

Technology Shocks and Hours Revisited: Evidence from Household Data

Hikaru Saijo*

University of California, Santa Cruz

June 11, 2017

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.

*I am 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.

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.¹ 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.

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,² conclusions based on aggregate hours could thus be misleading. Indeed, I show in a stylized model that, 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, I also show that 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 the technology shock series constructed in Fernald (2014). The series controls

¹For 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).

²For example, in the Consumer Expenditure Survey dataset I use, about 14% of all households hold stocks.

for heterogeneity across different types of inputs and variations in factor utilization such as labor effort and the workweek of capital. 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 a more sophisticated classification of households using the Survey of Consumer Finance (SCF), an alternative identification of technology shocks by Galí (1999), controlling for business cycle conditions, and several other checks.

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 or income do not generate the heterogeneous impulse responses I find from the classification based on the households' stock-holding status. The evidence thus supports the transmission mechanism in the stylized model; stockholders increase their hours in order to reap the benefit of the higher return on investment.

An additional concern of the main finding is that, since non-stockholders tend to be less educated and younger than stockholders, the heterogeneous impulse responses may reflect factors such as skilled vs. unskilled labor or life-cycle effects. To address this issue, I explore classifications based on education and birth cohorts. I show that the difference in impulse responses among households whose heads have college degrees and those who do not is much less pronounced than when households are split based on their stock holding status. In addition, I show that both stockholders with college degrees and stockholders without raise hours by a similar amount in response to a technology improvement. This evidence shows that the heterogeneous impulse response I find for stockholders and non-stockholders is not driven by the fact that stockholders tend to be more educated. I also find that, according to the classification based on birth cohorts, no household type increases the labor supply in response to a positive technology shock, indicating that life-cycle effects cannot explain my main result.

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). 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, I use a stylized model with limited stock market participation to illustrate the key qualitative predictions regarding the hours response to technology shocks. In Section 3, 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 4, I explore the quantitative implications of my empirical findings by estimating DSGE models with limited stock market participation. Finally, Section 5 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 (forthcoming) and Cloyne et al. (2015), who exploit household heterogeneity and study the role of liquidity constraints, proxied by the housing tenure status, on 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. (2016), who explore the effects of monetary policy shocks on consumption and income inequality, Anderson et al. (forthcoming), 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).³ 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.⁴ 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.⁵ 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 A stylized example

In this section, I use a simple two-periods-two-agents model to illustrate the key qualitative effects of limited stock market participation on household-level and aggregate employment. The intuition I develop in this section helps motivate the empirical analysis and also provides a basis for the DSGE model I develop later.

There is a continuum of households of measure one. A fraction $1 - \chi$ of households are stockholders and the remaining χ fraction are non-stockholders. The stockholders live for two periods and their

³See 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.

⁴See, 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.

⁵Recent works that use a similar framework include Broer et al. (2016) and Walsh (2016). 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.

lifetime utility is given by

$$U = \left[\frac{(C_1^s)^{1-\gamma}}{1-\gamma} - \frac{(H_1^s)^{1+\eta}}{1+\eta} \right] + C_2^s,$$

where $1/\gamma$ is the intertemporal elasticity of substitution and $1/\eta$ is the Frisch elasticity of labor supply. At $t = 1$, stockholders work (H_1^s) , consume (C_1^s) , and invest (K_2) for capital used in the second period. Their budget constraint is given by

$$C_1^s + K_2 \leq wH_1^s,$$

where w is the real wage. They are also subject to the cash-in-advance (CIA) constraint:

$$P_1(C_1^s + K_2) \leq M_1^s,$$

where P_1 is the price level and M_1^s is the nominal money holdings. At $t = 2$, they simply consume the capital rental income:

$$C_2^s \leq rK_2.$$

I assume that the second-period consumption C_2^s enters the utility linearly. This assumption allows me to explicitly characterize the solution.⁶

Non-stockholders live for only one period. This assumption allows me to completely turn off the intertemporal substitution effect on labor supply for non-stockholders. They maximize utility

$$U = \frac{(C_1^n)^{1-\gamma}}{1-\gamma} - \frac{(H_1^n)^{1+\eta}}{1+\eta},$$

subject to the budget constraint

$$C_1^n \leq wH_1^n,$$

and the CIA constraint

$$P_1 C_1^n \leq M_1^n,$$

where M_1^n is the nominal money holdings.

At $t = 1$, perfectly competitive firms produce according to the production function

$$Y_1 = z_1 H_1,$$

where z_1 is the technology level at $t = 1$, and H_1 is the aggregate hours worked: $H_1 = (1-\chi)H_1^s + \chi H_1^n$. At $t = 2$, firms produce according to

$$Y_2 = z_2 K_2,$$

where z_2 is the technology level.

I consider a surprise permanent increase in technology at $t = 1$; $z_1 = z_2 = z'$ and $z' > \bar{z}$, where \bar{z}

⁶The assumption also implies that investment unambiguously increases when the return increases.

is the steady-state level of technology. I study the household-level and aggregate hours (H_1^s, H_1^n , and H_1) at $t = 1$.

First, consider the case where prices are perfectly flexible. Non-stockholders' equilibrium labor supply is given by

$$\hat{H}_1^n = \frac{1 - \gamma}{\eta + \gamma} \hat{z}_1,$$

where variables with hats denote log-deviations from the steady state. When $1/\gamma < 1$, that is, when intertemporal elasticity of substitution (IES) is below one, the income effect on wage dominates the static substitution effect on wage, and hence non-stockholders reduce their labor supply. Stockholders' labor supply is given by

$$\hat{H}_1^s = \frac{2}{\eta} \hat{z}_1 > 0.$$

Stockholders always increase their labor supply because they want to take advantage of the increase in the return on investment. Aggregate hours worked decline if (i) $1/\gamma < 1$ and (ii) the share of non-stockholders is sufficiently high.

Next, consider the case where prices are sticky. I assume that price and money holdings are chosen before the realization of the technology shock. The assumption implies that, through the CIA constraints, the goods demand in the first period is fixed at the steady-state level. Because firms can produce the same amount of goods with less labor thanks to higher productivity, labor demand declines. In turn, households can buy the same amount of goods with less hours worked because the real wage is higher. Thus, equilibrium hours of both stockholders and non-stockholders decline:

$$\hat{H}_1^s = \hat{H}_1^n = \hat{H}_1 = \hat{z}_1 < 0.$$

The findings of the stylized model can be summarized as follows. With limited stock market participation, aggregate hours could decline in response to a positive technology shock in both flexible-price and sticky-price models. In contrast to aggregate data, household-level data are useful in discriminating among alternative models. In flexible-price models, stockholders always increase their labor supply while non-stockholders reduce their labor supply as long as the income effect dominates the static substitution effect. In sticky-price models, both stockholders and non-stockholders reduce their labor supply.

3 Empirical evidence

3.1 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 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 psuedo-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”. As illustrated in the stylized model presented above, intertemporal substitution made possible through capital accumulation is an important channel through which technology shocks affect labor supply in standard business cycle models. 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. Nevertheless, as a robustness check I also consider a more sophisticated definition of stockholders by supplementing the CEX data with probit analysis using household data from the Survey of Consumer Finance (SCF).

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

⁷In 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.

⁸For 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.

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 $(3600/4195 \approx 86\%)$ 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 $(749/1031 \approx 73\%)$ and households whose heads belong to the older birth cohort $(885/1044 \approx 85\%)$ are non-stockholders. I later examine the labor supply responses of households when they are classified based on their educations and birth cohorts.

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. Note that because 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. I define annual total hours worked as a product of (a) hours worked per week, averaged over 4 quarters, and (b) number of weeks worked in the last 12 months, which I use only the information collected in the fourth quarter.¹¹ A potential way to aggregate annual total hours worked would be

⁹As discussed in Vissing-Jørgensen (2002), the percentage of stockholders are 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.

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

¹¹Households are interviewed either in October, November, or December and it is possible to use only information collected in December to measure annual number of weeks worked. However, practical considerations regarding the sample size forces me to use information collected in all months during the fourth quarter.

to take the product of (a) and (b) for each household first and then take the cross-sectional average. However, because the entrance and exit of households to the survey happens every month, we need to exclude a large number of households who did not participate in the second through fifth interviews for all quarters in a calendar year to use this method.¹² In order to avoid the small sample problem, I thus take the cross-sectional averages of (a) and (b) separately across households¹³ and then take the product of the household averages of (a) and (b).

To further 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 I obtain the time series of stockholders' annual hours worked.

Annual hours for all households tracks the business sector hours series provided by the Bureau of Labor Statistics 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.

3.2 Measuring the household-level responses to technology shocks

To measure changes in aggregate technology, I 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. I later check the robustness of the main result by using the technology shock series identified from a long-run restriction from a structural VAR (SVAR) as in Galí (1999). The sample period of the analysis is 1981–2011.¹⁴

To measure 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

¹²For example, if a household participated in the second through fifth interviews 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).

¹³As recommended by the BLS, I use population weights.

¹⁴I lose the first year of the CEX data because I use growth rates.

Figure 1: Aggregate hours decline in response to a positive technology shock



Notes: Impulse responses of the levels of aggregate 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. See text for details on the construction. The sample period is 1981–2011.


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 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.

Figure 1 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 0.3 percent and gradually recovers to 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).¹⁵ Figure 2 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.

A potential caveat for using the simple definition of stockholders described above is that we may be misclassifying households who only hold bonds (but not stocks) as stockholders. To address this concern, I follow Malloy et al. (2009) and use the SCF to refine the categorization of stockholders with a predicted probability of households' stock ownerships from a probit model. The SCF from 1989,


¹⁵The 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.

Figure 2: 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, 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. See text for details on the construction. The sample period is 1981–2011.

Figure 3: Alternative definition of stockholders using CEX and probit from SCF



Notes: Impulse responses of the levels of stockholders' 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). Households are classified as owning stocks if they satisfy the simple definition of stockholders using CEX and have a predicted probability of stock ownership from the probit model greater than 0.5. The dark and light shaded areas are 68 and 90-percent confidence intervals, respectively. See text for details on the construction. The sample period is 1981–2011.

Figure 4: Technology shocks identified using a long-run restriction



Notes: Impulse responses of the levels of hours to a one-standard-deviation positive technology shock identified using a long-run restriction as in Galí (1999). The dark and light shaded areas are 68 and 90-percent confidence intervals, respectively. See text for details on the construction. The sample period is 1981–2011.


1992, 1995, 1998, 2001, 2004, 2007, and 2010 contains the entire information regarding an individual household’s wealth decomposition, including direct and indirect stock holding statuses. I use the SCF to estimate a probit model for a probability that households own stocks based on observable characteristics that are also asked in the CEX (age, education level, race, year, income, savings and checking accounts, and dividend income). I then use the estimated coefficients to form predicted probabilities of stock ownership in the CEX data. The probit estimate is reported in Appendix A. Under the more sophisticated definition of stockholders, households are classified as owning stocks if they satisfy the simple definition of stockholders *and* have a predicted probability of stock ownership from the probit model greater than 0.5. Figure 3 shows that under this more sophisticated definition of stockholders, the increase in hours is stronger than under the baseline definition. For example, under the alternative definition, according to the point estimate hours increase by 1.7 percent on impact while under the baseline definition it increases by 1.1 percent on impact.

I also estimate household labor supply responses when the technology shock is identified using a long-run restriction as in Galí (1999). Specifically, assuming there are technology shocks ν_t^z and non-technology shocks ν_t^x , the vector of log labor productivity growth Δpr_t and the growth rate of the variable of interest Δx_t can be expressed in a MA form as

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

where I impose $C^{12}(L) = 0$, i.e., only the technology shock has a permanent effect on labor productivity. I estimate the bivariate VAR on one lag using annual data from 1981–2011 where x is either

Figure 5: Impulse responses controlling for business cycle conditions



Notes: Impulse responses of the levels of hours to a one-standard-deviation positive technology shock, estimated from a regression on its own lags, current and lagged technology growth, and current and lagged real GDP 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. See text for details on the construction. The sample period is 1981–2011.

hours worked by stockholders or hours worked by non-stockholders. Figure 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. Interestingly, in contrast to Figure 2, non-stockholders’ hours do not show an overshooting.

To ensure the robustness of the results against a possible correlation of technology shocks and business cycle conditions, I re-run the regression (1) by adding current and lagged real GDP growth to the independent variables:

$$\Delta x_t = a + \sum_{i=1}^I b_i \Delta x_{t-i} + \sum_{j=1}^J c_j \Delta z_{t+1-j} + \sum_{k=1}^K d_k \Delta y_{t+1-k} + e_t,$$

where I set $K = 5$ and Δy is the real GDP growth.¹⁶ Figure 5 shows that the key finding of the heterogeneous impulse responses continues to hold after controlling for business cycle conditions. As in Figure 4, non-stockholders’ hours do not show an overshooting.

I conduct two additional robustness checks. The impulse responses are collected in Appendix B. First, I estimate impulse responses restricting the sample to single households. Figure 18 shows that, while the initial response of non-stockholders is not statistically different from zero, stockholders increase their hours much more than the baseline impulse response in Figure 2. Second, I consider alternative numbers of lags for lagged hours growth in (1). Specifically, I set $I = 1$ and $I = 3$. Figure

¹⁶The result is robust to alternative lag lengths.

19 shows that for both specifications, the impulse responses look very similar to Figure 2.

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.

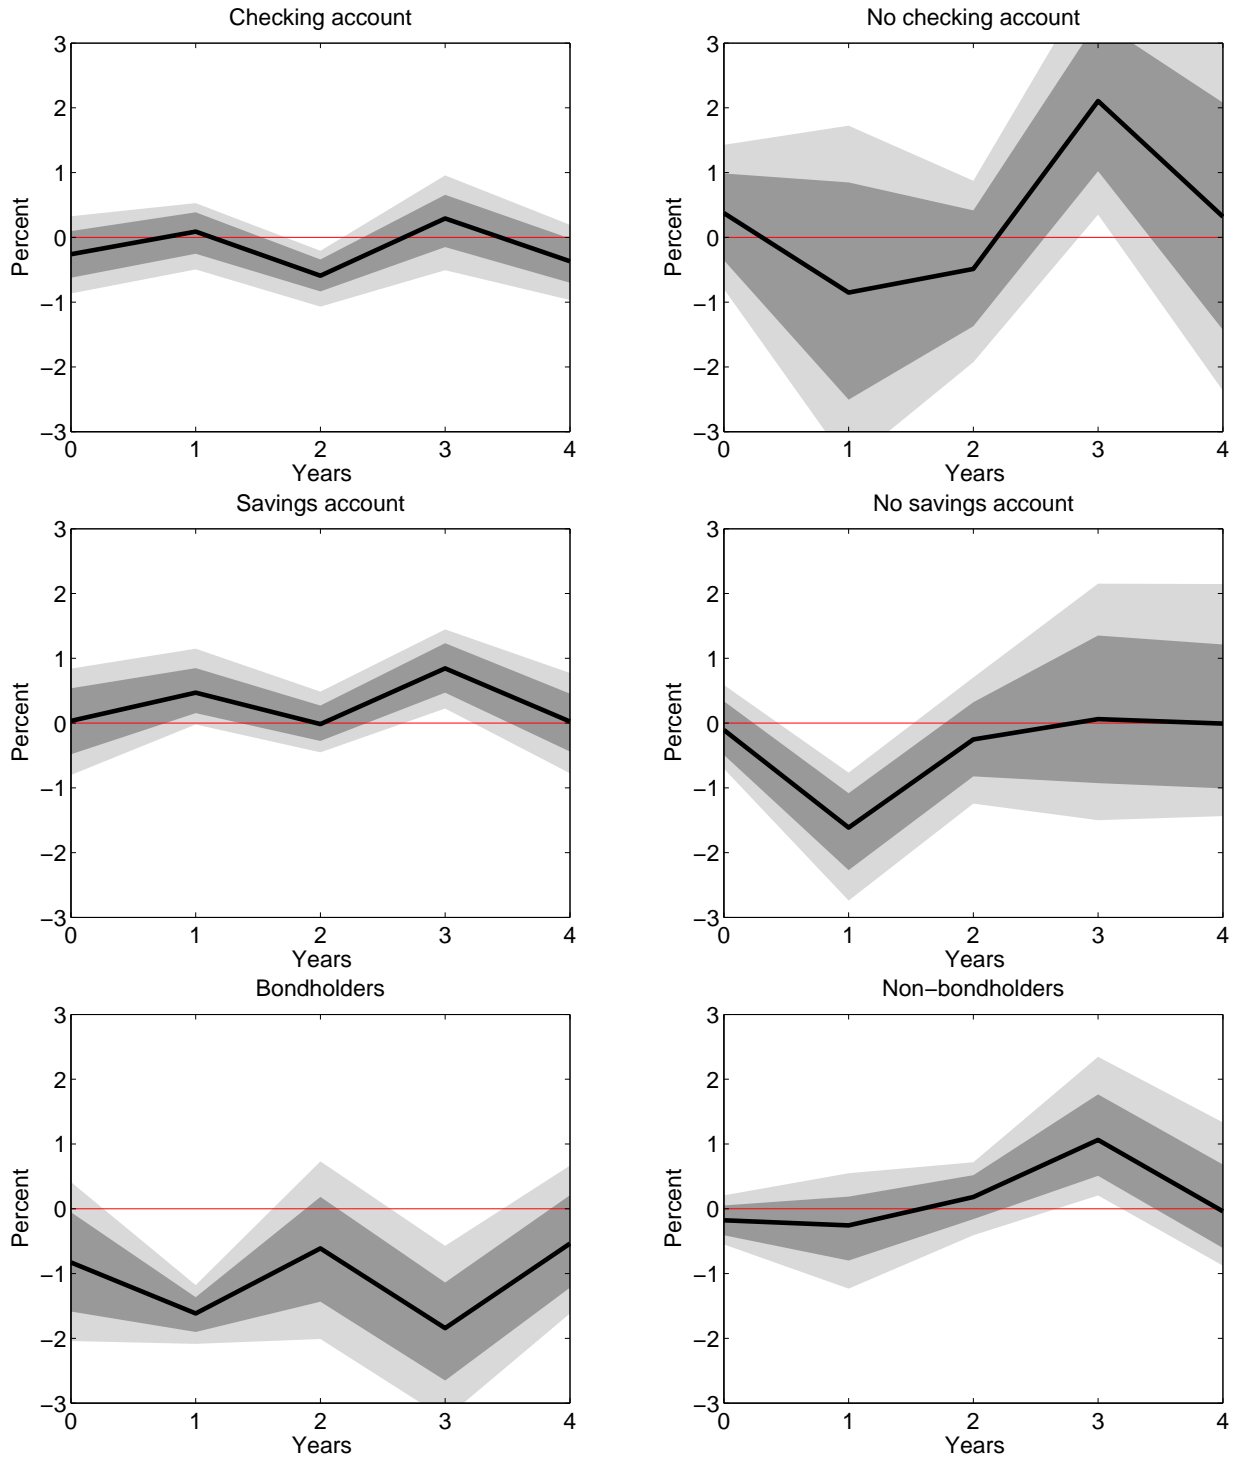
3.3 Splits based on holdings of other assets

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. Figure 6 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, recall in the stylized model presented above that stockholders increase their hours in order to reap the benefit of the higher return on capital. Figure 7 shows that the annual stock return, measured using S&P 500 index, indeed increases when technology improves.¹⁷ In contrast, returns on treasury bills or bonds do not increase after a positive technology shock (Figure 8). 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.

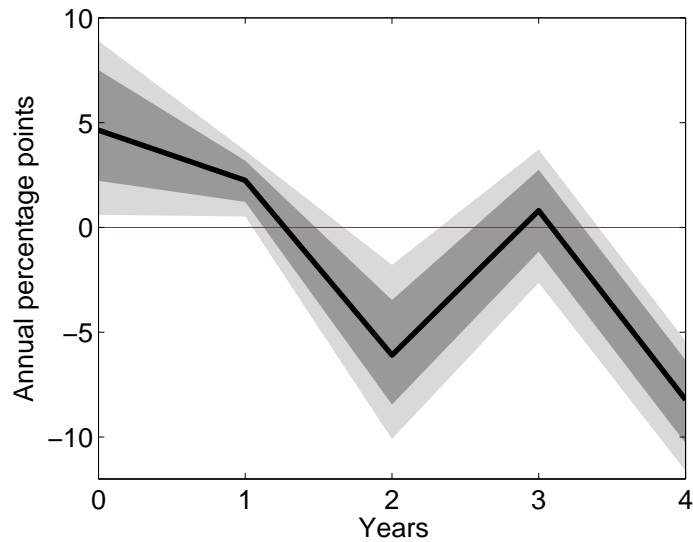
¹⁷It is important to point out that the increase in the stock return in response to a positive technology shock does not necessarily imply that the market is inefficient. Indeed, there is an extensive literature documenting the predictability of stock returns. Fama and French (1988), for example, find 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.

Figure 6: Splits based on holdings of other assets



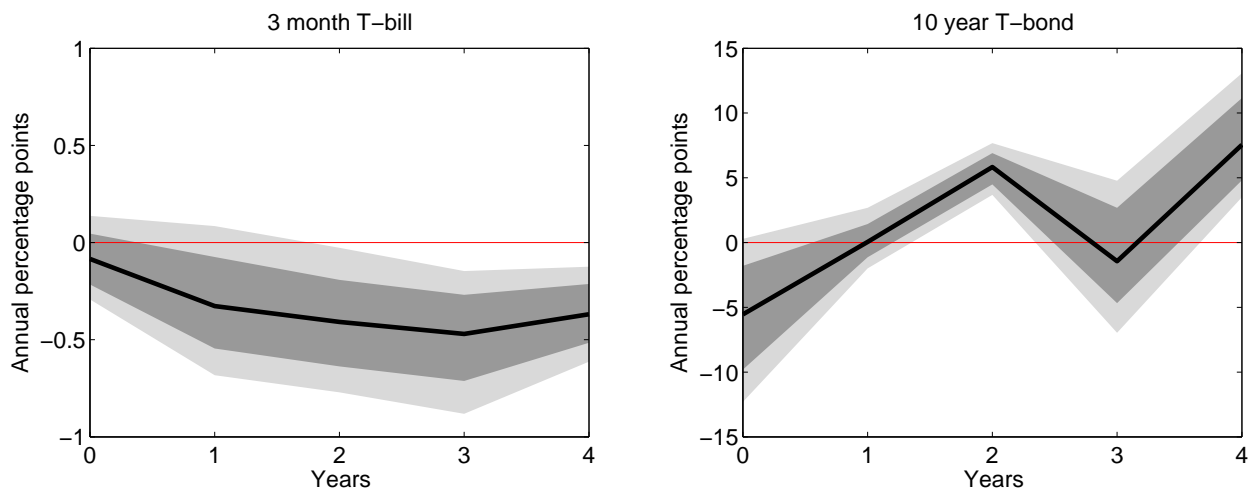
Notes: Impulse responses of the levels of hours to a one-standard-deviation positive technology shock, estimated from a regression on its own lags and current and lagged technology growth. Technology growth is measured using the utilization-adjusted TFP series constructed in Fernald (2014). The dark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

Figure 7: 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. See text for details on the construction. The sample period is 1981–2011.

Figure 8: 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. See text for details on the construction. The sample period is 1981–2011.

Figure 9: 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 are below the cutoff, including those with zero and negative wealth. The dark and light shaded areas are 68 and 90-percent confidence intervals, respectively. See text for details on the construction. The sample period is 1981–2011.

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). Figure 9 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.

Finally, related to the above issue regarding wealth, I also consider a classification based on the amount of total before-tax income. An important drawback to keep in mind regarding this classification is that current income is likely quite endogenous to the labor supply.¹⁸ As in the split based on total wealth, I set the cutoff so that the share of high income households is equal to the share of stockholders. Figure 10 shows that high income households do not increase hours after a

¹⁸In contrast to the asset holdings, it is not possible to know from the CEX the household income before the technology shock.

Figure 10: Split based on the total before-tax income



Notes: Impulse responses of the levels 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, \$37,220 in 1980 dollars (CPI for all urban consumers), is set so that the share of households that are classified to have high before-tax total income is equal to the share of stockholders (14% of the sample). The dark and light shaded areas are 68 and 90-percent confidence intervals, respectively. See text for details on the construction. The sample period is 1981–2011.

positive technology shock and, if anything, appear to reduce hours. Thus, the income difference between stockholders and non-stockholders cannot explain the heterogeneous impulse response.


3.4 Traditional sample splits

In comparison with the literature, I explore classifications based on standard demographic variables such as education and birth cohorts.¹⁹ This allows me to check whether the heterogeneity in the impulse responses I found in the baseline analysis is due to factors such as skilled vs. unskilled labor or life-cycle effects.

Figure 11 reports the movements of hours after a technology improvement for households whose heads have college degrees and those who don't. According to the point estimates, households with college degrees increase their hours by 0.5 percent on impact while households without reduce them by 0.5 percent. However, the heterogeneity in the impulse responses is less pronounced than when households are classified based on their stock holding statuses. For example, the initial response of the households with college degrees is not statistically significant at the 90 percent level. This suggests that the response of the college-educated households may be driven by the fact that many, but not all, of those households own stocks, as reported in Table 1. To further understand this result, in Figure 12, I classify stockholders based on whether the heads have college degrees or not.

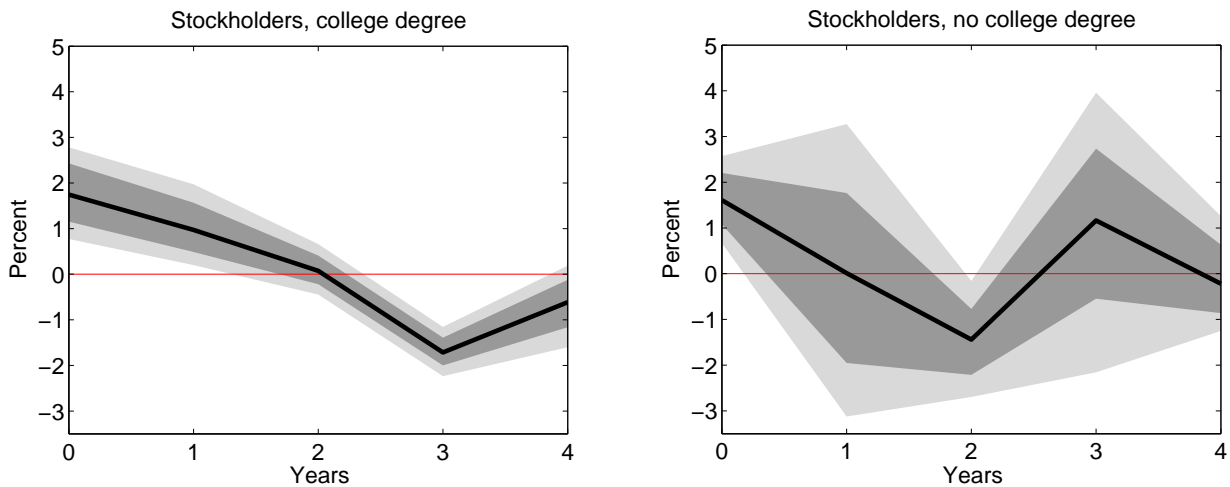
¹⁹I use classification based on birth cohorts instead of age following the microeconomic literature.

Figure 11: Split based on household heads' education (all households)




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. See text for details on the construction. The sample period is 1981–2011.

Figure 12: Split based on household heads' education among stockholders



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. See text for details on the construction. The sample period is 1981–2011.

Figure 13: Split based on household heads' birth cohorts



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). “Young”, “Middle-aged”, and “Old” refer to households whose heads were born 1955–, 1935–1954, and –1934, respectively. The dark and light shaded areas are 68 and 90-percent confidence intervals, respectively. See text for details on the construction. The sample period is 1981–2011.

On impact, both households with or without college degrees raise their hours by a similar amount and the responses are statistically significant. This evidence shows that the heterogeneity in the impulse responses I find for stockholders and non-stockholders is not because stockholders tend to be more educated. Finally, Figure 13 shows the classification based on whether the household head is born after 1954, between 1935 and 1954, and before 1935. For all cohorts, the initial responses are not statistically different from zero, suggesting that life-cycle effects cannot explain my main result.

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

4 A structural analysis

4.1 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 dynamic stochastic general equilibrium (DSGE) model with limited stock market participation based on Galí et al. (2007). 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. For the baseline analysis, I assume that non-stockholders do not hold any intertem-

poral smoothing device such as bonds. This is obviously a stark assumption but it highlights the difference in the strength of intertemporal substitution of labor supply between stockholders and non-stockholders in the cleanest way. I later extend the model to allow non-stockholders to hold an interest-bearing asset.

In each period, stockholders choose consumption C_t^s , nominal wage $W_{i,t}$, share holdings S_t^s , and bond holdings B_t^s to maximize utility:

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

where $H_{i,t}^s$ 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

$$P_t C_t^s + P_t S_t^s + B_t^s \leq W_{i,t} H_{i,t}^s + P_t D_t^s + P_t S_{t-1}^s + R_{t-1} B_{t-1}^s + Q_{i,t}^s + \Pi_t + T^s,$$

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 is the nominal wage, R_{t-1} is the gross nominal interest rate on risk-free bonds from period $t - 1$ to t . I assume that households buy securities, whose payoff $Q_{i,t}^s$ is contingent on whether they can re-optimize their wage.²⁰ Π_t is the combined profit of all the intermediate-goods firms 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

²⁰The existence of state-contingent securities ensures that stockholders are homogeneous with respect to consumption and asset holdings, even though they are heterogeneous with respect to the wage rate and hours because of the idiosyncratic nature of the timing of wage re-optimization. I assume that both stockholders and non-stockholders trade the securities among themselves. This implies that consumption and asset holdings could be different between the two groups.

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 $\delta(u_t)$.²¹ 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 face a simpler problem and choose consumption C_t^n and nominal wage $W_{i,t}$ to maximize utility:

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

where $H_{i,t}^n$ is hours worked, subject to the budget constraint:

$$P_t C_t^n \leq W_{i,t} H_{i,t}^n + Q_{i,t}^n + T^n.$$

Note that, because of the consumption habit, the income effect of wage changes could dominate the static substitution effect.²²

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

²¹I assume the function form: $\delta(u_t) = \delta_0 + \delta_1(u_t - 1) + \frac{\delta_2}{2}(u_t - 1)^2$, where $\delta_0 > 0, \delta_1 > 0, \delta_2 > 0$.

²²Without habit, the income and the static substitution effect would cancel out and hence non-stockholders' labor supply would be fixed.

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.²³

In each period t , a perfectly competitive representative employment agency hires labor from households to produce an aggregate labor service, H_t , using 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 $H_{i,t}$ denotes the time t input of labor service from household i and θ_w controls the price elasticity of demand for each household's labor service. The agency sells the aggregated labor input to the intermediate firms for a nominal price of W_t per unit. Households (both stockholders and non-stockholders) face a Calvo-type wage-setting friction: In each period t , a household can re-optimize its nominal wage with probability $(1 - \xi_w)$. Households that cannot re-optimize index their wage according to the steady-state wage growth rate, $\gamma\pi$.

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.²⁴ 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 $C_t + I_t = Y_t$.

4.2 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 are 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

²³In 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.

²⁴Galí et al. (2007) also make this assumption.

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).²⁵ 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

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

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

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

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.²⁶ These

²⁵Thus, the data and the theoretical model are treated symmetrically in the estimation (Kehoe 2006).

²⁶Output is defined as real GDP, consumption is defined as (expenditures on nondurables)+(expenditures on ser-

Table 2: Estimated parameters

	Description	Prior			Posterior mode		
		Type	Mean	Std.	Baseline	Extended	Rep. agent
α	Capital share	B	0.35	0.03	0.36 (0.03)	0.36 (0.03)	0.34 (0.03)
δ_2	Convexity of depreciation	G	0.50	0.20	0.43 (0.21)	0.40 (0.19)	0.41 (0.21)
η	Inverse Frisch elasticity	G	0.60	0.30	0.13 (0.24)	0.07 (0.08)	0.42 (0.29)
b	Consumption habit	B	0.50	0.20	0.88 (0.06)	0.89 (0.15)	0.54 (0.21)
κ	Bond holding cost	G	0.10	0.05	–	0.02 (0.08)	–
ξ_p	Calvo price	B	0.50	0.20	0.25 (0.15)	0.46 (0.17)	0.76 (0.18)
ξ_w	Calvo wage	B	0.50	0.20	0.14 (0.14)	0.16 (0.10)	0.50 (0.19)
ρ_R	Interest smoothing	B	0.60	0.10	0.60 (0.10)	0.51 (0.09)	0.66 (0.09)
ϕ_π	Inflation response	N	1.70	0.30	1.53 (0.23)	1.09 (0.27)	1.40 (0.20)
ϕ_Y	Output response	N	0.15	0.05	0.15 (0.04)	-0.01 (0.02)	0.17 (0.05)
$\phi_{\Delta Y}$	Output growth response	N	0.15	0.05	0.15 (0.05)	0.14 (0.05)	0.16 (0.05)
ρ_z	Technology shock	N	0.60	0.20	0.99 (0.13)	0.97 (0.01)	0.89 (0.09)
100σ	Technology shock	IG	1.00	1.00	0.27 (0.04)	0.23 (0.04)	0.28 (0.07)

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.

macro impulse responses are displayed in Figure 14. In response to a technology improvement, output, investment, consumption move little initially and increase over time. Nominal variables such as inflation and the 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.

4.3 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. Figure 14 shows the model impulse responses along with the data. 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.

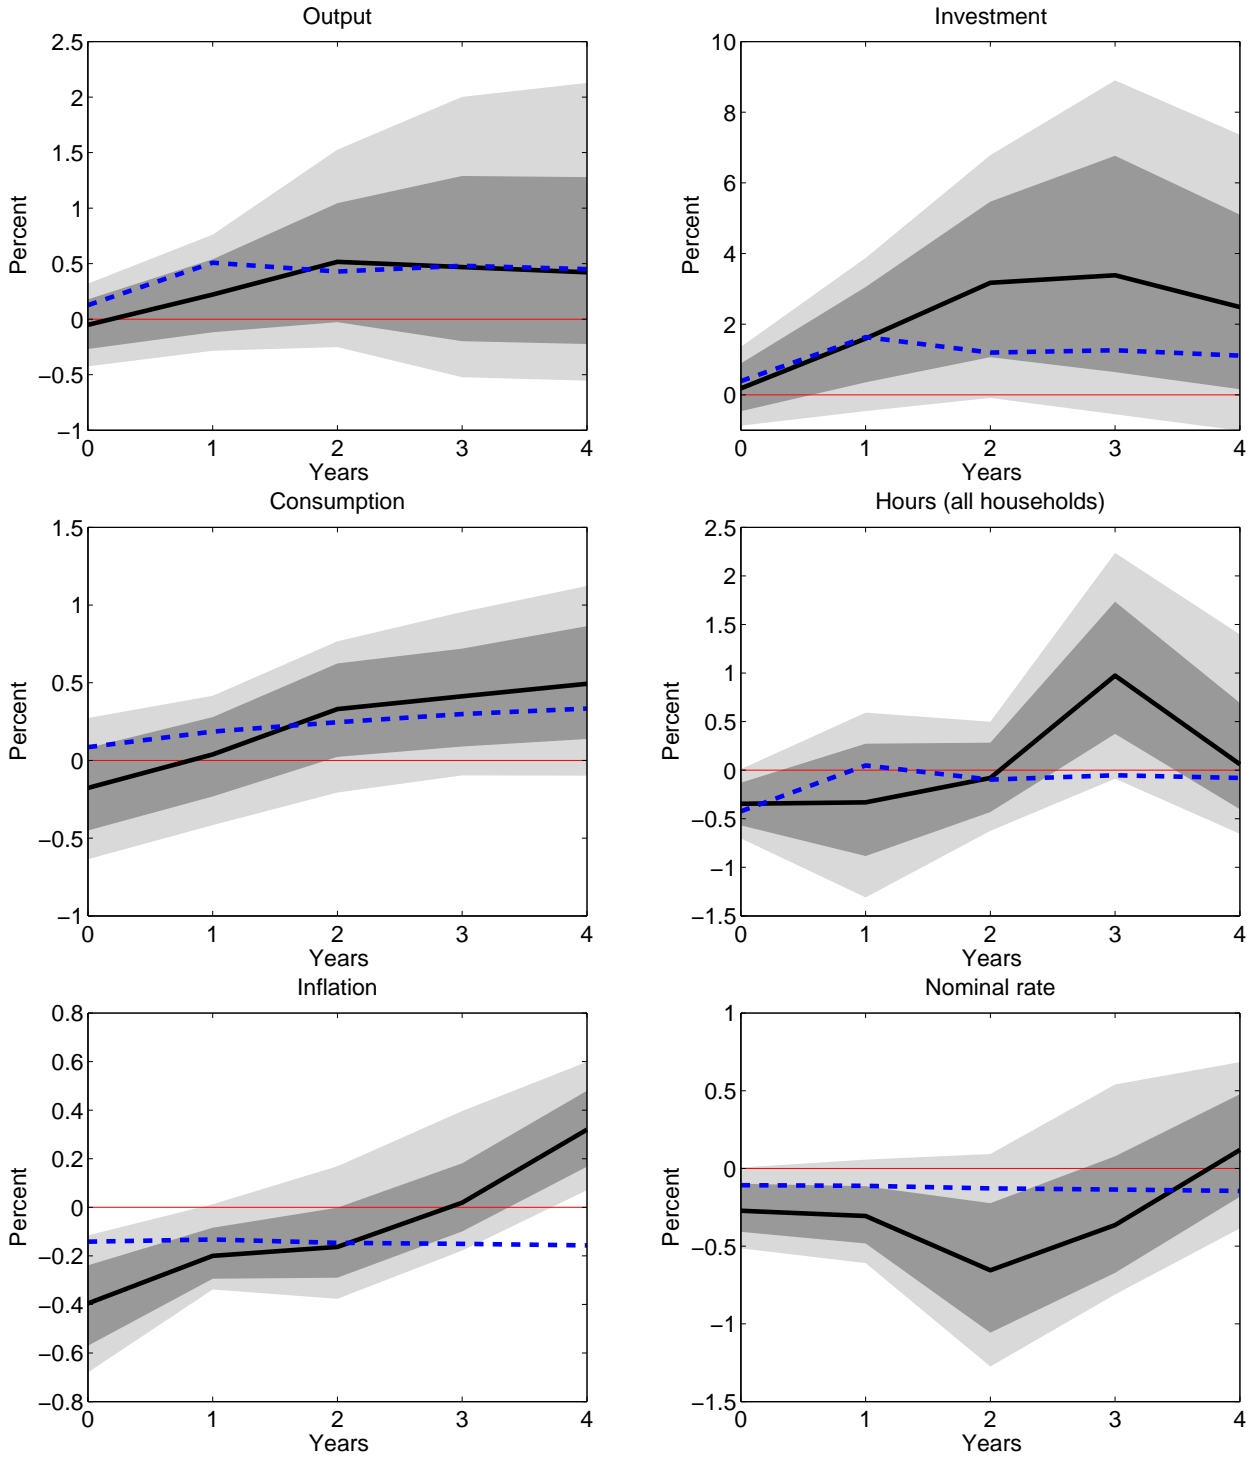
Next, consider the posterior of the baseline model (labeled “Baseline” in Table 2). The estimates of Calvo price and wage parameters are 0.25 and 0.14, respectively and are significantly smaller than the estimates from the full stock market participation model. The values imply that firms and households adjust their prices and wages roughly every quarter.²⁷

The blue dashed lines in Figure 15 show that the model is able to reproduce the heterogeneous response to a technology shock. The intuition is similar to the stylized model I presented earlier. 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 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 solid lines marked with circles in Figure 15). In this case, stockholders reduce their hours because a positive technology shock reduces the return on investment due to countercyclical markups. Figure 16 shows that, in

vices), 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. To ease comparison with the household-level results reported above, the macro impulse responses are estimated using the baseline regression (1).

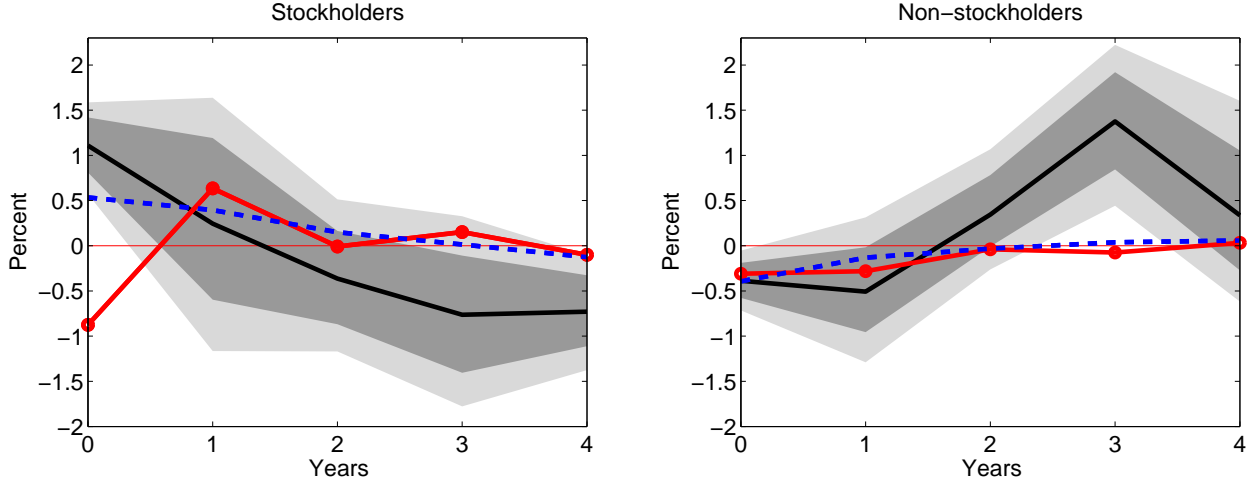
²⁷The estimated Frisch elasticity, $1/\eta$, is higher than the usual estimates. To explore this further, I have re-estimated the model by fixing η to 0.5, which implies a Frisch elasticity of 2, and found the main results unaffected. Further details are available upon request.

Figure 14: Data and model impulse responses: full stock market participation



Notes: Impulse responses from the data (black solid lines) and the model (blue dashed lines), estimated from a regression on its own lags and current and lagged technology growth.

Figure 15: Data and model impulse responses: limited stock market participation



Notes: Impulse responses from the data (black solid lines) and the model (blue dashed lines), estimated from a regression on its own lags and current and lagged technology growth. The red solid lines marked with circles 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.

in addition to the household-level hours, the heterogeneous-agent model is able to replicate the responses of other standard macro variables.

4.4 Extension: Allowing non-stockholders to hold assets

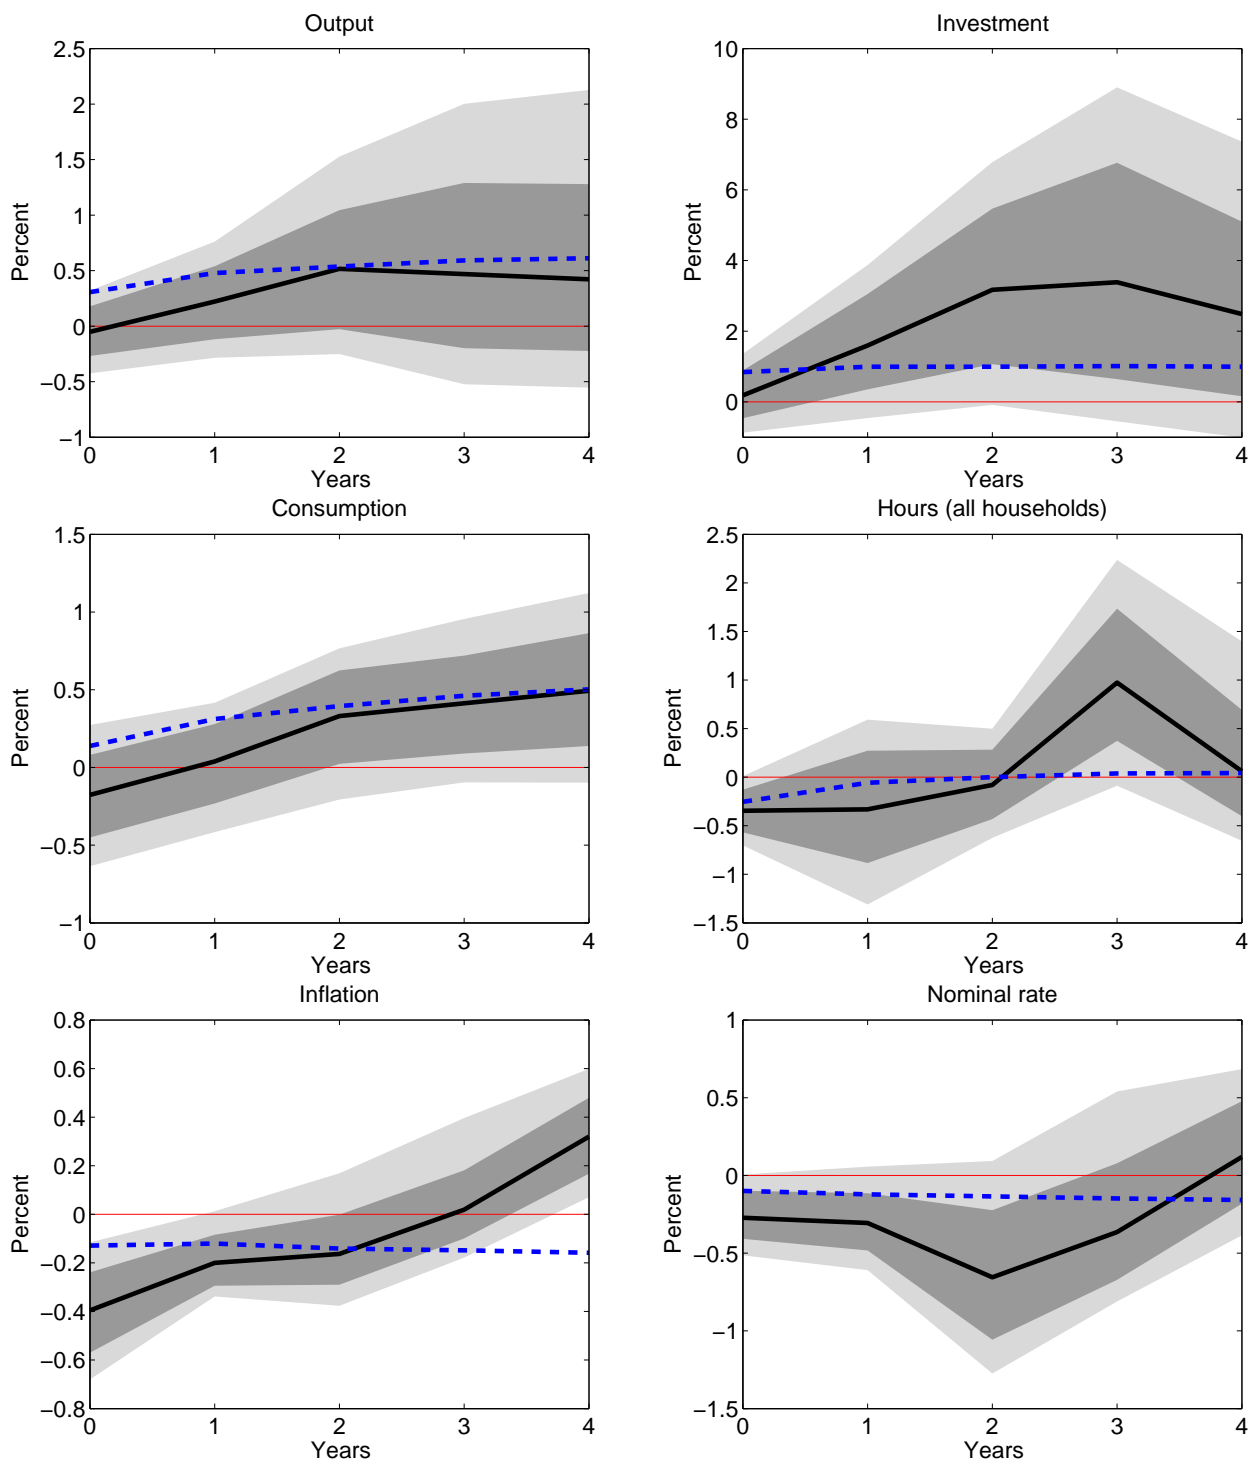
In the baseline model, non-stockholders do not have access to any intertemporal smoothing device. I now extend the model to allow for non-stockholders to hold assets.

Empirical results from Section 3.3 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. To capture this feature in the model, I assume that non-stockholders hold assets A_t^n that earn a constant interest rate \bar{r} . The budget constraint for a non-stockholder i is then

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

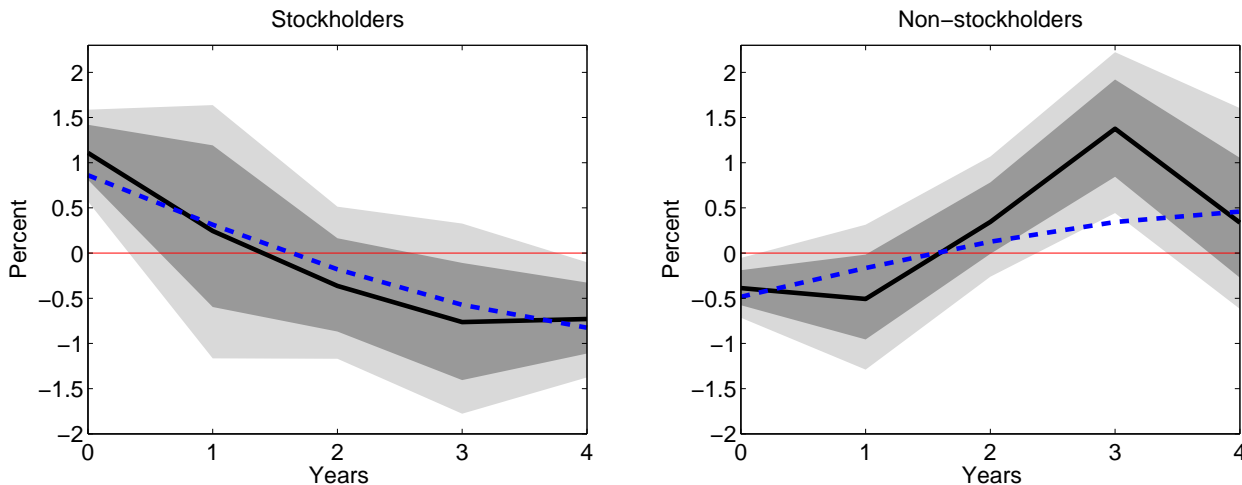
where the last term is the bond-holding cost, which I assume is transferred back to non-stockholders in a lump-sum manner. $\kappa > 0$ is a parameter that controls the size of the bond-holding cost. The cost eliminates non-stationarity otherwise built into models with exogenous interest rates and induces 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

Figure 16: Data and model impulse responses: limited stock market participation



Notes: Impulse responses from the data (black solid lines) and the model (blue dashed lines), estimated from a regression on its own lags and current and lagged technology growth.

Figure 17: Data and model impulse responses: extended model with limited stock market participation



Notes: Impulse responses from the data (black solid lines) and the model (blue dashed lines), estimated from a regression on its own lags and current and lagged technology growth.

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.²⁸ 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.²⁹

Column “Extended” in Table 2 reports the estimation result. Although the estimated Calvo price and wage parameters are larger than the estimates of the baseline limited stock market participation model, prices and wages are still considerably more flexible than the representative agent model; the estimated frequencies of price and wage adjustments are $(1/(1 - 0.46) \approx) 2$ and $(1/(1 - 0.16) \approx) 1$ quarters. Figure 17 shows that the extended model fits the household-level impulse responses quite well. In particular, the model is now able to replicate the overshooting of non-stockholders’ labor supply two years after the shock. Thus, the main results from the baseline limited stock market participation model – that the transmission mechanism of the heterogeneous-agent model is consistent with the micro evidence and its estimates of the sizes of nominal rigidities are smaller than the estimates from the representative-agent model – are robust to the extension that allows non-stockholders to hold assets.

²⁸The 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.

²⁹I have also tried a version where stockholders and non-stockholders share the same discount factor and the habit parameter and obtained similar results.

5 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 and cannot be explained by demographic factors. 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 found in Coibion et al. (2016) 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**, “Read All about It!! What Happens Following a Technology Shock?,” *American Economic Review*, 2011, *101*, 1144–1179.
- Altig, David, Lawrence J. Christiano, Martin Eichenbaum, and Jesper Lindé**, “Firm-Specific Capital, Nominal Rigidities and the Business Cycle,” *Review of Economic Dynamics*, 2011, *14*, 225–247.
- Anderson, Emily, Atsushi Inoue, and Barbara Rossi**, “Heterogenous Consumers and Fiscal Policy Shocks,” *Journal of Money, Credit, and Banking*, forthcoming.
- Attanasio, Orazio P., James Banks, and Sarah Tanner**, “Asset Holding and Consumption Volatility,” *Journal of Political Economy*, 2002, *110*, 771–792.
- Balvers, Ronald J., Thomas F. Cosimano, and Bill McDonald**, “Predicting Stock Returns in an Efficient Market,” *Journal of Finance*, 1990, *45*, 1109–1128.
- Basu, Susanto, John G. Fernald, and Miles S. Kimball**, “Are Technology Improvements Contractionary?,” *American Economic Review*, 2006, *96*, 1418–1448.

- Bilbiie, Florian**, “Limited asset markets participation, monetary policy and (inverted) aggregate demand logic,” *Journal of Economic Theory*, 2008, *140*, 162–196.
- Bocola, Luigi, Marcus Hagedorn, and Iourii Manovskii**, “Identifying Neutral Technology Shocks,” 2014. Manuscript.
- Broer, Tobias, Niels-Jakob Harbo Hansen, Per Krusell, and Erik Oberg**, “The New Keynesian Transmission Channel: A Heterogeneous-Agent Perspective,” 2016. Manuscript.
- Canova, Fabio, David Lopez-Salido, and Claudio Michelacci**, “The Effects of Technology Shocks on Hours and Output: A Robustness Analysis,” *Journal of Applied Econometrics*, 2010, *25*, 755–773.
- Chang, Yongsung and Jay H. Hong**, “Do Technological Improvements in the Manufacturing Sector Raise or Lower Employment?,” *American Economic Review*, 2006, *96*, 352–368.
- Chari, V.V., Patrick J. Kehoe, and Ellen R. McGrattan**, “Are Structural VARs With Long-Run Restrictions Useful in Developing Business Cycle Theory?,” *Journal of Monetary Economics*, 2008, *55*, 1337–1352.
- Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans**, “Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy,” *Journal of Political Economy*, 2005, *113*, 1–45.
- , – , and **Robert Vigfusson**, “What Happens After a Technology Shock?,” 2003. NBER Working Paper.
- , – , and – , “Unemployment and Business Cycles,” 2013. NBER Working Paper.
- , **Mathias Trabandt, and Karel Walentin**, “Involuntary Unemployment and the Business Cycle,” 2010. NBER Working Paper.
- Cloyne, James and Paolo Surico**, “Household Debt and the Dynamic Effects of Income Tax Changes,” forthcoming. *Review of Economic Studies*.
- , **Clodomiro Ferreira, and Paolo Surico**, “Monetary Policy When Households Have Debt: New Evidence on the Transmission Mechanism,” 2015. Manuscript.
- Coibion, Olivier and Yuriy Gorodnichenko**, “What Can Survey Forecasts Tell Us about Information Rigidities?,” *Journal of Political Economy*, 2012, *120*, 116–159.
- , – , **Lorenz Kueng, and John Silvia**, “Innocent Bystanders? Monetary Policy and Inequality in the U.S.,” 2016. Manuscript.

- Colciago, Andrea**, “Rule-of-Thumb Consumers Meet Sticky Wages,” *Journal of Money, Credit, and Banking*, 2011, *43*, 325–353.
- De Giorgi, Giacomo and Luca Gambetti**, “Business Cycle Fluctuations and the Distribution of Consumption,” *Review of Economic Dynamics*, 2017, *23*, 19–41.
- Dupor, Bill, Jing Han, and Yi-Chan Tsai**, “What Do Technology Shocks Tell Us About the New Keynesian Paradigm?,” *Journal of Monetary Economics*, 2009, *56*, 560–569.
- Erceg, Christopher J., Luca Guerrieri, and Christopher Gust**, “Can Long-Run Restrictions Identify Technology Shocks?,” *Journal of the European Economic Association*, 2005, *3*, 1237–1278.
- Fama, Eugene F. and Kenneth R. French**, “Permanent and Temporary Components of Stock Prices,” *Journal of Political Economy*, 1988, *96*, 246–273.
- Fernald, John G.**, “Trend Breaks, Long-Run Restrictions, and Contractionary Technology Improvements,” *Journal of Monetary Economics*, 2007, *54*, 2467–2485.
- , “A Quarterly, Utilization-Adjusted Series on Total Factor Productivity,” 2014. Manuscript.
- Francis, Neville and Valerie A. Ramey**, “Is the Technology-Driven Real Business Cycle Hypothesis Dead? Shocks and Aggregate Fluctuations Revisited,” *Journal of Monetary Economics*, 2002, *110*, 1379–1399.
- Furlanetto, Francesco and Martin Seneca**, “Rule-of-Thumb Consumers, Productivity, and Hours,” *Scandinavian Journal of Economics*, 2012, *114*, 658–679.
- Galí, Jordi**, “Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?,” *American Economic Review*, 1999, *89*, 249–271.
- , **J. David López-Salido, and Javier Vallés**, “Rule-of-Thumb Consumers and the Design of Interest Rate Rules,” *Journal of Money, Credit, and Banking*, 2004, *36*, 739–763.
- , – , and – , “Understanding the Effects of Government Spending on Consumption,” *Journal of the European Economic Association*, 2007, *5*, 227–270.
- Gorodnichenko, Yuriy and Michael Weber**, “Are Sticky Prices Costly? Evidence from the Stock Market,” *American Economic Review*, 2016, *106*, 165–199.
- Guvenen, Fatih**, “Reconciling Conflicting Evidence on the Elasticity of Intertemporal Substitution: A Macroeconomic Perspective,” *Journal of Monetary Economics*, 2006, *53*, 1451–1472.
- , “A Parsimonious Macroeconomic Model for Asset Pricing,” *Econometrica*, 2009, *77*, 1711–1740.
- Investment Company Institute**, “Equity and Bond Ownership in America, 2008,” 2008. Available at <http://www.ici.org/research/investors>.

- Kehoe, Patrick J.**, “How to Advance Theory with Structural VARs: Use the Sims-Cogley-Nason Approach,” 2006. NBER Working Paper.
- Lindé, Jesper**, “The Effects of Permanent Technology Shocks on Hours: Can the RBC-Model Fit the VAR Evidence?,” *Journal of Economic Dynamics and Control*, 2009, *33*, 597–613.
- Liu, Zheng and Louis Phaneuf**, “Technology Shocks and Labor Market Dynamics: Some Evidence and Theory,” *Journal of Monetary Economics*, 2007, *54*, 2534–2553.
- Malloy, Christopher J., Tobias J. Moskowitz, and Annette Vissing-Jørgensen**, “Long-Run Stockholder Consumption Risk and Asset Returns,” *Journal of Finance*, 2009, *64*, 2427–2479.
- Nakamura, Emi and Jon Steinsson**, “Fiscal Stimulus in a Monetary Union: Evidence from US Regions,” *American Economic Review*, 2014, *104*, 753–792.
- Newey, Whitney and Kenneth D. West**, “A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix,” *Econometrica*, 1987, *55*, 703–708.
- Romer, Christina D. and David H. Romer**, “A New Measure of Monetary Shocks: Derivation and Implications,” *American Economic Review*, 2004, *94*, 1055–1084.
- Rotemberg, Julio J.**, “Stochastic Technical Progress, Smooth Trends, and Nearly Distinct Business Cycles,” *American Economic Review*, 2003, *93*, 1543–1559.
- Schmitt-Grohe, Stephanie and Martin Uribe**, “Closing small open economy models,” *Journal of International Economics*, 2003, *61*, 163–185.
- Shea, John**, “What Do Technology Shocks Do?,” in “NBER Macroeconomics Annual 1998,” Vol. 13, MIT Press, 1999, pp. 275–322.
- Uhlig, Harald**, “Do Technology Shocks Lead to a Fall in Total Hours Worked?,” *Journal of the European Economic Association*, 2004, *2*, 361–371.
- Vissing-Jørgensen, Annette**, “Limited Asset Market Participation and the Elasticity of Intertemporal Substitution,” *Journal of Political Economy*, 2002, *110*, 825–853.
- Walsh, Carl E.**, “Workers, Capitalists, Wages, Employment and Welfare,” 2016. Manuscript.


Appendix (for online publication)

A Probit estimate from the SCF

	Parameter	Standard error
Age	0.044	0.002
(Age) ²	-0.001	0.000
Highschool	0.044	0.010
College	0.340	0.011
Nonwhite	-0.331	0.010
<i>Log</i> (income)	0.314	0.009
<i>Log</i> (chk+sav)	0.141	0.003
Chk+sav = 0	0.105	0.025
Dividend > 0	1.241	0.017
Y ₁₉₉₂	0.171	0.020
Y ₁₉₉₅	0.349	0.020
Y ₁₉₉₈	0.566	0.020
Y ₂₀₀₁	0.630	0.019
Y ₂₀₀₄	0.620	0.020
Y ₂₀₀₇	0.695	0.020
Y ₂₀₁₀	0.625	0.018
Constant	-5.501	0.080
Observations	17,7565	
Pseudo R^2	0.306	

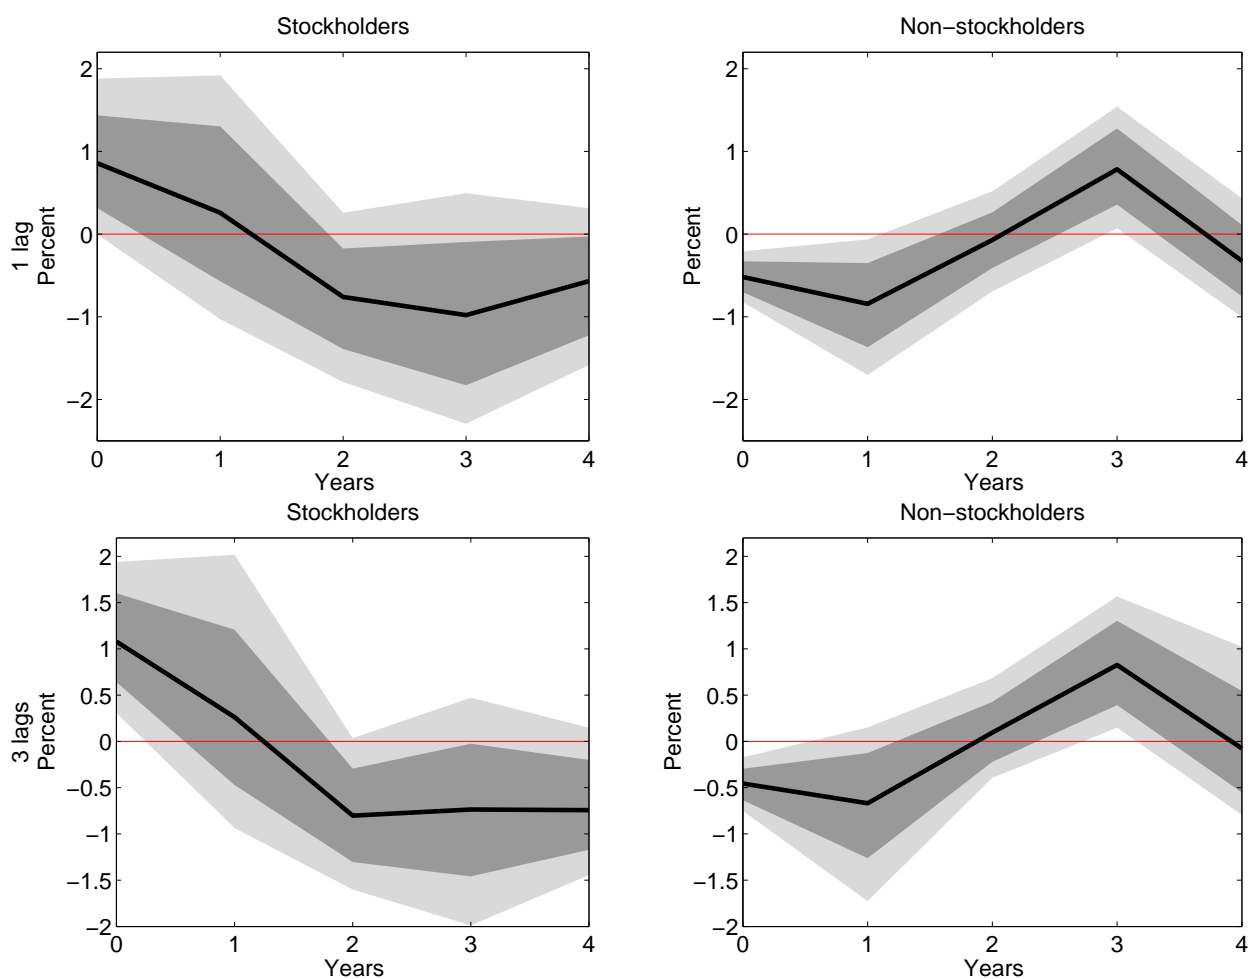
B Additional impulse responses from the CEX

Figure 18: Impulse responses for single households



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. The sample is restricted to single households. The dark and light shaded areas are 68 and 90-percent confidence intervals, respectively. See text for details on the construction. The sample period is 1981–2011.

Figure 19: Impulse response under alternative lags



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. The dark and light shaded areas are 68 and 90-percent confidence intervals, respectively. See text for details on the construction. The sample period is 1981–2011.