, 2007. Technology shocks and labor market dynamics: some evidence and theory. Journal of Monetary Economics 54, 2534–2553.
Nakamura, Emi, Steinsson, Jon, 2014. Fiscal stimulus in a monetary union: evidence from US regions. The American Economic Review 104, 753–792.
Newey, Whitney, West, Kenneth D., 1987. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. Econometrica 55, 703–708.
Romer, Christina D., Romer, David H., 2004. A new measure of monetary shocks: derivation and implications. The American Economic Review 94, 1055–1084.
Rotemberg, Julio J., 2003. Stochastic technical progress, smooth trends, and nearly distinct business cycles. The American Economic Review 93, 1543–1559.
Saijo, Hikaru, 2018. Redistribution and fiscal uncertainty shocks. Manuscript.
Schmitt-Grohe, Stephanie, Uribe, Martin, 2003. Closing small open economy models. Journal of International Economics 61, 163–185.
Shea, John, 1999. What do technology shocks do? In: NBER Macroeconomics Annual 1998, vol. 13. MIT Press, pp. 275–322.
Smets, Frank, Wouters, Raf, 2007. Shocks and frictions in US business cycles: a Bayesian DSGE approach. The American Economic Review 97, 586–606.
Uhlig, Harald, 2004. Do technology shocks lead to a fall in total hours worked? Journal of the European Economic Association 2, 361–371.

Vissing-Jørgensen, Annette, 2002. Limited asset market participation and the elasticity of intertemporal substitution. Journal of Political Economy 110, 825–853.
Walsh, Carl E., 2016. Workers, capitalists, wages, employment and welfare. Manuscript.