

Demand Elasticities, Nominal Rigidities and Asset Prices

Job Market Paper

Nuno Clara *

November, 2019

Abstract

I study heterogeneity in demand elasticities as a source of risk in asset pricing. I use high-frequency product price and quantity data from Amazon to overcome the classical endogeneity challenge and estimate demand elasticities for a large cross-section of firms. I find that firms facing more elastic demands are riskier and this is reflected in higher equilibrium average stock returns (6.2% annual return premium). I show that price stickiness contributes to these differences in risk: firms are slow to react to competitors' price changes (compared to the predictions of a Calvo model). My empirical findings show the importance of having heterogeneous demand elasticities and degrees of price stickiness in macro-finance models.

*Department of Finance, London Business School. I am particularly indebted to João Cocco, Francisco Gomes, Howard Kung and Anna Pavlova for their invaluable guidance and advice. I am also grateful to Simcha Barkai, Svetlana Bryzgalova, John Y. Campbell, Roberto Cram, Marco Grotteria, Nick Hirschey, Minsoo Kim, Thilo Kind, Christopher Lako (discussant), Sérgio Rebelo, Roberto Rigobon, Arkodipta Sarkar, Pasquale Schiraldi, Lukas Schmid (discussant), Henri Servaes, Varun Sharma, Dejanir Silva (discussant), Mikhail Tirkish, Vikrant Vig, Michael Weber, and seminar participants at the Society of Economic Dynamics, the ECB Young Economists Forum, Transatlantic Doctoral Conference, the QMUL Economics and Finance Workshop, the Lubrafin, the European Finance Association, the HEC PhD Workshop, for their insightful comments. This paper received the 2019 AQR Asset Management Institute Fellowship Award and the Best Paper award in the ECB Young Economists' competition. I thank the Deloitte Institute of Innovation and Entrepreneurship for financial support.

I. Introduction

Marshall's (1890) seminal work developed the concept of demand elasticity: the extent to which demand varies in response to a price change. Since then the concept has been central in economics (e.g. Gali (1994), Raith (2003)). This paper studies the role of demand in asset pricing, namely the relation between demand elasticities and frequency of price adjustment at the firm level and their joint implications for firm fundamentals and asset prices. The asset pricing literature, often assumes that firms face the same elasticity which overlooks the relation between demand elasticity, nominal rigidities and asset prices. In a standard macro-finance model, nominal rigidities generate operational leverage in firms and create a role for demand elasticity. The main empirical contribution of my paper is to estimate demand elasticities for a large cross-section of products and firms and understand the risk implications of differences in elasticity. I overcome the classical endogeneity challenge in estimating demand elasticities by using a novel high-frequency product level dataset from Amazon.

While the relation between sticky prices and asset pricing has been studied in the literature (namely Weber (2015) and Gorodnichenko and Weber (2016)), the role that demand elasticity plays in either nominal frictions or asset pricing has not. Further, most asset pricing accounts for risk and returns of firms focus on supply side factors, such as investment, leverage, profitability, value, liquidity, among others.¹ This paper instead focuses on the demand side.

The first building block of my analysis is to develop a New-Keynesian model where firms are heterogeneous in the demand elasticity that they face, potentially correlated with their degree of nominal rigidities. To do so, I generalize the standard multi-sector New-Keynesian model (e.g. Carvalho (2006), Nakamura and Steinsson (2010) and Weber (2015)), in which firms face nominal rigidities in adjusting their prices and allow for both heterogeneous degrees of price stickiness and heterogeneous demand elasticities. I calibrate the benchmark model to generate reasonable macroeconomic and asset pricing moments.

In the model when faced with a shock, firms' optimal prices change and, all else equal, move more in the sector with high elasticity of demand. The existence of nominal frictions leads firms in this sector to be further away from their optimal reset price, thus making them riskier. In equilibrium, their markups co-move more with marginal utility, thus yielding higher

¹There is a vast literature studying supply side factors and the cross-section of returns: Fama and French (2015), Zhang (2005), Carlson, Fisher, and Giammarino (2004), Kogan and Papanikolaou (2013), Pástor and Stambaugh (2003), Gomes and Schmid (2010), among many others.

expected returns. In good times, firms in the high-elasticity sector pay more dividends than firms in the low-elasticity sector, precisely because their markups co-move more with the business cycle. This makes firms in the high-elasticity sector pay more when the marginal utility of wealth is lower, so they are riskier in equilibrium. The benchmark calibration generates a 7% return spread between firms facing high elasticities of demand versus firms facing more inelastic demands.

In the model expected returns weakly line up with CAPM betas - firms facing high elasticities have higher returns and higher CAPM betas - this is because their cashflows co-move more with aggregate market cashflows.

The model also predicts that firms in the high-elasticity sector have lower equilibrium markups: there is a one-to-one mapping between steady-state markups and demand elasticity parameters, with firms with higher elasticity having lower markups. This is an important model prediction, as it allows me to distinguish its mechanism from that of Weber (2015) where heterogeneity is only in the degree of nominal rigidities. In fact, I show that if heterogeneity among firms comes only from nominal rigidities, then firms with more sticky prices have higher returns and higher markups. This is due to a precautionary savings motive, as firms with higher degrees of nominal rigidities fear locking a low markup more and therefore charge slightly higher markups.

To test the model predictions, I need firm-level estimates of demand elasticity. There are several well-known challenges in estimating them. The first is the standard endogeneity problem: firms' decisions to change prices are endogenous and empiricists only observe equilibrium prices and quantities, which makes it difficult to estimate the slopes that generated the equilibrium outcomes. This classical challenge can be overcome either by relying on parametric assumptions regarding the shape of the demand (e.g. Broda and Weinstein (2010) and Feenstra (1994)), or by using instruments to trace out demand (e.g. Berry, Levinsohn, and Pakes (1995), Hausman (1996) and Nevo (2001)). Further, even when instruments can be found, it is hard to do so for a large cross-section of firms.

The main empirical contribution of this paper is the development of a novel approach to estimating demand elasticities for a large cross-section of firms. I use high-frequency micro-level product data (prices and quantities) provided by **Keepa**, one of the largest Amazon US product trackers. I construct a database from scratch, with the price and quantity data of more than 280 thousand products sold by US public firms. I address the identification challenges

by looking into how quantity moves in a very narrow window around a price change. In particular, I estimate elasticity by measuring quantity demanded right before the price change and I see how quantity evolves within a short time-frame (12 hours) after the price has changed. This identification strategy to measure demand elasticities draws on the vast literature of high frequency identification of monetary policy shocks (e.g. Chodorow-Reich (2014), Gertler and Karadi (2015), Nakamura and Steinsson (2018), among others).

Naturally, price changes are endogenous, and my identification could fail if changes in prices are happening in response to a shift in the demand curve. To address this issue, I ensure that in these narrow windows, there are no demand shifts or expected demand shifts, and that therefore only the price shift is affecting the quantity demanded. I exclude price changes that occur when demand shifts are predictable (such as holidays and sales periods).

Second, my strategy could fail if firms are quick to react to competitors' price changes, which would trigger a shift in the demand curve faced by the firm for which I am estimating demand elasticity. I ensure that in such narrow windows, there are no competitors moving their prices. In fact, I find that firms are slow to adjust their prices when competitors change theirs. The degree of price synchronization across products that are close substitutes is 4%: an order of magnitude lower than the Calvo (1983) model would imply, where price changes are purely random. Even in such a case, the degree of price synchronization is higher than what is observable in the data.

The fact that firms are slow to react to competitors price changes is an important empirical result, because it sheds light on the nature of nominal frictions. Take for example a menu cost model: in such a model, price synchronization is one - if a firm changes its price, all firms that produce goods that are close substitutes change their prices. The same is not true for models such as the Calvo model, where the degree of price synchronization is lower than one, but even such a model fails to match the observable low degree of price synchronization in the data.

To further address endogeneity of price changes, I use standard instrumental variable methods to estimate demand elasticity. The main drawback of using instruments to trace demand is the reduction in sample size, as it is not possible to find an instrument for every product in the sample: an issue that does not arise with the high-frequency identification of demand elasticities. In particular, I follow Hausman (1996) and Nevo (2001) and use prices of the same products on Amazon Canada to instrument for demand. These instruments rely on the as-

sumption that the demand shocks are uncorrelated across the two markets.² Using instruments instead of the high-frequency strategy to estimate demand elasticity does not change the main results of the paper.

The estimated elasticities have several reasonable economic properties. First, elasticities are negative, as standard microeconomic theory would predict. Second, consumer goods industries such as clothing and non-durables have a more inelastic demand than durables and manufacturing industries, which is consistent with the fact that agents' have the choice of holding a durable good for longer and waiting for a price decrease (e.g. Gowrisankaran and Rysman (2012)). Third, elasticities seem to be fairly stable over time.

To study the impact of heterogeneous demand elasticities and nominal rigidities on firm fundamentals and asset prices, I merge the estimated elasticities with CRSP and Compustat. My sample contains an average of a thousand products per firm and 250 public firms. Returns monotonically increase for stocks sorted on demand elasticity: there is a 6.2% return differential between high- and low-elasticity firms. This return premium is statistically and economically meaningful. I show that the return spreads of elasticity-sorted portfolios are weakly explained by systematic risk or CAPM betas. In addition, I show that firms in highly elastic sectors have lower markups over marginal costs. This is in line with the prediction of the theoretical model: firms with a lower degree of monopoly power charge lower prices relative to costs. This is an important result, because a potential alternative explanation for the difference in expected returns is differences in the frequency of price adjustment (Weber (2015)). I show that heterogeneity in price stickiness and heterogeneity in elasticities have similar implications for expected returns: both higher elasticities and higher degrees of price stickiness imply higher expected returns. The implications are, however, opposite for markups. Higher elasticities imply lower equilibrium markups whereas higher degrees of price stickiness imply higher equilibrium markups. Empirically, the elasticity-sorted portfolios have similar degrees of price stickiness and higher-elasticity portfolios have lower markups.

The last section of this paper is dedicated to showing that the asset pricing results are robust. The empirical finding that firms with higher elasticities earn higher returns holds in a large out-of-sample period. Also, I show that the results are not driven by industry-specific

²In particular, Hausman assumes that a price in one market is a function of the average marginal costs of a product plus a markup amount that the firm is able to charge due to customers differing willingness to pay for that product in that specific market. If demand shocks for a product in Canada and the US are uncorrelated, then the price of the product in Canada will be correlated with the price of the same product in the US, due to common marginal costs.

characteristics. Even within industries, there is significant heterogeneity in the elasticity of demand across firms; using standard panel regressions, I show that firms with higher elasticities have higher returns even after controlling for time- and industry-fixed effects.

The paper is structured as follows: Section II offers a literature review, Section III lays down a theoretical framework to guide the empirical tests, Section IV describes the data and identification strategy, Section V shows the empirical results and Section VI concludes.

II. Literature Review

This paper is related to the literature on sticky prices and the intersection of sticky prices and asset pricing.

First, my paper is related to the empirical literature on sticky prices in retail markets. Cavallo (2017) undertook the first large-scale comparison of prices simultaneously collected from online stores and physical stores (e.g. Walmart online and Walmart stores), in 56 multi-channel retailers in 10 countries. He finds that online and offline prices are similar 72% of the time and that online and offline price changes have similar frequencies and magnitudes. He also finds that, 40% of the time, Amazon prices are identical to those of the products in stores, which is surprising as Amazon is a different retailer. This means that Amazon price data can be used to make inferences about physical retail markets and not only about online retail.

Furthermore, Cavallo (2016) finds that standard product-level datasets such as the Consumer Price Index (CPI) dataset and scanner datasets (for example the Nielsen dataset) suffer from several biases. On the one hand, price imputation and substitutions of temporarily missing products in the CPI dataset do not allow researchers to correctly know the quantity and price of a given product. On the other hand, the Nielsen dataset also suffers from a bias due to weekly averaging of individual product prices, which misses intra-week temporary shifts in prices, such as discounts and stock availability. Monthly and weekly data collection (a characteristic of these datasets) makes it difficult to disentangle shifts in the demand curve from shifts along the demand curve due to price changes. Scraped data such as the data collected by *Keepa*, which I use in this study, makes it possible to circumvent the issues of price imputation and averaging, as one can observe quantity and prices at higher frequencies. These high-frequency observations also make it easier to identify shifts along the demand curve and consequently to estimate demand elasticities.

Second, my paper relates to the literature on demand elasticity estimation. One common

way to estimate elasticities is to rely on instruments (e.g. Berry, Levinsohn, and Pakes (1995), Nevo (2001) and Hausman (1996)) to trace out demand. Due to the difficulty of finding good instruments, this approach is rarely applied to a large cross-section of products. An alternative approach is to make parametric assumptions regarding the shape of demand and supply. For example, Feenstra (1994) and Broda and Weinstein (2010) assume that demand and supply curves are linear in logarithm and that elasticities of products within a given product group are the same. This allows them to estimate constant demand and supply elasticities within product groups using panel data. I propose an alternative approach to estimating demand elasticities, which makes use of high-frequency data and allows me to make use of the full sample. As a robustness test, I run the same estimation using Hausman (1996) instruments to trace demand.

Finally, my paper is related to the asset pricing literature in the presence of nominal rigidities. The closest papers to mine are Gorodnichenko and Weber (2016) and Weber (2015), which look at how price stickiness relates to asset prices.

III. Framework

I develop a general equilibrium model to study the relation between demand elasticities and asset prices. The model will guide the empirical analysis, and will allow me to investigate potential alternative channels and test them in the data. In the appendix AI, I provide the basic intuition and qualitatively describe the main mechanism at play in a static one-period model. In the main text below, I develop a dynamic stochastic sectoral equilibrium model to evaluate the mechanism quantitatively.

The model features firms that exogenously face different elasticities of demand which may be correlated with the degree of nominal frictions.

A. Dynamic Model

I develop a quantitative neoclassical multi-sector equilibrium model in which sectors face different elasticities of demand and evaluate differences in risk among firms in the different sectors. The model is a generalization of the New-Keynesian multisector models (e.g. Carvalho (2006), Nakamura and Steinsson (2010) and Weber (2015)). In the aforementioned models, the heterogeneity among sectors comes from differences in the degree of price stickiness. In my model, firms in different sectors can face both different demand elasticities and different degrees of nominal rigidity. In this section, I lay down the main model equations and leave details and

derivations for the appendix.

The model delivers sharp quantitative implications that can be tested in the data. Sections A.1, A.2 and A.3 lay out the set of agents in the model and their optimization problems. Section A.4 calibrates the model. Section A.5 describes the model fit, its implications and the main mechanism underlying the results.

A.1. Households

There is a continuum of households indexed by $i \in [0, 1]$. Each household derives utility from consumption goods and supply differentiated labor services ($n_{t,i}$). The representative household has an utility function separable in a consumption bundle (C_t) and labor and maximizes:

$$E_t \sum_{s=0}^{\infty} \beta^s \left[u(C_{t+s} - \nu C_{t+s-1}) - \int_0^1 v(n_{t+s,i}) di \right] \quad (1)$$

where β is the time discount factor. The utility exhibits external habits in consumption (as in Christiano, Eichenbaum, and Evans (2005)), with intensity governed by the parameter ν . Each agent has some monopoly power in the labor market and posts the wage at which he/she is willing to supply labor services to firms that demand them.

Households consume goods produced in two sectors $k \in \{1, 2\}$.³ In each sector k there is a continuum of heterogeneous goods $j \in [0, 1]$ being produced. The output of each sector is given by the Dixit-Stiglitz aggregator over the different varieties:

$$C_{k,t} = \left[\int_0^1 C_{t,k,j}^{\frac{\eta_k-1}{\eta_k}} dj \right]^{\frac{\eta_k}{\eta_k-1}} \quad k \in \{1, 2\} \quad (2)$$

The sectors are heterogeneous in terms of the elasticity of demand that they face, i.e. holding everything else constant, the demand elasticity of any product j from sector k is given by:

$$\frac{\partial C_{t,k,j}}{\partial P_{t,k,j}} \frac{P_{t,k,j}}{C_{t,k,j}} = -\eta_k \quad (3)$$

Finally, households bundle the consumption from each sector into an overall consumption basket, C_t , using an upper Dixit-Stiglitz aggregator with elasticity of substitution between the two sectors equal to η . Without loss of generality, I assume that the elasticity of substitution between sectors is lower than within sectors: $\eta_1 > \eta_2 > \eta$. The representative agent Euler

³The model can be easily generalizable to have an arbitrarily larger number of sectors

equation is given by:

$$1 = \beta R_t E_t \left[\frac{1}{\pi_{t+1}} \frac{(C_{t+1} - bC_t)^{-\gamma}}{(C_t - bC_{t-1})^{-\gamma}} \right] \quad (4)$$

A.2. Wage Setting

I follow Erceg, Henderson, and Levin (2000) and Woodford (2013, chapter 4.1), and I model staggered wage contracts à la Calvo (1983): in each period, only a fraction $1 - \theta_w$ of households, drawn randomly from the population, reoptimize their posted nominal wage. There is a single labor market, with producers of all goods facing the same wages. However, the labor used to produce each good is a CES aggregate of the continuum of types of labor supplied by the representative household:⁴

$$N_t \equiv \left[\int_0^1 n_{i,t}^{\frac{\eta_w-1}{\eta_w}} di \right]^{\frac{\eta_w}{\eta_w-1}} \quad (5)$$

where η_w is the elasticity of substitution across different labor types. Consider a household resetting its wage, $w_{t,i}$, in period t . Given the household marginal utility of wealth λ_t , he/she will choose $w_{t,i}$ to maximize:

$$E_t \sum_{s=t}^{\infty} (\beta \theta_w)^{s-t} [\lambda_s w_{t,i} n_{s,i} - v(n_{s,i})] \quad (6)$$

subject to labor demand on the part of firms. Wage rigidity is important in the model to match the volatility of the ratio of labor hours to output. In absence of wage rigidities, labor hours in the high-elasticity sector would be too volatile. The presence of wage rigidities makes it harder to move labor between the two sectors.

A.3. Firms

Firm j from sector k hires labor services to produce its output using a linear constant returns to scale technology:

$$Y_{t,k,j} = A_t n_{t,k,j} \quad k \in \{1, 2\} \quad (7)$$

⁴It follows that demand for labor of type i on the part of wage-taking firms is given by: $n_{t,i} = N_t \left(\frac{w_{i,t}}{W_t} \right)^{-\eta_w}$, where W_t is aggregate average wage.

where $n_{t,k,j}$ is an aggregate of the different types of labor supplied by households and hired by firm j in sector k . A_t is aggregate productivity, it follows an AR(1). The firms in this economy face a nominal rigidity, which I model using the standard Calvo (1983) time-dependent price changes: in each period, firms receive an opportunity to change their prices at no cost with probability $(1 - \theta_k)$, but otherwise price changes are infinitely costly.⁵ The firms' objective is to maximize the expected real present value of a dividend flow, $E_t[\sum_{t=0}^{\infty} m_{0,t}d_{t,k,j}]$, where $d_{t,k,j}$ denotes the real dividend and $m_{0,t}$ is the stochastic discount factor. Given the monopolistically competitive product markets, firms' maximization problem is subject to a demand constraint. Formally, firms solve:

$$\max_{P_{t,k,j}^*} E_t \left[\sum_{t=0}^{\infty} (\beta\theta_k)^t m_{0,t} [P_{t,k,j}^* Y_{t,k,j} - w_t n_{t,k,j}] \right] \quad (8)$$

subject to:

$$Y_{t,k,j} = \left(\frac{P_{t,k,j}^*}{P_{k,t}} \right)^{-\eta_k} Y_{t,k} (\omega_k)^{-1} \quad (9)$$

$$Y_{t,k,j} = A_t n_{t,k,j} \quad (10)$$

I close the model by assuming that a monetary authority sets the one-period nominal interest rate $r_t \equiv \log(R_t)$ according to a Taylor (1993)-type policy rule:

$$r_t = \phi_{\pi} \pi_t + \phi_y \log \frac{Y_t}{Y_{t-1}} + \log \left(\frac{1}{\beta} \right) + \epsilon_t^r \quad (11)$$

where π_t is the inflation level, β is the impatience level of households and ϵ_t^r is a monetary policy shock. I do not explicitly model a zero lower bound in the model. In the context of the zero lower bound, the monetary policy shock should be interpreted as forward guidance shocks.

A.4. Calibration

I calibrate the model at quarterly frequency in order to match standard macroeconomic and asset pricing moments in the data. Section AIII and Table A1 from the appendix summarize the calibration choices, as well as the moments I target in the data. Below, I briefly outline the calibration.

⁵Up to a first order approximation modelling nominal rigidities using a fixed menu cost of changing prices or the Calvo (1983) method yields the same results (e.g. Sims and Wolff (2017))

Household Parameters

I set the household impatience level to 0.99, which implies a risk-free rate of 1%. The habit adjustment parameter, v , is 0.66, as estimated by Galí, Smets, and Wouters (2012). This value helps to match the level of the equity premium. The risk aversion, γ , is set to 10, a value that is commonly used in the asset pricing literature (Kung (2015) and Bansal, Kiku, and Yaron (2010)). I set ϕ_l , the weight on disutility of labor, so that steady-state labor hours are around 1. I set σ , the inverse Frisch labor supply elasticity, to 1 as in Rabanal and Rubio-Ramírez (2005). The elasticity of substitution between labor types η_w is set to 21 and the degree of wage stickiness θ_w is set to 0.64, which implies an average labour contract duration of 2.8 quarters, consistent with the evidence of Christiano, Eichenbaum, and Evans (2005).

Firms' Parameters

The elasticity of demand between the two sectors is assumed to be different. I set η_1 to 3, which implies a steady state markup of 50%, and η_2 to 13, which implies a steady state markup of 8%, which are respectively the first and last deciles of the markup distribution estimated by Epifani and Gancia (2011). These markups imply that firms in sector 1 have more monopoly power. I calibrate the correlation between demand elasticity and the degree of price stickiness to zero. This is empirically plausible, as: (i) I estimate a statistically zero correlation between demand elasticity and frequency of price adjustment at both the firm level and the industry level and, (ii) portfolios sorted on demand elasticity do not exhibit any difference in terms of frequency of price adjustment. The empirical section of the paper will provide evidence on the relation between the degree of nominal rigidities and demand elasticity. Thus, I set the degree of price stickiness θ_k to 0.77, which implies an average duration of price contracts of 4.49 quarters, a value that is consistent with the degree of price stickiness in my data and that is in line with the estimates of Rabanal and Rubio-Ramírez (2005), but slightly below the estimate of 0.92 from Christiano, Eichenbaum, and Evans (2005). At the end of this section, I analyze the sensitivity of the results to changes in these elasticity parameters as well as the remaining parameters. This will also allow me to pinpoint the exact forces behind the main result.

Other Parameters

Finally, the coefficients on the Taylor Rule are set to $\phi_\pi = 1.17$, to match the standard

deviation of inflation, and $\phi_y = 0.6$, as in Olivei and Tenreyro (2007). Technology and monetary policy shocks follow an AR(1). I set the coefficients on the auto-regressive processes to be the same as in Weber (2015).

I solve the model using second-order perturbation methods around the deterministic steady-state and simulate the model for 500 firms per sector and 500 periods. I follow the approach of Gorodnichenko and Weber (2016) and, for each time-period t and firm j in sector k , I compute their cum-div value $V(P_{t,k,j})_{cum-div}$ and ex-div value $V(P_{t,k,j})_{ex-div} = V(P_{t,k,j})_{cum-div} - d_{t,k,j}$, and use these two values to compute implied net returns of firms.

A.5. Model Fit, Mechanism and Implications

I now investigate the implications of heterogeneity in demand elasticities for firm fundamentals and asset prices, as well as the mechanism that drives the results. The baseline model generates reasonable macroeconomic moments (see Table A1 in the appendix). The annual volatility of consumption growth is 1.1% and the first order autocorrelation of consumption growth is 0.24. The model generates a volatility of inflation of 0.028 and a first order autocorrelation of 0.706. All of these values are in line with the data. By construction, the baseline calibration allows me to match the degree of nominal rigidities in my sample as well as the steady state markups of the low- and high-elasticity sectors. Finally, the model generates an equity premium of 0.05 and a Sharpe ratio of 0.31: both slightly lower than their data counterparts. Finally, the average price-dividend ratio and the volatility of the price-dividend ratio are 30.1 and 16% respectively.

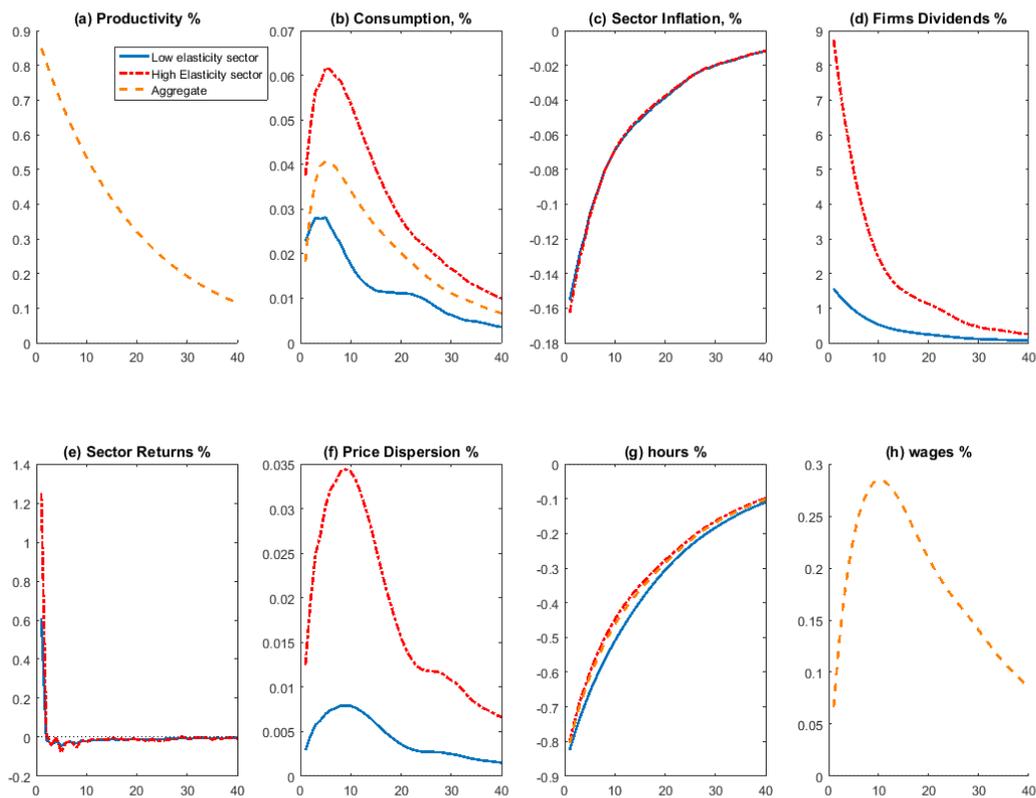
Mechanism

To understand the main forces at work, Figure 1 plots the impulse response functions of the model key variables to a one-standard deviation productivity shock $\log(A_t)$.

The first panel of the figure plots the log-level of productivity, which increases one standard deviation (0.85%) and then slowly decays to its steady state level. Firms are now more productive and therefore the level of output is increased in both sectors (Panel (b)). However, the output response is higher for firms that face a higher elasticity of demand. The reason for this is straightforward. Marginal costs have gone down for firms in both sectors. Firms' optimal reset price is a markup over marginal costs. This markup is lower for firms that face a bigger elasticity of demand and therefore, if allowed to adjust prices, these firms reduce prices more (Panel (c)). This means that the relative prices of goods in the sector with higher elasticity

versus sectors with lower elasticity have now decreased, so demand is relatively higher.

Figure 1: Impulse Response Function to a One Standard Deviation Positive Productivity Shock



y-axis shows % deviations from steady state

The sluggish increase in consumption is due to habits in utility. If there were no habits, consumption would jump and then steadily decrease. The presence of habits leads agents to smooth changes in consumption, which yields the hump-shaped pattern seen in Panel (b). Firms in the model have no savings technology and must distribute all profits as dividends. The higher production level in the sector with higher elasticity leads this sector to distribute more dividends (Panel (d)). The differences in covariance between consumption and dividends of each sector are important for explaining the heterogeneity in the risk of firms. Firms in

the high-elasticity sector pay more when marginal utility of wealth is lower and are therefore riskier. Panel (f) plots the price dispersion in each sector, $pd_{t,k}$ with $k \in \{1, 2\}$. I define price dispersion as the average price of each firm divided by the sector price index:

$$pd_{t,k} = \int_0^1 \left(\frac{P_{t,k,j}}{P_{t,k}} \right)^{-\eta_k} dj \quad (12)$$

Price dispersion in the sector facing a higher (lower) demand elasticity increases more (less). The more elastic sector wants to decrease prices relatively more, so that the staggered mechanism of price setting makes them on average further from the optimal price. This makes this sector riskier than the sector with a more inelastic demand. To demonstrate this, consider a claim over the dividend next period. Denote $d_{t+1,k}$ the aggregate dividend of sector k and $V_{t,k}^1$ the value of the claim over that dividend. Assume that the log-pricing kernel and asset log-returns at the sector level follow normal distributions (this is similar to the expositional assumption made by Li and Palomino (2014)). The return spread of the one-period dividend claim between sectors with high (H) and low (L) elasticities can be written as:

$$E_t[r_{t+1}^H - r_{t+1}^L] = -cov_t(m_{t,t+1}, d_{H,t+1} - d_{L,t+1}) \quad (13)$$

where r_{t+1}^k is the log-return on a claim over next period dividends from sector k , $m_{t,t+1}$ is the log stochastic discount factor and $d_{t+1,k}$ are the log-dividends. Using a log linear approximation, equation (13) can be written as:

$$\begin{aligned} E_t[r_{t+1}^H - r_{t+1}^L] &= -[(1 - \eta)cov_t(m_{t,t+1}, p_{H,t+1} - p_{L,t+1}) \\ &+ (\eta_H - 1)cov_t(m_{t,t+1}, \log(\mu_{H,t+1})) \\ &- (\eta_L - 1)cov_t(m_{t,t+1}, \log(\mu_{L,t+1}))] \end{aligned} \quad (14)$$

where $\mu_{k,t+1}$ is the average markup of prices over marginal costs. The first factor in this return decomposition implies a positive spread between the two returns, as the covariance between marginal utility and the price of dividend claims is larger for the sector with high elasticity (dividends in this sector respond more to shocks, as seen in Panel (b) of Figure 1). However, most of the spread comes from the covariance between marginal utility and markups. Markups are mainly driven by price dispersion, which is larger for the sector with high elasticity.

Furthermore, the covariance between marginal utility and markups is scaled by the demand elasticity η_k , which further enhances the effect.

The mechanism to generate a spread in returns through heterogeneity in demand elasticities is not isomorphic to that in a model with different frequencies of price adjustment, as in Carvalho (2006) and Weber (2015). In their models, the inefficiency comes from *how long* firms stay away from their optimal price, whereas in my model, the friction is related to *how far* posted prices are from optimal prices. This has different implications for firms' markups, which I discuss below.

Implications and Sensitivity to Calibration

With a fair understanding of the model, I will now describe its implications, which can be directly tested in the data. Table 1 reports annualized mean excess returns over the risk-free rate of portfolios of firms in the low-elasticity sector and in the high-elasticity sector, the market equity premium and the Sharpe ratio. The first row of the table reports the results for the baseline calibration.

Firms in the high-elasticity sector are riskier and earn an average excess return of 8.7%, whereas firms in the low-elasticity sector are safer and therefore earn an average excess return of 1.1%. The return spread between the two sectors is 7.6%. The overall market return excess return is 4.74% with a Sharpe ratio of 0.31. Returns in the high- (low-) elasticity sector are also more (less) volatile with a standard deviation of 0.2 (0.15). In the model, the CAPM holds weakly. The portfolio of firms in the high-elasticity sector has a CAPM beta of 1.2 and the portfolio of firms in the low-elasticity sector has a CAPM beta of 0.8. Although CAPM betas line up with expected returns, the difference in betas between the low-elasticity portfolio and the high-elasticity portfolio is not enough to explain the return spreads, this is due to the presence of risk factors not captured by the CAPM, such as habits and nominal rigidities. This is something that I also test in the data.

Table 1: Model Implied Stock Returns

This table reports annualized asset pricing moments for the model described in Section III.A. The second and third columns report average excess returns for stocks in the low-elasticity and high-elasticity sectors respectively. The fourth column reports the spread in returns between the two sectors. The fifth and sixth columns report the average excess market return and the Sharpe Ratio. The first line of the table reports these moments for the baseline calibration and the remaining lines report sensitivity to several model parameters.

	Low Elasticity	High Elasticity	Spread in Returns	Market Return	Sharpe Ratio
Benchmark	1.10%	8.70%	7.60%	4.74%	0.311
No Productivity Shocks	0.71%	3.40%	2.69%	2.05%	0.176
No Monetary Policy Shocks	0.04%	0.44%	0.40%	0.24%	0.083
No Sticky Prices	1.09%	1.09%	0.00%	1.09%	0.069
Higher Stance on inflation, $\phi_\pi = 1.24$	0.82%	1.71%	0.89%	1.27%	0.090
Higher Stance on consumption, $\phi_c = 0.7$	0.77%	9.03%	8.27%	4.90%	0.315
Higher Dispersion of Elasticities $\eta_2 = 14$	0.78%	11.56%	10.74%	6.17%	0.423
Higher Elast. of Substitution Across Sectors $\eta = 3$	0.76%	10.27%	9.52%	5.52%	0.326

The remaining rows of the table describe sensitivity to parameter choices. Rows 2 and 3 of Table 1 report the results for the baseline model without productivity shocks and monetary policy shocks, respectively. Compared with the baseline model, monetary policy shocks are quantitatively more important for both the spread in returns and the market equity premium. This is in line with the findings of Weber (2015) and De Paoli, Scott, and Weeken (2010). In particular De Paoli, Scott, and Weeken (2010) argue that nominal rigidities only enhance risk premia if they are coupled with monetary policy shocks. Rows 5 and 6 report the results if the monetary authority takes a stronger stance on either inflation or output. A stronger stance on inflation dampens the equity return and strongly decreases the spread among the two sectors. A stronger stance on output has limited effects on the outcomes of the model. Finally, the last two rows report the effects of a higher dispersion in elasticities and a higher elasticity of substitution across the two sectors. As expected, both an increase in the gap of elasticities between the two sectors and an increase in the substitutability of goods across the two sectors increase the spread in returns among the sectors.

Firms with a higher degree of monopoly power have higher markups.⁶ This is in line with what a simple reduced form monopoly power model would deliver: higher monopoly power (lower elasticity) implies that firms charge higher markups over marginal costs and consequently have higher monopoly rents. This is an important prediction of this model, as it allows me to disentangle the cross-sectional implications for returns of different demand elasticities from other cross-sectional mechanisms proposed in the literature. In particular, Weber (2015) and Gorodnichenko and Weber (2016) argue that firms with a higher degree of price stickiness earn higher returns.

In order to disentangle the heterogeneous sticky prices mechanism from the heterogeneous elasticities mechanism, I solve a two-sector version of Weber’s (2015) model. Table 2 reports results for markups and returns of my baseline model, where firms are heterogeneous in the elasticity they face (Panel A), as well as the results of Weber’s (2015) model, where firms are heterogeneous in the degree of price adjustment (Panel B).

In my model, as described above, the more elastic sector earns higher returns and has lower markups (Panel A). On the other hand, and in line with the results of Weber (2015), sectors with a higher degree of price stickiness earn higher expected returns (Panel B). In Weber’s

⁶In the absence of nominal rigidities firms would set the optimal price to be a markup over marginal costs, i.e.: $p^* = \frac{\eta}{\eta-1} \frac{A_t}{w_t}$. Therefore, firms with more inelastic demands have higher market power and set higher markups over marginal costs (lower η).

Table 2: Heterogeneity in Sticky Prices Versus Heterogeneity in Elasticities of Demand in the Model

Panel A: Heterogeneity in Elasticities only		
	Returns	Markups
Low Elasticity	1.10%	1.531
High Elasticity	8.70%	1.138

Panel B: Heterogeneity in Price Stickiness only		
	Returns	Markups
Low Stickiness	0.72%	1.092
High Stickiness	6.34%	1.149

This table reports returns and markups for the baseline model described in Section III.A (Panel A) and a two-sector version of the model developed by Weber (2015) (Panel B). The calibration of Weber (2015) is identical to the calibration of the baseline model with the exception of the stickiness parameters θ_1 and θ_2 which are set to 0.3 and 0.9, respectively, and the elasticity parameters η_1 and η_2 , which are identical and equal to 8 (the same value as in Weber (2015)).

(2015) model, in the deterministic steady state, firms in either sector have the same markups. However, in the stochastic steady state, firms in the high price stickiness sector fear selling at a loss more and charge slightly higher markups for precautionary motives. This implication is the opposite of the one in the my model. When firms face different demand elasticities, firms with low markups earn higher returns, whereas when the heterogeneity is in the nominal rigidity, firms with low markups earn lower returns. This is an implication that I also take to the data.

IV. Data and Identification

A. Data

The data used in this study are from **Keepa.com** (hereafter **Keepa**). **Keepa** is a private online price tracker that tracks over 250 million products sold on Amazon USA, UK, France, Japan, Canada, China, Italy, Spain, Mexico and Brazil. It started operating in January 2011. In this section, I will describe the **Keepa** data and leave the details regarding the data gathering process for Appendix AIV.

The **Keepa** database stores several fixed attributes of the products it tracks: the product name, category node⁷, Universal Product Codes (UPC), International Article Numbers (EAN),

⁷e.g. Home Care & Cleaning

product brand, label, model, color and size, among other attributes. **Keepa** also stores several time series for time-varying quantities such as prices and a proxy for quantity called “sales rank” (an ordinal ranking describing the quantity sold of a product).

The frequency of update and the sales rank indicator are two of the distinctive features of this database compared to other product price datasets (such as the Nielsen dataset or the BLS dataset). First, the **Keepa** database is updated several times a day. For most products, it is updated once every hour. In contrast, the Nielsen Dataset and the BLS dataset are updated weekly and monthly (respectively). Even most scraped datasets used in the sticky prices literature (e.g. Cavallo (2016)) have at most daily update frequencies. The second distinguishing feature of this dataset is the sales rank, which is a very good proxy for quantity sold (more on this relationship below). Neither the BLS datasets nor the scraped datasets have quantity information. The Nielsen scanner datasets, which do have information on quantities, are an exception. The **Keepa** dataset fills this void in the micro-level product datasets, having both quantity information and high update frequency.

Keepa is one of the largest Amazon online price trackers available. It provides free access to interactive product price and sales rank charts, and it charges a fee for a user to access and download the data. It also allows users to monitor any Amazon product and receive alerts once the price drops. Most users of this database are concerned with product sales rank or product prices. Sellers looking for new product markets to enter rely on the sales rank of a product, i.e., how frequently the product is sold, to decide whether or not they should enter a market, while consumers looking to take advantage of product deals and price drops rely on price alerts to make purchase decisions. Therefore, for both sellers and buyers, the high update frequency of **Keepa** is crucial. This is important to me, as the high-frequency nature of these data will be extremely helpful for identifying demand elasticities.

Keepa can potentially track all Amazon products (except for e-books). Once a product enters the database, tracking begins and will continue indefinitely. Even if the product is no longer sold on Amazon, it still appears in the database, albeit with no associated price. **Keepa**’s database is constantly growing as new products are added on a daily basis. At any given point in time, **Keepa** ensures that it tracks all best-sellers within each Amazon product category and if anyone searches for a product that is not being tracked, then **Keepa** starts tracking it.

Unlike standard datasets, it is not possible to download the full **Keepa** database at once due to its large size. Each Amazon product has a unique identifier called the Amazon Standard

Identification Number (hereafter, **Asin**) which can be used to obtain the data for a specific product. The challenge here is to identify products sold by Amazon that are produced by publicly traded firms. I manage to get data on 278 thousand products sold by public firms. I use a two-tier approach to identify public firms that sell products on Amazon. First, I rely on prior literature that uses Amazon data for public firms, namely Huang (2016), who identified 246 public firms that sell their products on Amazon. Out of those firms, 59 were taken private or filed for bankruptcy before 2011. These are therefore dropped from my sample. Second, I take the full list of brands sold by Amazon under each product category, manually match each brand to a corporation and check if the company is publicly traded. This increases the sample from 187 to 250 firms. For each firm in the sample, I manually search for each brand it sells on Amazon and use a web-crawling algorithm to retrieve the **Asins** of their products, which I can then “insert” into the **Keepa** database to download the data (see details in Appendix AIV).

Overall, the sample is comprised of approximately 278 thousand products, with an average of a thousand products per firm. The data range from January 2011 to March 2017.⁸ This a fairly large database in terms of both the time series and the cross-section (e.g. the Cavallo (2016) scraped USA sample contains 170 thousand products over the course of two years). The sample contains firms belonging to nine of the thirteen Fama-French industries - Manufacturing, Health, Shops, Business Equipment, Durables, Telecom, Non-Durables, Chemicals and Others. The three industries that are not represented in my sample are Energy, Utilities and Finance.

I aggregate the product-level data to the firm level and merge it with the standard equity prices database (CRSP) and firm fundamentals database (Compustat). Panel A of Table 3 reports the main characteristics of the firms in the sample. On average, each firm sells around a thousand products on Amazon. The median number of products per firm is 298, which reflects the fact that a few firms sell far more products than others (e.g., clothing firms sell many products whereas manufacturing firms sell fewer products). The average market capitalization of the firms in the sample is US\$26 billion, as most of the firms that have their products sold by Amazon are large. They also have high average revenues of US\$3 billion and high average EBITDA of US\$700 million. This reflects that fact that most of these firms are not only large, but also mature and well established. Panel B reports the same metrics for the overall Compustat sample over the same time period. As expected, the firms included in my sample are larger in all the metrics (market capitalization, revenue and EBITDA), and have

⁸The Sales Rank data starts in February 2015 as **Keepa** only started keeping track of Sales Rank at this date.

lower book-to-market capitalization. Therefore, while this is not a representative sample of US public firms, the granularity of the product data might still allow us to draw important empirical conclusions regarding product level dynamics and asset prices.

Table 3: Summary Statistics *Keepa* Sample and Compustat Sample

Panel A: Summary Statistics					
	Mean	Std. Deviation	25th	Median	75th
Number of Products	1109.1	2593.3	223.0	298.0	612.0
Frequency of Price adjustment (months)	4.93	1.82	3.92	4.50	5.39
Market Capitalization (M\$)	26969.8	61048.6	1408.0	5576.3	19414.0
Book-to-Market	4.2	8.2	1.8	2.9	4.6
Revenues (M\$)	3450.2	6239.6	314.1	945.1	3456.5
EBITDA (M\$)	695.2	1643.8	34.8	147.2	545.8

Panel B: Full Compustat Sample					
	Mean	Std. Deviation	25th	Median	75th
Market Capitalization (M\$)	2745.0	14096.4	25.9	135.1	823.3
Book-to-Market	11.0	2771.9	0.5	1.3	2.8
Revenues (M\$)	743.4	4023.6	0.7	24.5	211.7
EBITDA (M\$)	156.8	1087.4	-0.5	2.3	32.6

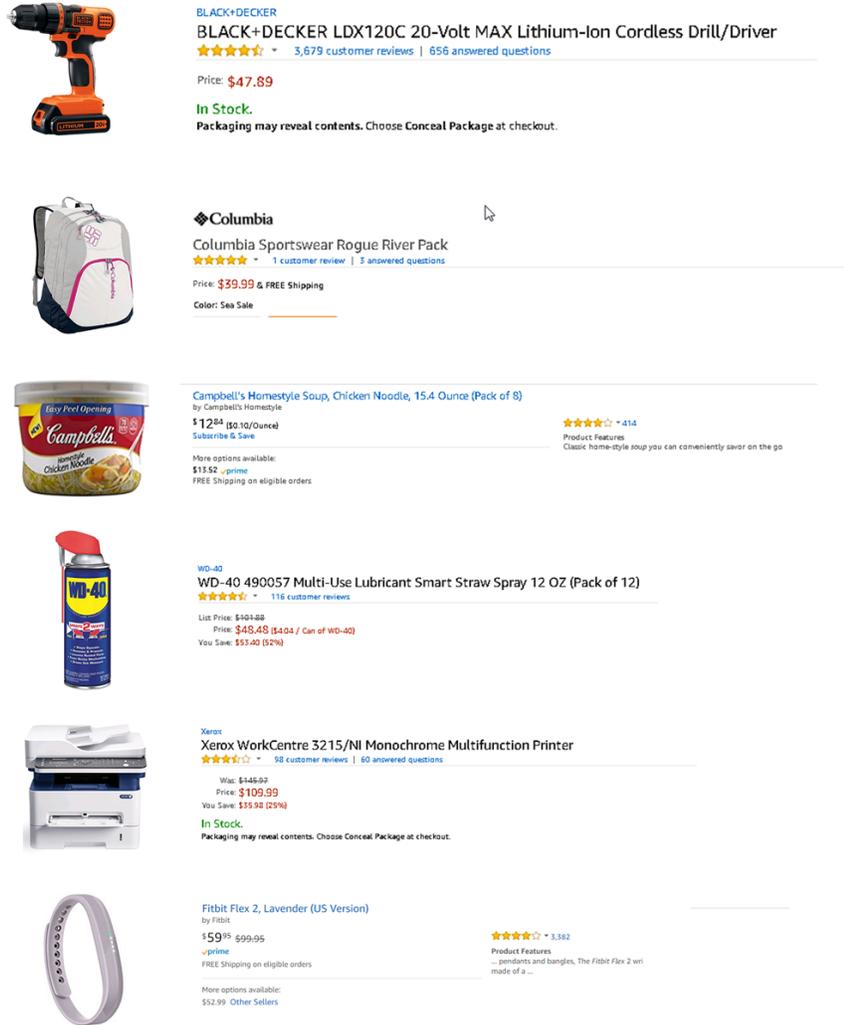
This table reports summary statistics for the firms included in the *Keepa* sample (Panel A) and the full Compustat sample (Panel B). The first row of Panel A reports the statistics for the number of products per firm. The second row reports the frequency of price adjustment. The third to sixth rows report average market capitalization, book-to-market, revenues and EBITDA per firm. Panel B reports the same financial statistics for the entire Compustat sample.

B. Amazon Prices and Sales

I use some examples to explain the nature of the data. Figure 2 plots screenshots of 6 products in the sample, taken from *Amazon.com*. As the figure illustrates, Amazon sells a vast range of products: drills, sportswear, food, sprays, copying devices and, sports watches, among others. At any one point in time, Amazon only shows the current price and sales rank of each of these products. The advantage of using a database such as *Keepa* is that it includes historical time series of prices and sales rank for each product.

Consider the last product in Figure 2, the Fitbit Flex 2. This wristband is produced by Fitbit, Inc., a US public firm traded on the New York Stock Exchange (NYSE). Panel A of

Figure 2: Examples of Products in the Dataset



This picture plots six examples of products in the Keepa dataset. The figures and prices are screenshots taken from the website *amazon.com*.

Figure 3 shows the time series of the original price p_t for this product. The wristband was introduced to the market in September 2016, with a retail price on Amazon of 80.5 US dollars. There are a couple of stylized facts from the literature on sticky prices that can be seen in the figure. First, prices have different low- and high-frequency dynamics. At a lower frequency, prices are quite sticky and move around very little. In this specific example, the price is usually either 80.50 or 60 US dollars. At a higher frequency, prices experience a great deal

more movement, with shorter price spells. For instance, in December-2017 the price briefly dropped from 80.5 US dollars to 75 US dollars for a day, before reverting to 80.5 US dollars. Second, following a temporary price change, the price often reverts to the original nominal price. Midrigan (2011) and Kehoe and Midrigan (2015) document these two stylized facts and conclude that the standard Calvo and Menu Cost New-Keynesian models cannot generate both patterns: (i) very sticky (flexible) prices at low (high) frequencies, and (ii) that after a temporary price change, the nominal price often returns to the nominal pre-existing price. Although they use CPI datasets to document these, the same patterns appear in Amazon prices.

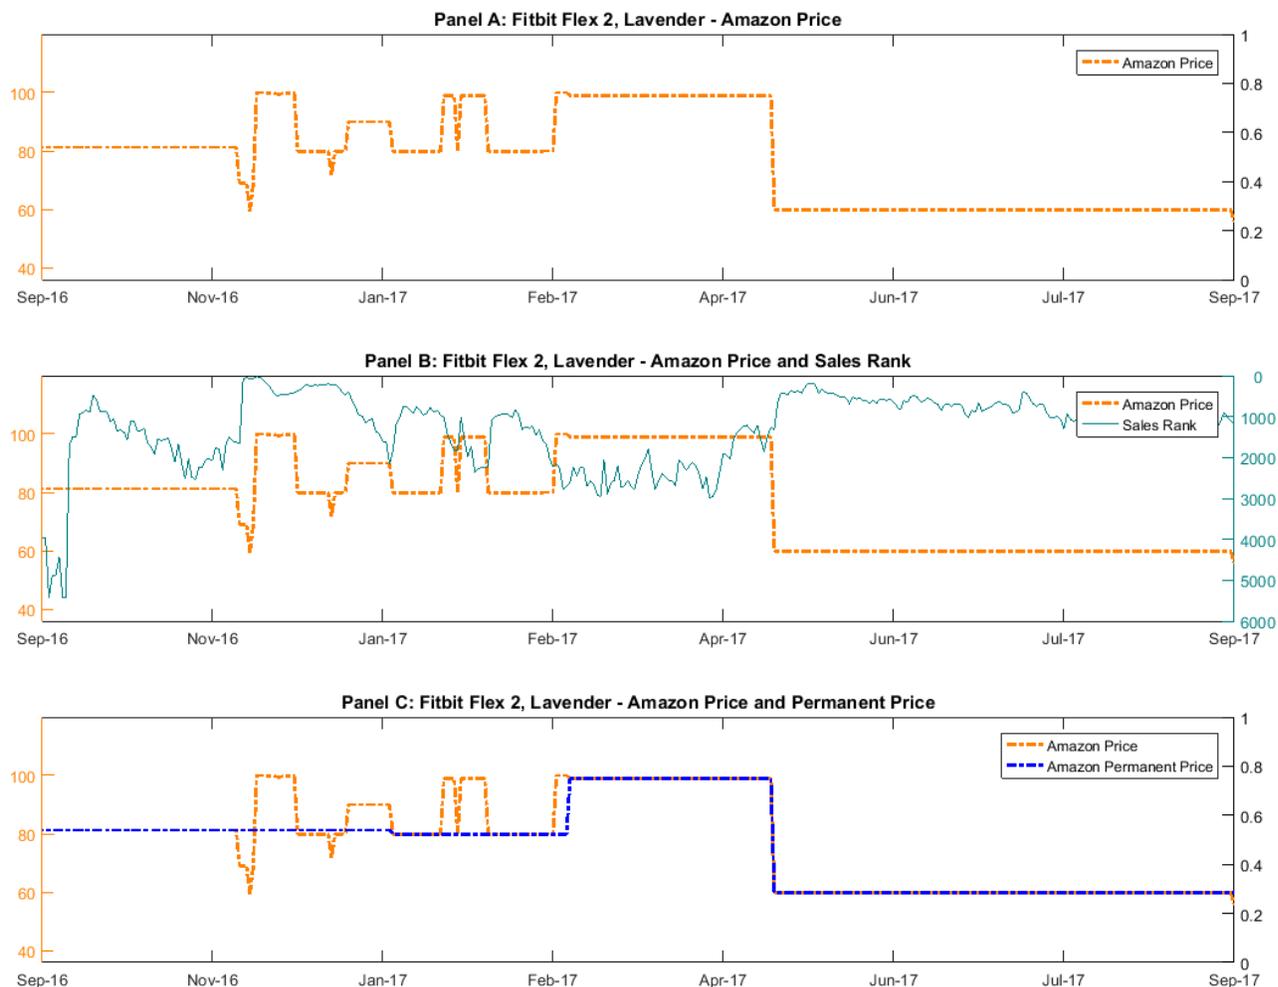
Panel B plots the Amazon price (in the left y -axis) and the Amazon Sales Rank (in the right y -axis). The Amazon Sales Rank is an ordinal ranking describing the quantity sold of a product within its product category. If a product is sold, the ranking increases and becomes closer to 1 (which would be the ranking of the best-selling product). If the product is not sold and other products are, the ranking starts drifting down from 1. Although prices are quite sticky, the figure shows that there is much more variation in Sales Rank, because products move up and down the ranking as they get sold more or less often than products in the same category. For instance, on Friday, November 25th, 2016, there is a spike down in the price of the Fitbit watch. This was the USA’s so-called Black Friday: the shopping day after Thanksgiving, when many retailers offer promotional sales. It is clear that there is a sharp spike up in the Sales Rank when Black Friday arrives, which then slowly reverts to its unconditional mean. All results in the paper are robust to the exclusion from the sample of weeks surrounding event days such as Black Friday, Christmas, New Year, Valentine’s day and Father’s/Mother’s day.

Several marketing studies look into the properties of the Amazon Sales Rank and conclude that there is an extremely tight relationship between Sales Rank and quantities sold (e.g. Chevalier and Mayzlin (2006), Brynjolfsson, Hu, and Smith (2006) and Chevalier and Goolsbee (2003)). In particular Chevalier and Goolsbee (2003) have data on both Amazon quantities sold and Sales Rank for 20,000 Amazon products, and they conjecture that a standard distributional assumption for rank data is a Pareto distribution (i.e., a power law). The Pareto distribution conjecture implies that one can translate the sales rank to quantity sold using the following log-relation:

$$\log(\text{SalesRank}) = c - \theta \log(\text{QuantitySold}) \tag{15}$$

where c and θ are a function of parameters of the power law distribution. They estimate

Figure 3: Amazon Price, Sales Rank and Permanent Price



This figure plots price and sales rank time series for the Fitbit Flex 2, a sports wristband sold by Fitbit, Inc. The top panel plots the end-of-day Amazon price of the watch. The middle panel plots the end-of-day Amazon price on the left y -axis and the sales rank on the right y -axis. The last panel plots the Amazon price and the permanent price computed using the Midrigan (2011) algorithm. The data is taken from *Keepa*.

equation (15) and find very robust estimates of θ and c across different samples, with R^2 's exceeding 0.95. This implies that sales rank movements do indeed have a strong relation with

quantities sold and that the log change in Sales Rank translates almost one-to-one with log change in quantities sold (just take first differences of equation (15)).

Therefore, I interpret movements in the Amazon Sales Rank as changes in quantities and do several tests below to confirm this hypothesis. During the remainder of the paper, I will use the terms *Sales Rank* and *quantity sold* interchangeably, and I will only refer to the difference between the concepts when appropriate. This approximation matters for the quantitative level of the estimated elasticities, but not for ranking products and firms according to the demand elasticity they face. In this paper, I am only concerned about how firms rank in terms of demand elasticity and not with the level of the estimates.

The sticky prices literature places great emphasis on the distinction between high-frequency and low-frequency price movements, and most frequencies of price adjustment statistics are computed using low-frequency price movements. It is observationally difficult to disentangle when a price movement is temporary versus when it is permanent. I employ the same algorithm as Midrigan (2011) and Kehoe and Midrigan (2015), which is based on the idea that a price is *regular* or *permanent* if the retailer frequently charges it in a window surrounding a particular week, provided the modal price is used sufficiently often. It is beyond the scope of this paper to describe this algorithm in detail; readers should refer to Midrigan’s (2011) online appendix for further details.

Panel C of Figure 3 plots the Amazon price, p_t , in the left y -axis and the regular price, pR_t in the right y -axis. As expected, regular price changes occur much less frequently than temporary price changes. The regular price is much stickier than the actual price.

In the literature, it is common to compute price stickiness metrics excluding changes in prices due to sales (e.g. Cavallo (2016), Nakamura and Steinsson (2008), Midrigan (2011) and Gorodnichenko and Weber (2016)). Therefore, I apply Midrigan’s (2011) algorithm to each product in the sample, to disentangle temporary from regular price changes, then I use the regular price series to compute frequency of price adjustment (*FPA*) at the firm level. I define frequency of price adjustment in the standard way. First, I obtain the frequency of price adjustment per individual good by calculating the number of price changes over the total valid observations for a particular product. Next, I take *mean* frequency per good at the firm level and then compute implied durations using: $-1/\ln(1 - frequency)$. The second line of Panel A in Table 3 reports the frequency of price adjustment at the firm level. On average, firms adjust their prices every 4.93 months (permanently). This value is comparable to previous estimates

in the literature (e.g. Nakamura and Steinsson (2008) and Bils and Klenow (2004)).

There is, however, significant heterogeneity in the permanent frequency of price adjustment among firms. The interquantile range is 1.47 months and the standard deviation is 1.82 months. Weber (2015) uses this heterogeneity to sort portfolios and argues that firms with stickier prices are riskier, and therefore have higher CAPM betas and higher expected returns when compared to firms with more flexible prices. I test this hypothesis in the online appendix of the paper and find that in the more recent sample, this result does not hold. This is consistent with Weber (2015), which shows that the relationship holds only before 2007. Furthermore, I estimate a relation between the elasticity of demand of a firm and the frequency with which firms alter their prices, and I fail to find such an association.

C. High-Frequency Identification

The main empirical prediction of the model described in Section III is that firms with higher elasticities of demand are riskier and therefore earn higher excess returns in equilibrium. In order to test this hypothesis, I need to estimate demand elasticities. Having both individual product quantity and price data is usually not enough to identify elasticities, due to the standard endogeneity problem: empiricists only observe equilibrium prices and quantities, therefore the slopes that generate the equilibrium outcomes are unknown.

In this section, I make use of the high-frequency *Keepa* data to estimate elasticities. Given the nature of my data, I can measure the level of demand right before the price change, then see how demand moves within a very narrow window after the price change. In reality, firms could be moving prices due to demand shocks or supply shocks, or they could be simply changing markups (as often occurs in sales). Price changes are endogenous, changes in markups or supply side shocks are not a problem for my identification, as they are shifts along the demand curve and allow me to trace demand. However, shifts of the demand curve are problematic for my identification strategy. Thus, for the identification strategy to be valid, I need to ensure that in such a narrow window around the price change, the demand curve is not moving. There are two main reasons for demand for a specific product to move. First, demand can move if there is an overall shift in demand for products in a certain product category. This would be the case for products such as flowers on Valentine’s day, for electronics on Cyber Monday, or for a wide range of other products in holiday periods such as Christmas and New Year. Second, demand for a given product should also move around if competitors are changing their prices. I address

these concerns below, then I detail how I estimate demand elasticities.

Narrow Windows with Stable Demands

To correctly identify demand elasticities, i.e. movements *along* the demand curve in response to price shifts, demand needs to be stable and consumers should not be able to predict that prices are going to change. Demand is likely to shift in the days surrounding holidays and occasions, so I remove from my sample the weeks around Christmas, New Year, Valentine’s, Thanksgiving, Black Friday, Father’s Day and Mother’s Day.⁹ Second, it is important that competitors are not changing prices when a firm changes its price.

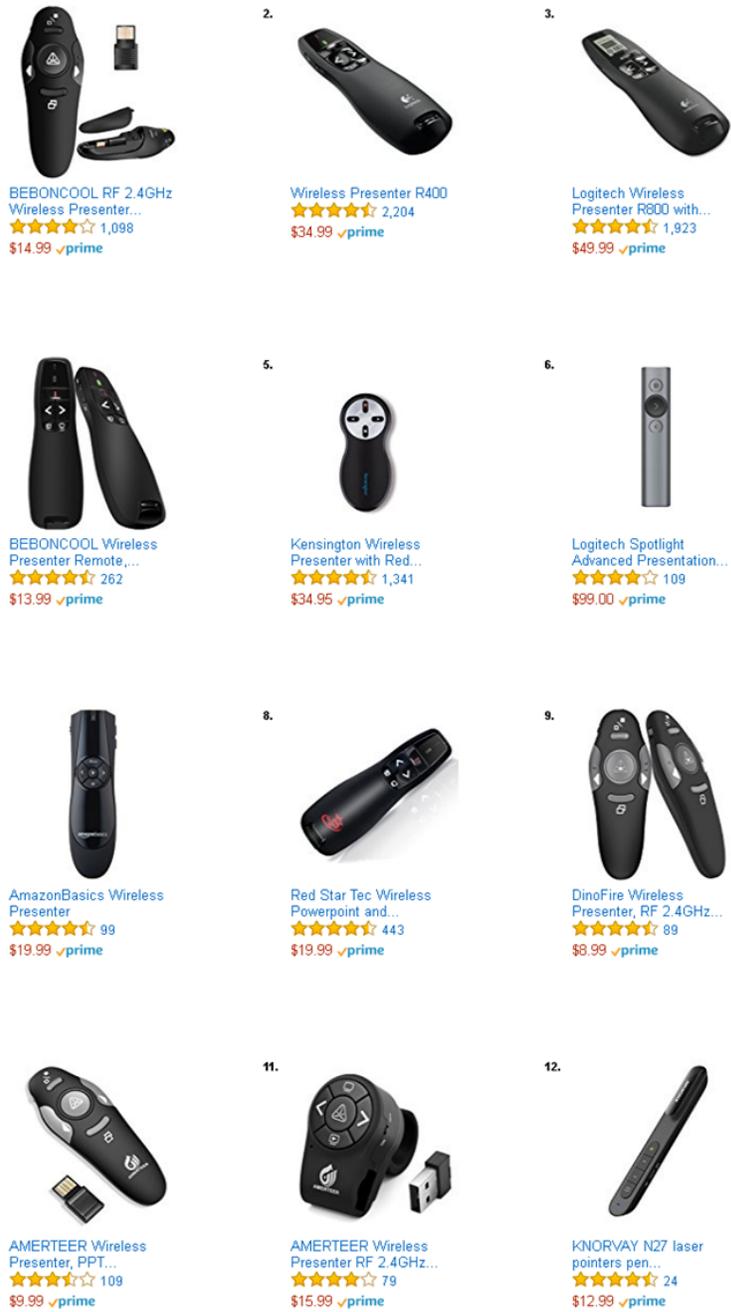
One of the advantages of using the *Keepa* database is that it is straightforward to identify competitor products. For each firm in my sample, I pick a random product that it sells, and extract its narrowest Amazon product category. A given product might belong to multiple product categories. For example, one of the products sold by Logitech Inc, a company in my sample, is a presentation clicker. This product belongs to the Amazon product category *Electronics* and the subcategories *Office Electronics > Presentation Products > Presentation Remotes*, so its narrowest product category is *Presentation Remotes*. Figure 4 illustrates a few products in this subcategory. All of the products are close substitutes, with similar price tags, so in a frictionless market, a change in the price of a product in such a narrow category should trigger a change in the competitor’s demand and consequently a change in the price of the competitor’s product.

For each product, i , in the sample, I take the competitors’ products and see if they are changing their prices when the price of product i changes. Figure 5 plots a histogram with the probability of a change in price by a competitor when product i changes its price. Unconditionally, the overall probability of a competitor’ changing a price when product i changes its price is 2.95% (in a two-sided 12-hour window). Further, I find that 50% of the time, only 2.5% of competitors change their price, and that 80% of the time only 5% of competitors change their price. The average price change of competitors is around 3.2%, which also implies that demand should be fairly stable in the windows where I am measuring elasticity.

Even if price changes are not synchronized at such high frequency, it is interesting to understand whether they are synchronized at lower frequencies. Price changes being correlated at lower frequencies are not a problem for my identification strategy. Nevertheless, it is em-

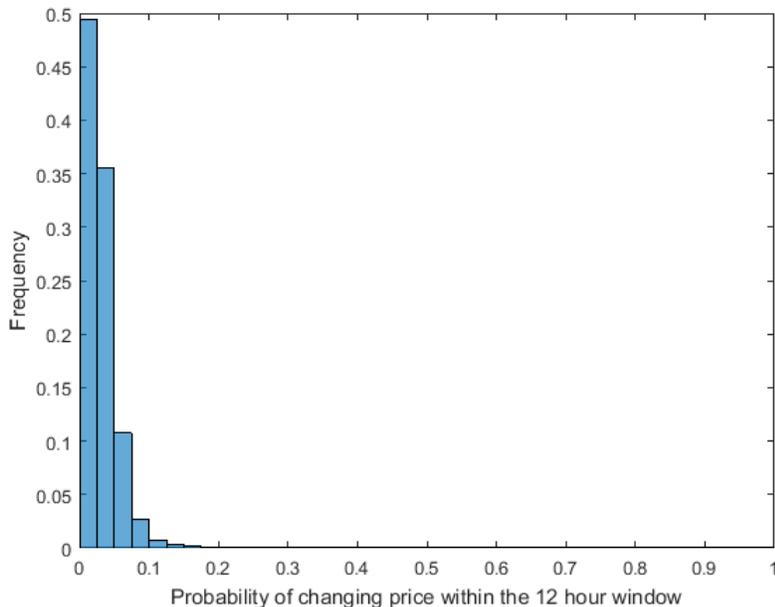
⁹The main results of the paper are not sensitive to using either the full sample or the sample with holidays and occasions included.

Figure 4: Close Substitutes - Example



This figure plots a screenshot of the *amazon.com* *Presentation Remotes* product subcategory.

Figure 5: Probability of a Competitors' Price Change in the 12-hour Windows



This figure plots a histogram with the probabilities of price change for a competitor's product in the 12-hour windows where I am estimating elasticities of demand. The x -axis shows the probability of a competitor's price changing and the y -axis the frequency at which each probability occurs.

pirically relevant to document price synchronization patterns across similar goods, as it may provide new evidence regarding the nature of nominal rigidities.

Formally, I define the synchronization of price changes as the mean share of sellers that change the price for a particular good when another seller of the same good changes its price.¹⁰ Consider the set \mathbb{S}_i of products that are close substitutes in product category i . If $s_{i,t \rightarrow T}^c$ is the number of products in \mathbb{S} that changed their prices between times t and T , then one can define the synchronization rate as:

$$sync_i = \frac{s_{i,t \rightarrow T}^c - 1}{\#\mathbb{S}_i - 1} \quad (16)$$

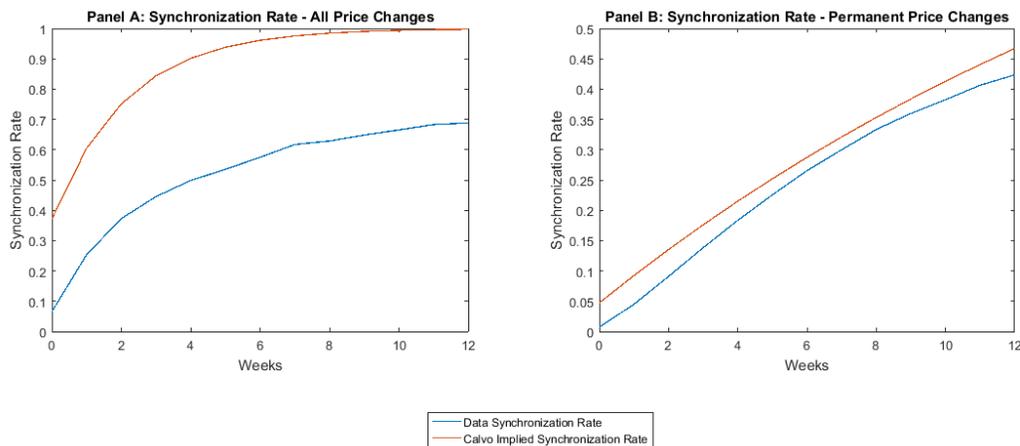
where $\#\mathbb{S}_i$ is the total number of products in subcategory i .

¹⁰This definition is similar to the definition of synchronization used by Gorodnichenko, Sheremirov, and Talavera (2014).

The synchronization rate defined above ranges between one if all products change price between t and T , and zero if price changes are perfectly unsynchronized. To put the synchronization rate defined above into perspective, I compute the Calvo theoretical synchronization rate. This is a useful benchmark: in a Calvo model, each firm is allowed to change the price randomly, therefore firms cannot synchronize price changes. However, given that some price changes coincide in time, the synchronization rate is not zero.

Figure 6 plots the synchronization rate at several weekly horizons.¹¹ The left panel plots the synchronization rate taking into account all price changes (including temporary price changes). The blue line is the synchronization rate implied by the *Keepa* data. As soon as a product changes its price, the probability of a competitor adjusting its prices is only 4%.

Figure 6: Synchronization Rates: Data and Calvo-implied



This figure plots the synchronization across products of price changes at the week- h horizon. The blue lines plot the synchronization rate in the data computed using equation (16). The red line plots the Calvo-implied synchronization rate. We compute the frequency of price adjustment at the subcategory level and use it to backout the Calvo parameter at a weekly frequency. The Calvo-synchronization rate is given by: $1 - (1 - \theta)^{h+1}$, where θ is the median frequency of price adjustment across sectors and h is the horizon in weeks. The left-hand panel plots the synchronization rates without filtering for temporary price changes and the right-hand panel plots the synchronization rate excluding temporary price changes. The permanent price changes are computed using the algorithm of Midrigan (2011).

This implies that firms are slow to react to demand shocks such as price changes and

¹¹It might be the case that sellers do not synchronize price changes immediately, but may be able to do so at lower frequencies.

therefore this should not be an issue for the identification strategy outlined above. The Calvo synchronization rate is depicted by a red line.¹² The Calvo synchronization rate always lies above the synchronization rate of the sellers in my *Keepa* sample, which is in line with the evidence found by Gorodnichenko, Sheremirov, and Talavera (2014). This is a very puzzling finding, given that we usually expect competitors to react fast to changes in prices of products that are close substitutes.

The right-hand panel of the figure plots the same synchronization rates, but excluding the temporary price changes.¹³ Again, the Calvo synchronization rate lies above the data synchronization rate, which again is evidence that firms are slow to react to changes in competitors' prices. Remember that the products are almost perfect substitutes and therefore price responses to a competitor's change in prices should be much quicker. Alternatively, if instead of using the Calvo model as a benchmark I use a menu cost model, then the differences would be even more striking. In fact, in a menu cost model, price changes of substitute goods are perfectly coordinated and the synchronization rate should be one.

I now turn to explaining the demand elasticity estimation in detail.

Demand Elasticity Estimation

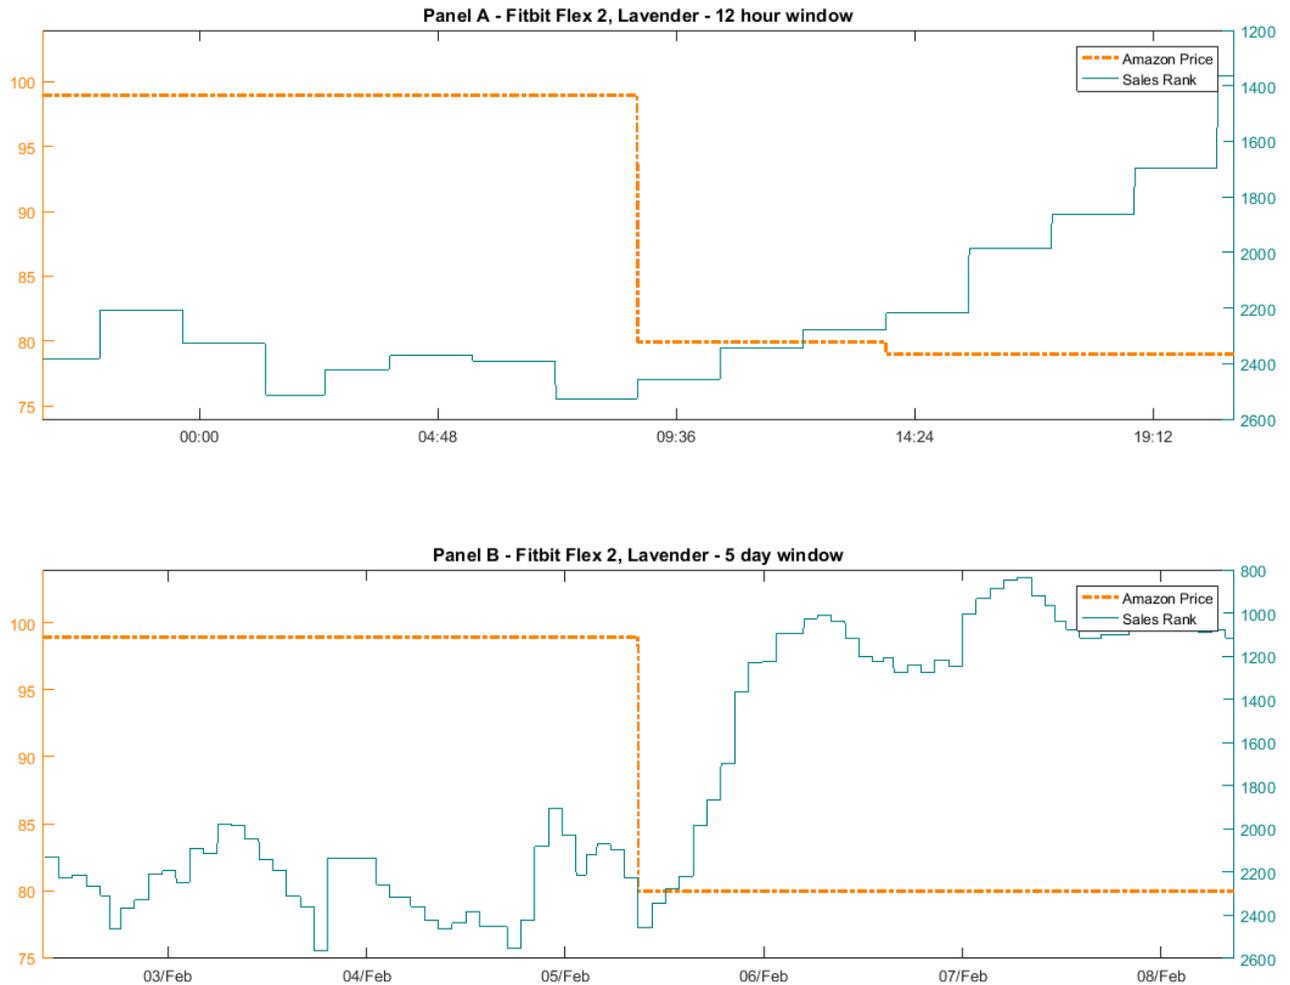
To better understand the identification strategy, I rely again on the example of the Fitbit Flex 2 wristband. On February 2nd, 2017 at 8:15 a.m. UTC -4, Amazon changed the price of this product from 99 to 79.5 US dollars. A decrease in price should lead to an increase in the quantity sold of this product: a movement *along* the demand curve. Panel A of Figure 7 illustrates this effect. The orange dashed line shows the Amazon product price decrease. The green line is the product sales rank, which can be interpreted as quantity sold. The sales rank was on average 2400, meaning that the product was ranking fairly low in comparison with products in its category. At 8:15 am, when Amazon changed its price, the sales rank immediately jumped to 1300, which means that the product sold well in that exact moment. Over the next 12 hours, the sales rank kept increasing towards one. In this narrow window, it is unlikely that events were affecting the demand for this product, other than a simple movement *along* the demand curve.

Panel B of Figure 7 plots the Amazon price and Sales Rank of this same price change, but

¹²Given the evidence from Weber (2015) that different sectors might have different frequencies of price adjustments, I compute the degrees of stickiness at the product category level and average them.

¹³To filter temporary price changes, I use the Midrigan (2011) algorithm.

Figure 7: High-Frequency Amazon Price and Sales Rank



This figure plots high-frequency price and sales rank for the Fitbit Flex 2, in two narrow windows around a price change. The top panel plots the Amazon price (left y -axis) and sales rank (right y -axis) of the product in a 12-hour window around a price change that occurred on February 5th, 2017. The bottom panel plots the same thing in a 5-day window.

in a 5-day window. Demand is quite stable until the price change on February 2nd when it increases sharply until February 7th. After the 7th, demand stabilizes again. This illustrates

that the only shock occurring at the time of the price change is indeed the price shock itself, and no other confounding effects should be at work.

For each price change Δp_t and for each product i of firm f , I calculate within a narrow window of $[t, t + w]$ the following:

$$\epsilon_{t \rightarrow t+w, i, f} = -\frac{\Delta \log SalesRank_{t \rightarrow t+w, i, f}}{\Delta \log p_{t \rightarrow t+w, i, f}} \equiv \frac{\Delta \log Q_{t \rightarrow t+w, i, f}}{\Delta \log p_{t \rightarrow t+w, i, f}} \quad (17)$$

which I interpret as an elasticity of demand for the product. Sales rank moves in the opposite direction to quantity; therefore the sign of the first equality is flipped to make it closer to the standard demand-price elasticity, which theory predicts to be negative. To be clear, I measure quantity demanded right before the price as changed (seconds before) and then again 12 hours after the price change. I use the change in quantity in this 12 hour asymmetric window to compute demanded elasticities. The estimated elasticities and main results are also robust to the choice of an 8-hour or a 16-hour window, or to the use of symmetric windows.

I average elasticities resulting from equation (17) at the product level, then at the firm level, to get an estimate of the average demand elasticity per product and the average elasticity that a firm faces. The averaging across products reduces measurement error and potential noise in the estimates of demand elasticity. Further, asset pricing data only exists at the firm level and not at the individual product level. This allows me to aggregate products at the firm level.

To address any remaining endogeneity concerns in the robustness section of the paper, I re-estimate elasticities using standard instrumental variable methods and replicate the main findings of the paper. This methodology comes at the cost of having to drop some products and firms from the sample due to the lack of appropriate instruments.

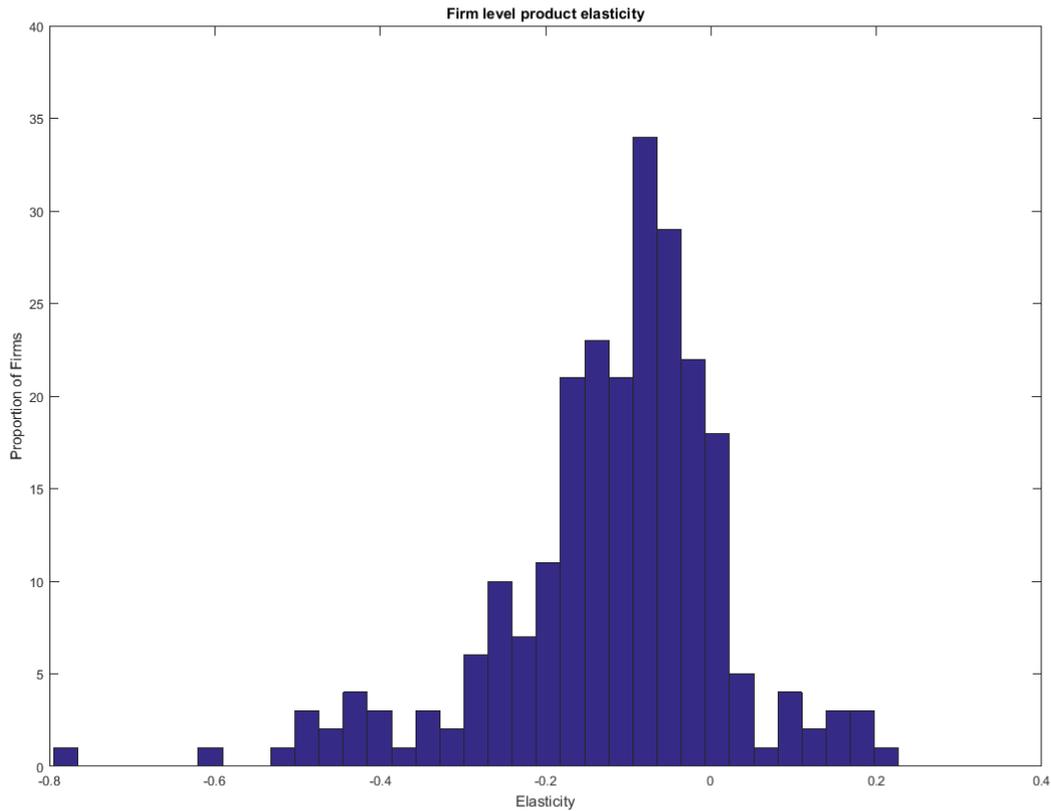
V. Empirical Results

A. Demand Elasticity

Figure 8 plots a histogram of the estimated firm demand elasticities. The average elasticity of demand faced by a firm is -0.124. In line with the predictions of economic theory, most firms in our sample face a negative elasticity, with significant heterogeneity across firms: the standard deviation of firms' elasticities is equal to 0.185. There are a few firms in our sample with a positive estimate for elasticity. This is likely to be due to noise in the data: conditional on having a positive estimate for elasticity, only 0.8% of firms have estimates that are statistically

different from zero at the 5% significance level.¹⁴ Therefore, I interpret these estimates as being close to zero. Most of the estimates are fairly low, this has to do with the fact that I am using sales rank as a proxy for quantity. However, the goal of this paper is just to rank firms according to the demand elasticity they face and not to document the levels of demand elasticity, thus in no analysis I will use the level of the estimated elasticity.

Figure 8: Average Elasticity Per Firm



This figure plots a histogram with the average estimated elasticities for the firms in the sample. For each product in the sample, I compute equation (17), then I average across products and firms to get an average elasticity per firm. The sample with which we compute elasticities spans from February 2015 to April 2017.

I categorize each firm in the sample according to the Fama-French 12-industries classification

¹⁴Conditional on a negative elasticity, 86% of the estimates are statistically different from zero

and compute the average industry elasticity. The first row of Table 4 reproduces the estimates for the 9 industries in my sample. All industries have negative average elasticities, and industries such as durables and manufacturing have higher demand elasticities than industries focusing on consumption goods, such as non-durables and shops. This is in line with the empirical macroeconomic evidence that in a recession (expansion), durable consumption falls (increases) more than output and non-durable consumption falls (increases) much less than output (see Kydland and Prescott (1982)).

Table 4: Average Elasticity per Industry

	Durables	Telecom	Manuf	Health	Other	Non-Durables	Chems	BusEq	Shops
Elasticity	-0.200	-0.170	-0.169	-0.138	-0.130	-0.123	-0.120	-0.119	-0.051
t-stats	-2.301	-5.283	-6.126	-4.674	-3.632	-4.348	-4.947	-5.806	-1.553

This table reports the average equally-weighted elasticities per industry. I use the Fama-French 12-industry classification. The elasticities are computed using equation (17) and the `Keepa` data. The sample period is February 2015 to April 2017. The second row of the table reports the t-statistics.

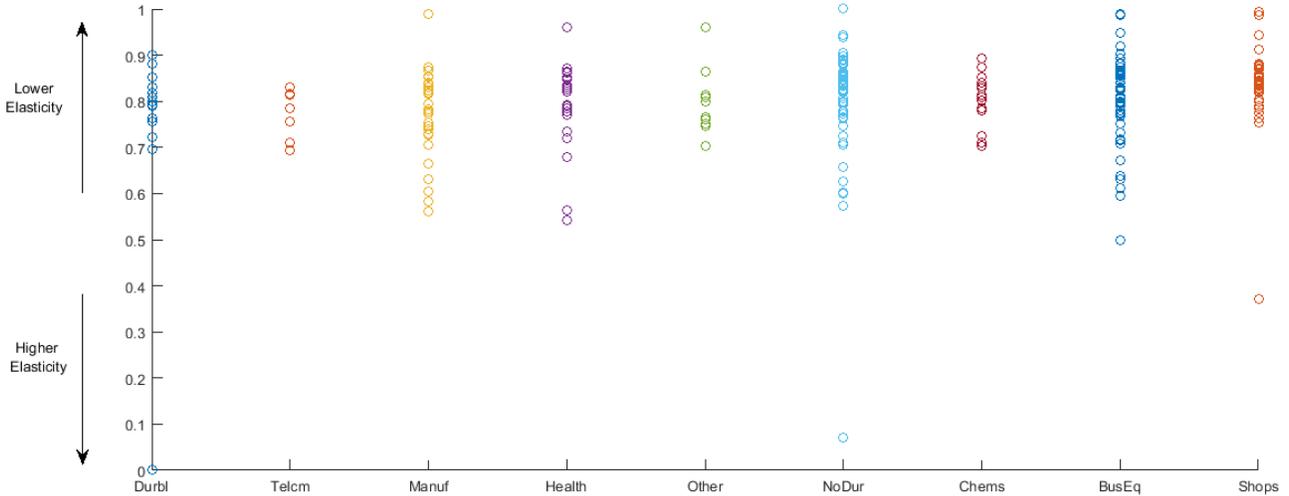
All industry elasticities are statistically negative except for the most inelastic industry (clothes), which is statistically indistinguishable from zero. The difference between the most elastic industry (Durables) and the least elastic (Clothes) is 0.149, which is statistically significant with 99% confidence. Figure 9 plots the ranking of elasticities per firm on the y -axis and the industry on the x -axis. The firm with the highest elasticity has a rank of 1 and the firm with the lowest elasticity has a rank of 0. Despite the large heterogeneity in elasticities across industries, there is still a great variation in elasticities within each industry.¹⁵ The effects reported in this paper are robust to industry level controls.

B. Portfolio Sorts and Systematic Risk

I test whether differences in the elasticity of demand that firms face are associated with differences in expected returns. I start by sorting stocks into five portfolios based on elasticity of

¹⁵There are three outliers in Figure 9. The three outliers are Leggett & Plat a premium furniture brand, Nautilus, a company that sells treadmills and other fitness equipment and AVG Technologies, a firm that sells security software. The first two outliers, have high elasticity, precisely because these companies sell luxury and expensive goods and therefore customers have the option to wait for price decreases before making purchases. The later company, also offers a free version of its software, and thus consumers have a readily available substitute.

Figure 9: Elasticities rank at the Firm-Industry Level



This scatter plot shows the rank of elasticities at the firm level (y -axis) conditional on the industry to which the firm belongs (x -axis). The highest elasticity firm has a rank of zero and the firm with the lowest elasticity has a rank of 1. Firms are classified into industries according to the Fama-French 13-industry classification. The elasticities are computed using equation (17) and the *Keepa* data.

demand. I measure returns at the daily level from the beginning of my sample in January 2011 to the end of the sample March 2017. The elasticity of the demand that firms' face is stable over time (more on this later). Therefore, I use an approach similar to Weber (2015), and do not rebalance portfolios, but only sort them once to minimize concerns about measurement error in firm elasticity estimates. In the robustness section, I show that the elasticity estimates are stable over time and that the results also hold for an out-of-sample period.

Table 5 reports the results. The demand elasticity of each portfolio is by construction monotonically increasing from low elasticity (close to zero) to high elasticity (close to -0.40). Panel A reports the results for equally weighted returns. The portfolio of firms that face high average elasticities of demand earns an average of 19.1% per year, whereas the low-elasticity portfolio earns 12.9% per year. The difference between the high- and low-elasticity portfolios is 6.2%, which is statistically and economically significant.

Panel B reports average value-weighted annual returns. The same pattern holds: returns increase monotonically from the portfolio with lowest elasticity to the one with highest elasticity,

Table 5: Sorted Portfolios on Demand Elasticity

Average Annual Excess Returns						
	Low Elasticity	2	3	4	High Elasticity	H-L
Panel A: Equally Weighted Returns						
$R_i - R_f$	0.129	0.156	0.168	0.180	0.191	0.062
$t\text{-stat}$	[2.01]***	[2.57]***	[2.48]***	[2.85]***	[2.71]***	[2.15]***
Panel B: Value-Weighted Returns						
$R_i - R_f$	0.137	0.142	0.160	0.153	0.189	0.053
$t\text{-stat}$	[1.97]**	[2.92]***	[2.76]***	[2.74]***	[2.92]***	[1.99]**

This table reports time-series averages for annual excess returns of elasticity-sorted portfolios. Firms are assigned to one of five portfolios based on the average elasticity of demand they face for their products. Portfolio low (high) has the stocks with the lowest (highest) demand elasticity. The last column of the table reports the return on a portfolio that is long high-elasticity stocks and short low-elasticity stocks. Panel A reports the results for equally-weighted portfolios and Panel B for value-weighted portfolios. The sample ranges from January 2011 to March 2017.

with a spread of 5.3% per annum.

To formally test the hypothesis that the relation is monotonically increasing, I use the monotonicity test of Patton and Timmermann (2010). This test considers the full time series of returns of each portfolio (not only the average return). The null hypothesis of the test is that there is no relation between the returns of the portfolios (i.e., a flat relation) and the alternative hypothesis can be specified as either an increasing or decreasing relationship. When the alternative hypothesis is specified as an increasing relationship, I reject the null that the relation is flat (with a p - value = 0). If, instead, the alternative hypothesis is specified as a decreasing relationship, I fail to reject the null hypothesis (with a p - value = 0.49). This is strong evidence that there is an increasing relation between demand elasticity and asset returns.¹⁶

In my theoretical framework, different exposures to the elasticity of demand are partially explained by systematic risk, i.e. CAPM betas line up weakly with stock returns. To test this prediction, I perform standard time-series tests and regress the returns of the elasticity-sorted portfolios on the market portfolio as well as on the Fama and French (1993) three factors. Let $R_{i,t}^e$ be the excess return on elasticity-sorted portfolio i , $R_{m,t}^e$ be the excess market return, and

¹⁶Table A2 from the appendix shows some of the financial properties of the firms in each portfolio.

X_t be a time-series vector with the Fama and French (1993) size and value factors. I run the following time-series regression using daily data over my sample:

$$R_{i,t}^e = \alpha_i + \beta_i R_{m,t}^e + \Gamma_i X_t + u_{i,t} \quad (18)$$

If the model is correctly specified, then exposure to market risk and to the remaining two factors should be enough to explain the elasticity-sorted portfolio's excess returns. This implies that the intercept α_i of the time-series regression (18) should be zero. Panel A of Table 6 shows the results for the Fama-French type regressions and Panel B reports the results for standard CAPM regressions. The first row of the table reports α 's for each portfolio in annualized terms. The intercepts range from 0.0 for the low-elasticity portfolio to 0.05 for the high-elasticity portfolio. Most estimates are statistically indistinguishable from zero. Furthermore, a portfolio that is long on firms that face a high elasticity of demand for their products and short on firms that face a low elasticity of demand earns an alpha of 4.6% per year, that is statistically zero. This implies that CAPM as well as the Fama-French model fully explain the cross-section of returns. As a robustness table A3 from the appendix runs the Fama-French 5-factor model, and compares the R^2 fit of different models. Around 84% to 90% of the variation in returns can be explained by a simple CAPM model. Adding additional factors, only marginally improves the explanatory fit.

The third row of Table 6 shows the market betas of each portfolio. Market betas are monotonically increasing, ranging from 0.92 for the low-elasticity portfolio to 1.047 for the high-elasticity portfolio. The fifth and seventh rows of the table show the loadings on book-to-market and size, respectively. Both of these factors also help to explain the cross-section of returns.

Panel B of the same table reports the results of the CAPM model, i.e. the estimates of running regression (18) using only the market as a factor. The results are similar to before: intercepts are statistically zero for most portfolios and CAPM betas line up well with expected returns. The low-elasticity portfolio has a beta of 0.97 and the high-elasticity portfolio has a beta of 1.12. The beta estimates are statistically significant at the 1% confidence level.

The results are robust to the exclusion of weeks surrounding holidays and occasions. These are days where there may be predictable changes in the prices of products. For example, retailers often discount their products on Black Friday. Therefore, if agents anticipate significant price changes during that day, this might bias the elasticity estimates, as demand would not be

Table 6: CAPM Regressions

Panel A: Fama-French Model

	Low Elasticity	2	3	4	High Elasticity	H-L
Annualized Alpha	0.000	0.028	0.030	0.046	0.046	0.046
<i>Std. Error</i>	[0.001]	[0.001]	[0.001]	[0.001]***	[0.001]***	[0.002]
Rm-Rf	0.921	0.928	0.989	0.967	1.047	0.126
<i>Std. Error</i>	[0.01]***	[0.009]***	[0.01]***	[0.008]***	[0.01]***	[0.013]***
Book-to-Market	0.424	0.198	0.414	0.276	0.355	-0.069
<i>Std. Error</i>	[0.018]***	[0.016]***	[0.017]***	[0.015]***	[0.018]***	[0.023]***
Size	-0.083	-0.030	-0.015	-0.058	0.010	0.094
<i>Std. Error</i>	[0.019]***	[0.017]**	[0.019]	[0.016]***	[0.019]	[0.024]***

Panel B: CAPM model

	Low Elasticity	2	3	4	High Elasticity	H-L
Annualized Alpha	-0.006	0.025	0.023	0.043	0.040	0.046
<i>Std. Error</i>	[0.002]	[0.001]	[0.002]	[0.001]***	[0.002]**	[0.002]
Rm-Rf	0.970	0.966	1.073	1.018	1.121	0.122
<i>Std. Error</i>	[0.010]***	[0.008]***	[0.010]***	[0.008]***	[0.010]***	[0.011]***

This table reports results for Fama-French 3-factor model (Panel A) and CAPM model (Panel B) time-series regressions. I regress returns of the 5 elasticity-sorted portfolios and the return on the High-Low Elasticity Portfolio on the market, size and book-to-market factors (Panel A) and on the market (Panel B). The regressions are done with daily data, and the factors data is taken from the Kenneth French data library. The first row of Panel A reports the intercept of the regression in annual terms. The third, fifth and seventh rows report the loadings on the market, size and book-to-market factors. The first row of Panel B reports the intercept of the CAPM regression in annual terms. The third row reports the loadings on the market. The standard errors are reported in brackets. The sample ranges from January 2011 to April 2017.

stable. To alleviate this concern, I run the same portfolio sorts excluding the major holidays and festivities along with a 7-day window around these days. Panel A of Table 7 reports the results of the portfolio sorts when Christmas, New Year, Valentine's, Thanksgiving, Black Friday, Father's day and Mother's day are excluded from the sample.¹⁷ The pattern is fairly similar to the pattern from Table 5. Higher elasticity portfolios earn higher excess returns with an economically and statistically meaningful spread of 7.2% (slightly higher than the 6.2%

¹⁷For Christmas and New Year, I also exclude the full set of days between these dates.

spread when holidays and occasions are not excluded).

Table 7: Elasticity-Sorted Portfolios - Robustness

Panel A: Portfolio Sorts Excluding Holiday Days and Occasions						
	Low Elasticity	2	3	4	High Elasticity	H-L
$R_i - R_f$	0.124	0.1694	0.1578	0.1738	0.196	0.072
t -stat	[1.96]**	[2.83]***	[2.35]***	[2.71]***	[2.76]***	[2.47]***
FPA (months)	4.79	4.83	5.23	4.7	5.03	

Panel B: Portfolio Sorts and Frequency of Price Adjustment						
	Low Elasticity	2	3	4	High Elasticity	H-L
$R_i - R_f$	0.129	0.156	0.168	0.18	0.191	0.062
t -stat	[2.01]***	[2.57]***	[2.48]***	[2.85]***	[2.71]***	[2.15]***
FPA (months)	4.89	4.72	5.2	4.92	4.86	

This table reports time-series averages of annual excess returns of elasticity-sorted portfolios. Panel A in this table is the same as the first Panel of Table 5, but with an additional line that reports the frequency of price adjustment at the portfolio level. I compute frequency of price adjustment for each firm using the methodology of Midrigan (2011). Panel B shows the results of the portfolio sorts when a week surrounding a holiday or occasion is excluded (except for Christmas and New Year, where I also exclude the full period between those two dates). I exclude the days around Christmas, New Year, Valentine’s, Thanksgiving, Black Friday, Father’s day and Mother’s day.

Finally, another important concern is whether heterogeneity in the frequency of price adjustment might be affecting the results. If the portfolios sorted by elasticity of demand correlate with the degree of nominal rigidities, then this might bias the results. Panel B of the same table reports the results of the elasticity-sorted portfolios along with the frequency of price adjustment within each portfolio. It seems that firms facing high elasticities of demand do not adjust their prices more often than firms facing low elasticities of demand. In the robustness section, I formally address this hypothesis and test whether there is any relation between the two.

C. Panel Regressions and Implications for Markups

The model from Section III.A predicts that firms with higher demand elasticities have lower markups. To test this hypothesis, I make use of the heterogeneity in demand elasticities across

firms and run standard panel regressions. Specifically, I run the following regression at an annual frequency:

$$Y_{f,t} = \alpha + \beta_{\epsilon_f} \times \epsilon_f + \gamma_t + u_{t,i} \quad (19)$$

where $Y_{f,t}$ is the outcome variable of interest (returns or markups), ϵ_f is a variable that indicates the elasticity quintile of firm f (higher value means more elastic demand) and γ_t are year-fixed effects.¹⁸ Table 8 shows the results. The first column of the table reports the results for the regression using returns as a dependent variable. Moving from a firm in the lowest quintile of elasticity (more inelastic demand) to the highest quintile of elasticity yields a return differential of 5.5%, which is in line with the estimates in Table 5.¹⁹

Table 8: Panel Regressions of Demand Elasticity on Firm Returns and Markups

	Return	Return	Return	Markups	Markups	Markups
Elasticity Quintiles	0.015	0.015	0.015	-0.054	-0.053	-0.031
(t-stat)	(2.25)***	(2.45)***	(2.25)***	(-4.23)***	(-4.21)***	(-2.55)***
Year FE	N	Y	N	N	Y	N
Year \times Industry FE	N	N	Y	N	N	Y
R^2	0.004	0.112	0.118	0.010	0.010	0.121
# Obs	1374	1374	1374	1564	1564	1564

This table reports the results estimating equation (19). The regression is done using annual data. Standard errors are clustered at the firm level and t -stats are reported in parentheses.

Adding time-fixed effects or time \times industry-fixed effects to the regression (column 2 of Table 8) does not significantly change the coefficient of the regression. This implies that industry-specific factors are not driving the results. The fourth column shows the results of regressing markups on the elasticities. Compustat accounting data does not allow direct estimation of markups. Therefore, as a proxy for markups, I use gross margin as in Bustamante and Donangelo (2017) and Corhay, Kung, and Schmid (2017) (in Appendix AVI, I detail how I compute markups and discuss the literature on the relation between markups and the cross-section of returns). As we move from a firm with a low elasticity to a higher elasticity quintile, markups

¹⁸The panel regression results are robust to the use of terciles, quintiles or deciles.

¹⁹Moving a quintile up in the elasticity yields a return differential of 1.1%; therefore moving from the first to the last quintile yields a return differential of $1.1\% \times 5$.

decrease by 5 percentage points. This yields a margin differential of around 25 percentage points between the lowest and highest quintiles.

This empirical evidence on markups allows me to differentiate the mechanism that generates the spread in returns from other mechanisms that have been proposed in the literature. Weber (2015) argues that the frequency of price adjustment is a dimension of heterogeneity that also generates a spread in returns: in particular, firms with a lower frequency of price adjustment earn higher expected returns. However, in his model, the predictions for markups would go in the opposite direction (see Table 2). The empirical evidence on markups is consistent with the theoretical implications of differences in elasticity of demand.

The model also has implications for firms' revenues and markups. In the model firms facing higher demand elasticities have higher revenues and lower profit margins (as a consequence of the lower markup they charge). I qualitatively test that relationship in the data (see table A5) and find that indeed firms with higher demand elasticity have lower profit margins and higher revenues.²⁰

D. Robustness

In this section, I run several robustness tests. First I replicate the main finding that firms with higher demand elasticity earn on average higher returns using instrumental variables to estimate demand. Next, I validate whether the results are robust to an out-of-sample period and whether elasticities are stable over time. Finally, I establish the relation between the degree of price stickiness and elasticity of demand.

D.1. Demand Elasticity using Instrumental Variables

I replicate the main result estimating demand elasticity using instrumental variables. Given the nature of my data, a feasible instrument to identify demand is prices of the same products in other markets. These are the so called "Hausman (1996)) Instruments". The underlying assumption is that demand for a product in a given market m is uncorrelated with demand for the same product in another market n . Therefore, using prices in market n to instrument for prices in market m captures supply shifts and traces demand.

I rely again on the *Keepa* database and extract time-series information for all the products in my sample, but in the *Canadian* Amazon market. In Appendix AIV, I detail how I match the

²⁰Moving a quintile up in the elasticity yields a decrease in profit margin of 0.1 percentage points and the same move would lead to an increase in revenues of around 6%

data between the US Amazon market and the Canadian Amazon market. If demand shocks are uncorrelated across the US market and the Canadian market, then prices on Amazon Canada can be reliably used as instruments. In particular I run the following specification:

$$\log(SR_{j,m,t}) = \delta_{j,m} \log(p_{j,m,t}) + w_{j,n,t} \quad (20)$$

where $SR_{j,m,t}$ ($p_{j,m,t}$) is the log of sales rank (price) of product j in market m at time t . I use price $p_{j,m,t}$ as an instrument for price $p_{j,n,t}$ using the price of the same product at time t in another market n . The coefficient of interest is $\delta_{j,m}$, which can be interpreted as the elasticity of demand. I average the estimates $\delta_{j,m}$ to get an average elasticity faced by a firm. The main drawback of using this methodology to compute demand elasticity is the reduction in sample size. The sample now contains 94 firms, with an average of 237 products and a median of 77.

I follow the same methodology as before and sort stocks into portfolios based on the elasticity of demand. Given the lower number of stocks in my sample, I sort stocks into four portfolios instead of five. This ensures that a significant number of stocks is included in each portfolio.

Panel A from Table 9 reports the results of the instrumental variable sort. The portfolio of firms that face high average elasticities of demand earns an average of 19.0% per year, whereas the low-elasticity portfolio earns 12.8% per year. These estimates are in line with those reported in Table 5. The difference between the high- and low-elasticity portfolios is 7%, which is statistically significant. Panel B from the same table reports the results using the high-frequency demand elasticity estimation method, where only the products and firms used in the IV estimation are included. Most firms are assigned to the same portfolio using either method (instrumental variables or high-frequency), so the results of the portfolio sorts are fairly similar. This confirms the hypothesis that the high-frequency elasticity estimation method delivers results that are consistent with standard instrumental variables methods for estimating demand.

D.2. Stability of Elasticities and Out-of-Sample

An important assumption underlying the results of the portfolio sorts is that demand elasticities are stable over time. It is important that this is indeed the case, because the estimated portfolios were not rebalanced. I test this assumption using two different approaches. First, I set a cutoff date that splits the sample exactly into two identically sized time series and I estimate elasticities at the firm level on both samples. If elasticities are stable over time, there should be

Table 9: Portfolios Sorted by Demand Elasticity estimated using Instrumental Variables

Average Annual Excess Returns - instrumental variables					
	Low Elasticity	2	3	High Elasticity	H-L
	<u>Average Returns</u>				
$R_i - R_f$	0.128	0.169	0.168	0.192	0.063
$t\text{-stat}$	[1.95]*	[2.45]***	[2.42]***	[3.03]***	[2.09]***

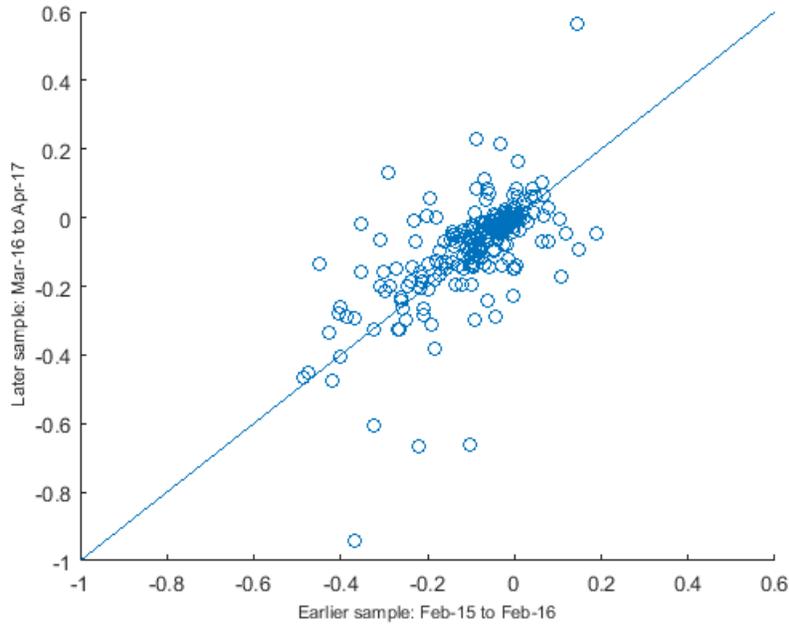
Average Annual Excess Returns - high frequency					
	Low Elasticity	2	3	High Elasticity	H-L
	<u>Average Returns</u>				
$R_i - R_f$	0.132	0.179	0.157	0.181	0.049
$t\text{-stat}$	[2.17]*	[2.87]***	[2.24]***	[2.54]***	[1.5]*

This table reports time-series averages of annual excess returns of elasticity sorted portfolios where I use Hausman Instruments to estimate elasticity of demand. In particular, I use prices of the same SKU (stock keeping unit) in Amazon Canada to instrument for prices in Amazon USA. Firms are then assigned to one of five portfolios based on the average elasticity of demand they face for their products. Portfolio low (high) has the stocks with the lowest (highest) demand elasticity. The last column of the table reports the return on a portfolio that is long high-elasticity stocks and short low-elasticity stocks. The sample ranges from January-2011 to March -2017.

no differences between the two samples. Figure 10 plots the results for the sample splits. The blue dots in the figure are the elasticities of each of the 250 firms in my sample. The x -axis displays the elasticities for the earlier part of the sample and the y -axis shows the elasticities for the latter part of the sample. If elasticities are indeed stable, the points should cluster around the 45-degree line. This seems to be the case. To formally test this hypothesis, I run a Wilcoxon signed rank test. This is an extension of the standard t -test for equality of means when there are multiple means and the normality of the data cannot be assumed. The null hypothesis of the test is that all the means of the two samples are pairwise identical. The test leads to a failure to reject the null hypothesis, meaning that elasticities are indeed stable over time (p - value of 0.50).

Second, I make an out-of-sample exercise, by repeating the portfolio sorts done in Section V.B, using a different sample from the one with which I estimated the elasticities. If elasticities are indeed stable over time, then the return spreads from Table 5 should be similar in an

Figure 10: Subsample Splits



In this figure, I split the sample into two. The first sample ranges from February 2015 to February 2016, and the second sample from March 2016 to April 2017. I compute average elasticities per firm for the two subsample periods. The resulting elasticities are plotted in a scatter plot with the earlier sample on the x -axis and the later sample on the y -axis.

out-of-sample period. Remember that I have estimated the elasticities using a sample that started in February 2015. Usually, out-of-sample exercises are forward-looking, meaning that, conditionally on an estimate, researchers look to the period ahead (e.g. Welch and Goyal (2007)). However, an out-of-sample exercise should also be robust to backward-looking, which is often more challenging. Given the data, I estimate the elasticities using the *Keepa* sample and look backward in time. I use an out-of-sample time series ranging from January 2001 to January 2010. The start date of this subsample is chosen to ensure that at least 30 firms are kept in each portfolio at any point in time. The last date is chosen just before *Keepa* started tracking Amazon products. The results are shown in Table 10 and are similar to the ones presented before. The portfolio of firms that face a higher elasticity of demand earn excess returns that are on average 4.55% larger than the low-elasticity portfolio. This is slightly lower than the previous estimate for the later sample, but economically and statistically significant.

Furthermore, although the returns on the five portfolios are not strictly monotonic, I still reject the null of a flat relationship in returns according to the Patton and Timmermann (2010) test.

Table 10: Portfolio Sorts Out-of-Sample

	Average Annual Excess Returns - out-of-sample					
	Low Elasticity	2	3	4	High Elasticity	H-L
$R_i - R_f$	0.174	0.191	0.200	0.186	0.219	0.046
$t\text{-stat}$	[2.50]***	[2.68]***	[3.07]***	[2.70]***	[2.98]***	[1.98]**

This table reports time-series averages of annual excess returns of elasticity-sorted portfolios. Stocks are assigned to one of five portfolios based on the average elasticity of demand they face for their products. The elasticity estimates are done using a sample from February 2015 to April 2017. The portfolio returns are out-of-sample with a sample ranging from January 2001 to January 2010. Portfolio low (high) has the stocks with the lowest (highest) demand elasticity. The last column of the table reports the return on a portfolio that is long high-elasticity stocks and short low-elasticity stocks.

D.3. Degree of Price Stickiness and Demand Elasticities

Up until now, I have argued that firms with higher demand elasticities are riskier in equilibrium and therefore earn higher expected returns. It would be plausible that these firms, with higher demand elasticity, would adjust prices more often (in fact, Weber (2015) shows that different firms adjust prices with different frequencies). If firms facing higher demand elasticities are riskier, then one could expect that these firms would change their prices more frequently. To test this relation formally, I estimate the following regression:

$$\epsilon_i = \alpha + \beta_{FPA} FPA_i + \sum_n \beta_n X_{t,i,n} + \mu_t + \epsilon_{i,t} \quad (21)$$

where ϵ_i is the elasticity of demand of firm i , FPA_i is the frequency of price adjustment, $X_{t,i}$ is a set of time-varying controls and μ_t are time-fixed effects. Table 11 reports the results. The first column shows the results of regressing elasticity on frequency of price adjustment. The coefficient is positive and marginally significant. This implies that there is a weak positive association between the degree of nominal rigidities and elasticity, i.e. firms facing more elastic demands adjust prices more often.

This result can also be seen graphically: Panel A of Figure 11 illustrates at the firm level

Table 11: Demand Elasticity and Frequency of Price Adjustment

	(1)	(2)	(3)	(4)	(5)	(6)
freqpradj	0.00431* (1.729)					0.00457 (1.648)
hhi		0.0259 (0.901)				0.0617** (2.047)
beta			-0.0810*** (-4.754)			-0.0974*** (-4.523)
size				-0.00245 (-1.192)		-0.00693** (-2.196)
b2m					0.0108** (1.995)	0.00785 (1.073)
Year Fixed Effects	Y	Y	Y	Y	Y	Y
Observations	1,456	1,449	1,358	1,623	1,546	890
R-squared	0.002	0.001	0.016	0.001	0.003	0.035

t-statistics in parentheses

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

This table reports the results of regressing demand elasticity faced by a firm on the degree of price stickiness and several firm characteristics. The regression is done using annual data. Standard errors are clustered at the firm level and $t - stats$ are reported in parentheses.

the relation between price stickiness and elasticity. The relation is weakly positive. Panel B of the same figure illustrates the relation but at the industry level. Again the relation is positive but weak, as the slope coefficient is not statistically different from zero.²¹

Remember from Table 7 that portfolios sorted on demand elasticity did not show differences in their degree of nominal rigidities. This is an important and puzzling finding that demand elasticity is if anything weakly related to sticky prices, and it might help future researchers to explain the drivers of heterogeneity in the degree of price adjustment. Another way of looking into this, is by double sorting portfolios, based on the degree of price elasticity and the degree of nominal rigidities. I do so in Table A4 from the appendix. Consistent with the results above,

²¹Once a set of standard controls from the literature are added to the panel regression above, the coefficient becomes insignificant. The last column of Table 11 reports the results when firm and industry controls are added to the regression, namely size, book-to-market, Herfindahl-Hirschman index (HHI), and the CAPM beta. Most of the variation in elasticities can be explained by differences in systematic risk (or beta). This is the same result as the one from Section V.B: heterogeneity in demand elasticities is fully captured by differences in systematic risk.

firms with higher elasticity have higher returns, but the degree of price adjustment does not seem to influence returns in my sample.

It is beyond the scope of this paper to understand why such correlation is so low, though I provide some plausible explanations. The first possible explanation for such a low correlation is that firms are uncertain about the demand elasticity they face and learn through their pricing decisions. This would be in accordance with the recent study of Argente and Yeh (2017), who find that the magnitude and likelihood of a price change is higher in the early stage of a product life-cycle. In their model, firms face an intertemporal trade-off between maximizing static profits and learning about demand. Firms with high elasticity of demand might indeed change their prices less often, as price changes can be very costly for them. Second, if not all consumers are informed about prices of goods, it might indeed be the case that firms have the incentive to charge different prices (even for homogeneous goods) and the degree of demand elasticity is related to the amount of informed/uninformed consumers in the market and not to how often a firm changes its price. This intuition can be seen in the models of Varian (1980), Burdett and Judd (1983) and Stahl (1989). In these models, price dispersion among firms arises endogenously, as there is no pure strategy price equilibrium. The proportion of informed consumers drives the degree of demand elasticity whereas frequency of price adjustment is purely determined by changes in marginal costs. A third possible explanation for why firms with higher elasticities adjust prices more often comes from the supply side. If firms face shocks that affect their marginal costs and the frequency or magnitude of these shocks is related to the demand elasticity of their products, then it is possible to empirically observe a lack of relation between nominal rigidities and demand elasticity.

VI. Conclusion

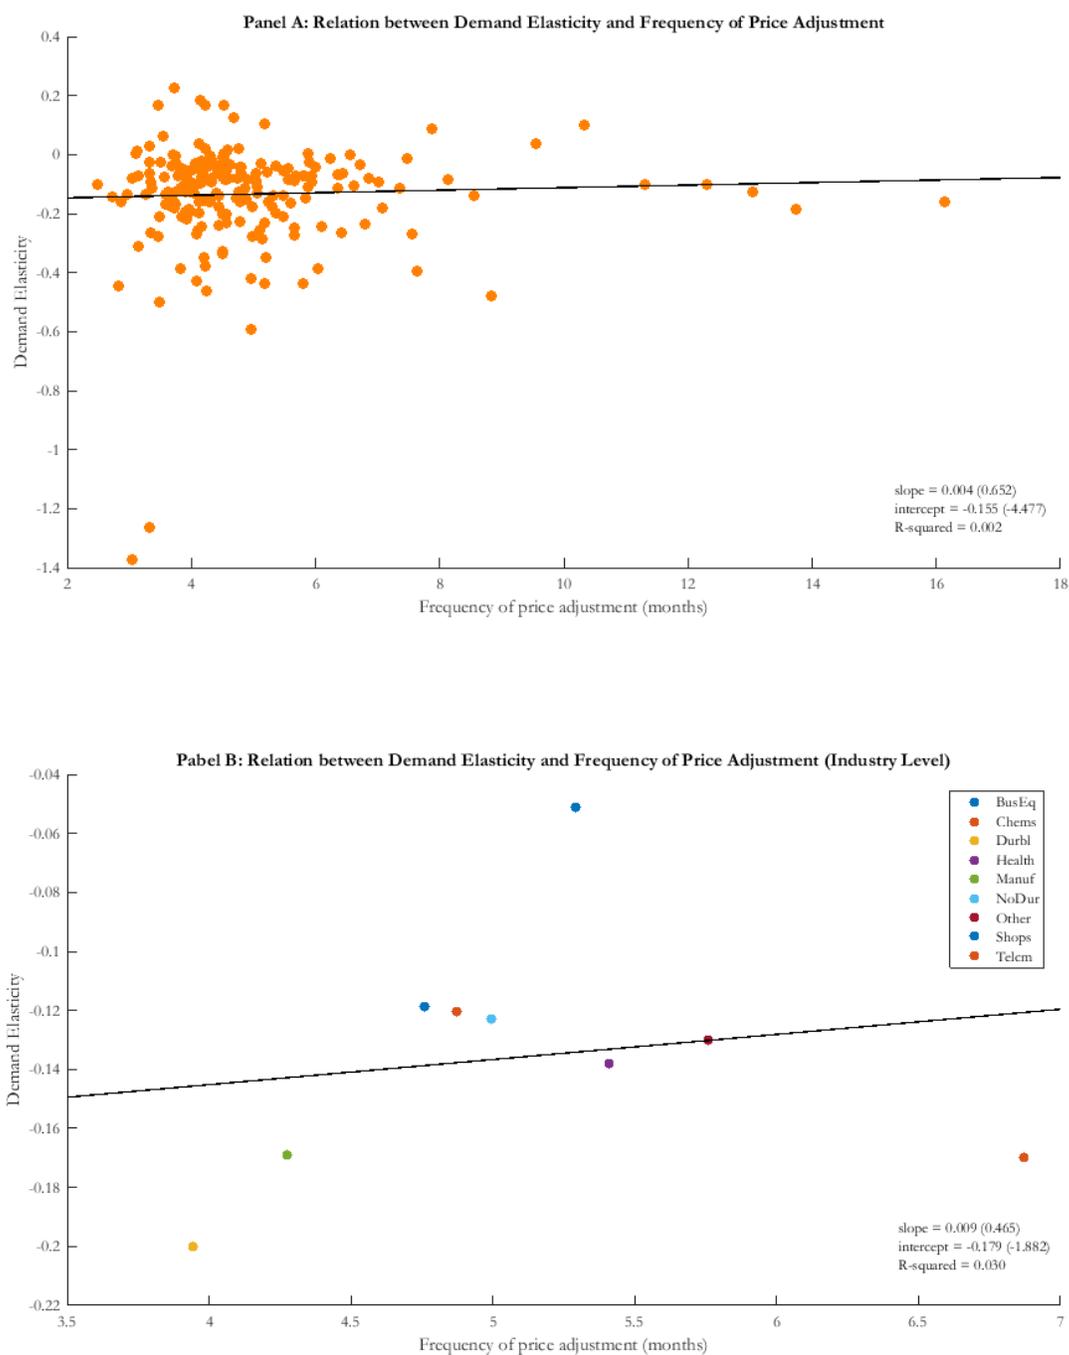
In this paper, I have studied the relationship between demand elasticities, sticky prices and asset prices. I have generalized a multi-sector New-Keynesian model to allow for differences in demand elasticity faced by firms. I have shown theoretically that firms with more elastic demands earn higher returns and have lower markups in equilibrium.

To test the theoretical predictions, I have developed a novel way of estimating demand elasticities at the firm level by exploiting high-frequency data on products sold on Amazon. I show that demand elasticity matters for cross-section market outcomes: firms facing higher elasticities of demand earn higher returns in equilibrium. The spread in returns is weakly

explained by systematic risk: firms with high demand elasticities also have higher CAPM betas. The results are not driven by industry characteristics and are consistent with a standard New-Keynesian model in which firms have different elasticities of demand.

Nominal rigidities are crucial for the heterogeneity in returns driven by different elasticities of demand. In the absence of sticky prices, there would be no differences in expected returns of firms with different demand elasticities. Interestingly, firms with high demand elasticities do not adjust prices more frequently. This evidence deepens the puzzle regarding the micro-foundations of sticky prices. Standard Calvo models and menu cost models cannot account for the low degree of price synchronization observed in online markets. In future work, I plan to use this high-frequency dataset to study what drives price stickiness and heterogeneity in price stickiness.

Figure 11: Cross-sectional Relation Between Price Stickiness and Demand Elasticity



50

This figure shows the relation between the degree of price stickiness (x -axis) and the demand elasticity (y -axis). Panel A plots the scatter plot at the firm level. The fitted line is an ordinary-least-squares regression. Panel B plots the same scatter plot and fitted line, but at the industry level.

References

- ARGENTE, D., AND C. YEH (2017): “Products Life Cycle, Learning, and Nominal Shocks,” *Working Paper*.
- BANSAL, R., D. KIKU, AND A. YARON (2010): “Long run risks, the macroeconomy, and asset prices,” *The American Economic Review*, 100(2), 542–546.
- BEELER, J., AND J. Y. CAMPBELL (2009): “The long-run risks model and aggregate asset prices: An empirical assessment,” Discussion paper, National Bureau of Economic Research.
- BERRY, S., J. LEVINSOHN, AND A. PAKES (1995): “Automobile prices in market equilibrium,” *Econometrica: Journal of the Econometric Society*, pp. 841–890.
- BILS, M., AND P. J. KLENOW (2004): “Some evidence on the importance of sticky prices,” *Journal of Political Economy*, 112(5), 947–985.
- BRODA, C., AND D. E. WEINSTEIN (2010): “Product Creation and Destruction: Evidence and Price Implications,” *The American Economic Review*, 100(3), 691–723.
- BRYNJOLFSSON, E., Y. J. HU, AND M. D. SMITH (2006): “From niches to riches: Anatomy of the long tail,” .
- BURDETT, K., AND K. L. JUDD (1983): “Equilibrium price dispersion,” *Econometrica: Journal of the Econometric Society*, pp. 955–969.
- BUSTAMANTE, M. C., AND A. DONANGELO (2017): “Product market competition and industry returns,” *The Review of Financial Studies*, 30(12), 4216–4266.
- CALVO, G. A. (1983): “Staggered prices in a utility-maximizing framework,” *Journal of Monetary Economics*, 12(3), 383–398.
- CARLSON, M., A. FISHER, AND R. GIAMMARINO (2004): “Corporate investment and asset price dynamics: Implications for the cross-section of returns,” *The Journal of Finance*, 59(6), 2577–2603.
- CARVALHO, C. (2006): “Heterogeneity in Price Stickiness and the Real Effects of Monetary Shocks,” *The BE Journal of Macroeconomics*, 6(3), 1–58.
- CAVALLO, A. (2016): “Scraped data and sticky prices,” *Review of Economics and Statistics*, (forthcoming).
- (2017): “Are online and offline prices similar? Evidence from large multi-channel retailers,” *The American Economic Review*, 107(1), 283–303.
- CHEVALIER, J., AND A. GOOLSBEE (2003): “Measuring prices and price competition online: Amazon. com and BarnesandNoble. com,” *Quantitative Marketing and Economics*, 1(2), 203–222.

- CHEVALIER, J. A., AND D. MAYZLIN (2006): “The effect of word of mouth on sales: Online book reviews,” *Journal of Marketing Research*, 43(3), 345–354.
- CHODOROW-REICH, G. (2014): “Effects of Unconventional Monetary Policy on Financial Institutions,” *Brookings Papers on Economic Activity*, (Spring), 155–204.
- CHRISTIANO, L. J., M. EICHENBAUM, AND C. L. EVANS (2005): “Nominal rigidities and the dynamic effects of a shock to monetary policy,” *Journal of Political Economy*, 113(1), 1–45.
- CORHAY, A., H. KUNG, AND L. SCHMID (2017): “Competition, markups and predictable returns,” .
- DE PAOLI, B., A. SCOTT, AND O. WEEKEN (2010): “Asset pricing implications of a New Keynesian model,” *Journal of Economic Dynamics and Control*, 34(10), 2056–2073.
- DIXIT, A. K., AND J. E. STIGLITZ (1977): “Monopolistic competition and optimum product diversity,” *The American Economic Review*, 67(3), 297–308.
- EPIFANI, P., AND G. GANCIA (2011): “Trade, markup heterogeneity and misallocations,” *Journal of International Economics*, 83(1), 1–13.
- ERCEG, C. J., D. W. HENDERSON, AND A. T. LEVIN (2000): “Optimal monetary policy with staggered wage and price contracts,” *Journal of Monetary Economics*, 46(2), 281–313.
- FAMA, E. F., AND K. R. FRENCH (1993): “Common risk factors in the returns on stocks and bonds,” *Journal of Financial Economics*, 33(1), 3–56.
- (2015): “A five-factor asset pricing model,” *Journal of financial economics*, 116(1), 1–22.
- FEENSTRA, R. C. (1994): “New product varieties and the measurement of international prices,” *The American Economic Review*, pp. 157–177.
- GALI, J. (1994): “Monopolistic competition, business cycles, and the composition of aggregate demand,” *Journal of Economic Theory*, 63(1), 73–96.
- GALÍ, J., F. SMETS, AND R. WOUTERS (2012): “Unemployment in an estimated new keynesian model,” *NBER Macroeconomics Annual*, 26(1), 329–360.
- GERTLER, M., AND P. KARADI (2015): “Monetary policy surprises, credit costs, and economic activity,” *American Economic Journal: Macroeconomics*, 7(1), 44–76.
- GOMES, J. F., AND L. SCHMID (2010): “Levered returns,” *The Journal of Finance*, 65(2), 467–494.
- GORODNICHENKO, Y., V. SHEREMIROV, AND O. TALAVERA (2014): “Price setting in online markets: does it click?,” Discussion paper, National Bureau of Economic Research.

- GORODNICHENKO, Y., AND M. WEBER (2016): “Are sticky prices costly? Evidence from the stock market,” *The American Economic Review*, 106(1), 165–199.
- GOWRISANKARAN, G., AND M. RYSMAN (2012): “Dynamics of consumer demand for new durable goods,” *Journal of Political Economy*, 120(6), 1173–1219.
- HAUSMAN, J. A. (1996): “Valuation of new goods under perfect and imperfect competition,” in *The Economics of New Goods*, pp. 207–248. University of Chicago Press.
- HERBST, E. P., AND F. SCHORFHEIDE (2015): *Bayesian estimation of DSGE models*. Princeton University Press.
- HOU, K., AND D. T. ROBINSON (2006): “Industry concentration and average stock returns,” *The Journal of Finance*, 61(4), 1927–1956.
- HUANG, J. (2016): “The customer knows best: The investment value of consumer opinions,” *Journal of Financial Economics*, (forthcoming).
- KEHOE, P., AND V. MIDRIGAN (2015): “Prices are sticky after all,” *Journal of Monetary Economics*, 75, 35–53.
- KOGAN, L., AND D. PAPANIKOLAOU (2013): “Firm characteristics and stock returns: The role of investment-specific shocks,” *The Review of Financial Studies*, 26(11), 2718–2759.
- KUNG, H. (2015): “Macroeconomic linkages between monetary policy and the term structure of interest rates,” *Journal of Financial Economics*, 115(1), 42–57.
- KYDLAND, F. E., AND E. C. PRESCOTT (1982): “Time to build and aggregate fluctuations,” *Econometrica: Journal of the Econometric Society*, pp. 1345–1370.
- LETTAU, M., AND S. LUDVIGSON (2001): “Resurrecting the (C) CAPM: A cross-sectional test when risk premia are time-varying,” *Journal of Political Economy*, 109(6), 1238–1287.
- LI, E. X., AND F. PALOMINO (2014): “Nominal rigidities, asset returns, and monetary policy,” *Journal of Monetary Economics*, 66, 210–225.
- MARSHALL, A. (1890): *Principles of Economics*, no. v. 1 in *Principles of Economics*. Macmillan and Company.
- MIDRIGAN, V. (2011): “Menu costs, multiproduct firms, and aggregate fluctuations,” *Econometrica*, 79(4), 1139–1180.
- NAKAMURA, E., AND J. STEINSSON (2008): “Five facts about prices: A reevaluation of menu cost models,” *The Quarterly Journal of Economics*, 123(4), 1415–1464.
- NAKAMURA, E., AND J. STEINSSON (2010): “Monetary non-neutrality in a multisector menu cost model,” *The Quarterly Journal of Economics*, 125(3), 961–1013.

- NAKAMURA, E., AND J. STEINSSON (2018): “High-frequency identification of monetary non-neutrality: The information effect,” *The Quarterly Journal of Economics*, 133(3), 1283–1330.
- NEVO, A. (2001): “Measuring market power in the ready-to-eat cereal industry,” *Econometrica*, 69(2), 307–342.
- OLIVEI, G., AND S. TENREYRO (2007): “The timing of monetary policy shocks,” *The American Economic Review*, 97(3), 636–663.
- PÁSTOR, L., AND R. F. STAMBAUGH (2003): “Liquidity risk and expected stock returns,” *Journal of Political economy*, 111(3), 642–685.
- PATTON, A. J., AND A. TIMMERMANN (2010): “Monotonicity in asset returns: New tests with applications to the term structure, the CAPM, and portfolio sorts,” *Journal of Financial Economics*, 98(3), 605–625.
- RABANAL, P., AND J. F. RUBIO-RAMÍREZ (2005): “Comparing New Keynesian models of the business cycle: A Bayesian approach,” *Journal of Monetary Economics*, 52(6), 1151–1166.
- RAITH, M. (2003): “Competition, risk, and managerial incentives,” *The American Economic Review*, 93(4), 1425–1436.
- SIMS, E., AND J. WOLFF (2017): “State-dependent fiscal multipliers: Calvo vs. Rotemberg,” *Economics Letters*, 159, 190–194.
- SMETS, F., AND R. WOUTERS (2007): “Shocks and frictions in US business cycles: A Bayesian DSGE approach,” *American Economic Review*, 97(3), 586–606.
- STAHL, D. O. (1989): “Oligopolistic pricing with sequential consumer search,” *The American Economic Review*, pp. 700–712.
- TAYLOR, J. B. (1993): “Discretion versus policy rules in practice,” in *Carnegie-Rochester Conference Series on Public Policy*, vol. 39, pp. 195–214. Elsevier.
- VARIAN, H. R. (1980): “A model of sales,” *The American Economic Review*, 70(4), 651–659.
- WEBER, M. (2015): “Nominal rigidities and asset pricing,” *Working Paper*.
- WELCH, I., AND A. GOYAL (2007): “A comprehensive look at the empirical performance of equity premium prediction,” *The Review of Financial Studies*, 21(4), 1455–1508.
- WOODFORD, M. (2013, chapter 4.1): *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press.
- ZHANG, L. (2005): “The value premium,” *The Journal of Finance*, 60(1), 67–103.

AI. Static Model

AI.A. Static Model

Consider a one-period (two-date) partial equilibrium model. There is a continuum of monopolistic competitive firms that maximize profits subject to the demand for their products. Demand for goods produced by firm i is given by:

$$Q_i = \left(\frac{P_i}{P}\right)^{-\eta_i} Y \quad (\text{I.1})$$

where Q_i is the quantity demanded, P_i is the price set by firm i , P is the price index of goods sold in the sector and Y is the aggregate sectoral demand.²² It follows from the above specification that the own-price elasticity of demand for good i is given by η_i . For simplicity, assume that firms have quadratic costs of producing goods such that total costs are given by: $C = c_i Q^2$, where $c_i > 0$ is a parameter. Firms face a nominal friction and with probability θ_i are unable to adjust their price P_i . In the optimal symmetric equilibrium, each firm i will set a price P_i^* that is a markup over marginal cost:

$$P_i^* = P = \frac{\eta}{\eta - 1} 2cY \quad (\text{I.2})$$

Now consider a shock that moves marginal costs c to all firms i . In an equilibrium model, this could be motivated by a shift in aggregate productivity. What would happen to the firm's profit if the firm could not adjust its price? Panel A of Figure A1 plots the profit loss due to price stickiness if a shock to c_i occurs. Profit loss is defined as the difference between the profit the firm makes when it is not allowed to change its price (π) and the profit it makes if allowed to change its price (π^*), divided by the latter:

$$\pi^{loss} = \frac{\pi - \pi^*}{\pi^*} \quad (\text{I.3})$$

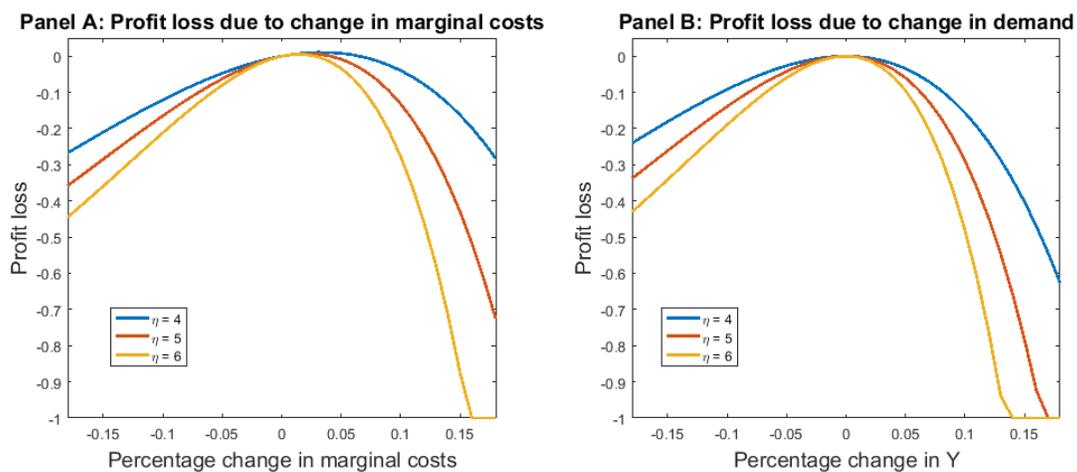
The measure is plotted for several values of η_i . If there is no change in costs, there is no profit loss. This is the middle point in the graph. If c moves then the profit losses start to increase. The larger the shock, the further away firm i is from the optimal price and the bigger its losses are. The loss is more pronounced for cost increases than for cost decreases,

²²This demand function can be micro-founded through a Dixit and Stiglitz (1977) aggregator over different varieties.

this is due to the fact that a cost increase reduces the markup the firm receives, whereas a cost decrease actually increases it. The loss in profit is an order of magnitude larger if the elasticity of demand is larger, meaning that for the same shock, firms with greater demand elasticity want to move their prices by more and therefore are further away from their optimal price. Notice that regardless of whether the shock is positive or negative, there is always an inefficiency driven by the existence of nominal rigidities. This is the key mechanism that will be present in the dynamic model in the next section: when facing a shock, firms with a greater elasticity of demand are more likely to be further away from their optimal price and therefore are riskier in equilibrium. The nominal rigidity plays a key role in this mechanism. Given that price stickiness is more costly for firms with higher elasticity of demand, these firms should adjust their prices more frequently. This should dampen the difference in risk among firms with different elasticities. In the quantitative model below, I allow for this possibility.

Figure A1: Comparative Statics on the Elasticity of Demand η

This figure plots the loss in profits when firms do not adjust their prices after a change in cost (Panel A) and demand (Panel B). The comparative statics is done concerning the elasticity of demand parameter η . The parameters underlying the figure are: $c = 5$, $Y = 5$ and $\eta \in \{4, 5, 6\}$.



AII. Dynamic Model

AII.A. Households

A representative agent has preferences given by:

$$E_t \sum_{t=0}^{\infty} \left[\frac{(C_t - \nu C_{t-1})^{1-\gamma}}{1-\gamma} - \phi_L \int_0^1 \frac{n_{t,i}^{1+\sigma}}{1+\sigma} di \right] \quad (\text{II.1})$$

Utility depends on current consumption C_t relative to lagged consumption (the external habit reference level), and $n_{t,i}$ the number of hours worked. Households maximize utility subject to their budget constraint:

$$P_t C_t = \int_0^1 w_{i,t} h_{i,t} di + R_{t-1} B_{t-1} - B_t + D_t \quad (\text{II.2})$$

where P_t is the aggregate price level, R_{t-1} is the nominal gross return on a zero-coupon bond which is in zero net supply and D_t are aggregate profits.

Let λ_t be the Lagrange multiplier of the constraint. The first-order conditions with respect to consumption and bond holdings are given by:

$$(C_t) : \quad (C_t - \nu C_{t-1})^{-\gamma} - \lambda_t P_t = 0 \quad (\text{II.3})$$

$$(B_t) : \quad -\lambda_t + E_t[\beta \lambda_{t+1} R_t] = 0 \quad (\text{II.4})$$

Combine the first-order conditions to get the Euler equation:

$$1 = \beta R_t E_t \left[\frac{1}{\pi_{t+1}} \frac{(C_{t+1} - \nu C_t)^{-\gamma}}{(C_t - \nu C_{t-1})^{-\gamma}} \right] \quad (\text{II.5})$$

Each type of household i is specialized in one type of labor that is supplied monopolistically, and in each period only a fraction of labor types can adjust their posted nominal wage.²³ Their reset wage $w_{t,i}^*$ is defined as the solution to the following problem:

$$\max_{w_{s,i}^*} \sum_{s=0}^{\infty} \left\{ (\beta \theta_w)^s \left[-\phi_L \frac{n_{s,i}^{1+\sigma}}{1+\sigma} + \lambda_s w_{s,i}^* n_{s,i} \right] \right\} \quad (\text{II.6})$$

²³This assumption can be interpreted as households/workers specialized in a given labor service represented by a trade union.

subject to labor demand on the part of firms:

$$n_{i,t} = \left(\frac{w_{s,i}^*}{W_t} \right)^{-\eta_w} \quad (\text{II.7})$$

where η_w is the elasticity of substitution between labor types and W_t is the aggregate wage rate: $W_t = \left[\int_0^1 W_{t,i}^{1-\eta_w} \right]^{\frac{1}{1-\eta_w}}$. In a symmetric equilibrium, the solution to the maximization problem (II.6) is the optimal reset wage:

$$(w_{t,i}^*)^{1+\eta_w\sigma} = \phi_l \frac{\eta_w}{1-\eta_w} W_t^{\eta_w\sigma} \frac{E_t \sum_{s=0}^{\infty} (\beta\theta_w)^s \left[\left(\frac{W_t}{W_{t+s}} \right)^{-\eta_w(1+\sigma)} N_{t+s}^{1+\sigma} \right]}{E_t \sum_{s=0}^{\infty} (\beta\theta_w)^s \left[\lambda_{t+s} \left(\frac{P_t}{P_{t+s}} \right) \left(\frac{W_t}{W_{t+s}} \right)^{-\eta_w} N_{t+s} \right]} \quad (\text{II.8})$$

which only depends on aggregate state variables. Equation II.8 has a very simple interpretation. If given the opportunity, the household resets its wage to the expected weighted sum of future marginal rates of substitution between consumption and labor.²⁴

III.B. Two-Tier Consumption Baskets

The household consumes an index of goods produced by two sectors $k = 1, 2$:

$$C_t = [\omega_1^\eta C_{1,t}^{\frac{\eta-1}{\eta}} + \omega_2^\eta C_{2,t}^{\frac{\eta-1}{\eta}}]^{\frac{\eta}{\eta-1}} \quad (\text{II.9})$$

There is a final good producer assembling this basket in a perfectly competitive manner. The price index associated with it is given by:

$$P_t = [\omega_1 P_{1,t}^{1-\eta} + \omega_2 P_{2,t}^{1-\eta}]^{\frac{1}{1-\eta}} \quad (\text{II.10})$$

and the demand for each basket of goods produced by sector k is given by:

$$C_{k,t} = \left(\frac{P_{k,t}}{P_t} \right)^{-\eta} C_t \omega_k \quad (\text{II.11})$$

²⁴Notice that if $\theta_w = 0$ then the optimal decision of the agent would be $W_t = \phi_L \frac{\eta_w}{\eta_w-1} \frac{N_t^\sigma}{(C_t - bC_{t-1})^{-\sigma}}$, i.e. marginal rate of substitution between labor and consumption equals the real wage, which is exactly what you would get in a frictionless benchmark.

In each sector $k = 1, 2$ there is a continuum of firms indexed by j who produce goods $C_{t,k,j}$. The individual goods produced by each firm are aggregated in a sectoral consumption basket according to the following aggregator:

$$C_{k,t} = \left[\int_0^1 C_{t,k,j}^{\frac{\eta_k-1}{\eta_k}} dj \right]^{\frac{\eta_k}{\eta_k-1}} \quad (\text{II.12})$$

The optimality conditions imply the following demands and price indexes:

$$C_{t,k,j} = C_{k,t} \left(\frac{P_{t,k,j}}{P_{k,t}} \right)^{-\eta_k} \quad (\text{II.13})$$

$$P_{k,t} = \left(\int_0^1 P_{k,t,j}^{1-\eta_k} dj \right)^{\frac{1}{1-\eta_k}} \quad (\text{II.14})$$

III.C. Goods Producers

Consider firm j in sector k . Firms maximize the present value of real dividends:

$$\max \quad E_0 \left[\sum_{t=0}^{\infty} m_{0,t} \frac{d_{t,k,j}}{P_t} \right] \quad (\text{II.15})$$

subject to:

$$\frac{d_{j,k,t}}{P_t} = \frac{P_{t,k,j}}{P_t} Y_{t,k,j} - \frac{w_t}{P_t} h_{t,k,j} \quad (\text{II.16})$$

$$Y_{t,k,j} = A_t h_{t,k,j} \quad (\text{II.17})$$

$$Y_{t,k,j} = \left(\frac{P_{t,k,j}}{P_{k,t}} \right)^{-\eta_k} Y_{k,t} \quad (\text{II.18})$$

$$a_{t+1} = \rho_a a_t + \sigma_a \epsilon_{t+1,a} \quad (\text{II.19})$$

Each period, only a fraction θ_p of firms are allowed to change their prices. Any reoptimizing firm solves the following problem:

$$\max_{P_{t,k,j}^*} E_0 \left\{ \sum_{t=0}^{\infty} (\beta \theta_p)^t m_{0,t} \frac{P_{t+1}}{P_t} \left[P_{t,k,j}^* Y_{t,k,j} - W_t h_{t,k,j} \right] \right\} \quad (\text{II.20})$$

Replace equation (II.17), (II.18) and make use of market clearing conditions to get:

$$\begin{aligned} \max_{P_{t,k,j}^*} E_0 \left\{ \sum_{s=0}^{\infty} (\beta \theta_p)^s m_{t,t+s} P_t \left[\left(\frac{P_{t,k,j}^*}{P_t} \right)^{1-\eta_k} \left(\frac{P_t^*}{P_{t+s}} \right)^{1-\eta_k} \left(\frac{P_{t+s,k}}{P_{t+s}} \right)^{\eta_k} \left(\frac{P_{k,t+s}}{P_{t+s}} \right)^{\eta} \right. \right. \\ \left. \left. \left(\frac{W_{t+s}}{P_{t+s}} \right) \left(\frac{1}{A_{t+s}} \right) \left(\frac{P_{t,k,j}^*}{P_t} \right)^{-\eta_k} \left(\frac{P_{k,t+s}}{P_{t+s}} \right)^{-\eta} \left(\frac{P_t}{P_{t+s}} \right)^{-\eta_k} \right] Y_{t+s} \right\} \end{aligned} \quad (\text{II.21})$$

Take first-order condition w.r.t. $P_{t,k,j}^*$ to get the firm's optimal reset price:

$$\frac{P_{t,k,j}^*}{P_t} = \frac{\eta_k}{\eta_k - 1} \frac{E_0 \sum_{s=0}^{\infty} (\beta \theta_p)^s m_{t,t+s} \left(\frac{W_{t+s}}{P_{t+s}} \right) \left(\frac{1}{A_{t+s}} \right) \left(\frac{P_{t+s,k}}{P_{t+s}} \right)^{\eta_k} \left(\frac{P_{k,t+s}}{P_{t+s}} \right)^{-\eta} \left(\frac{P_t}{P_{t+s}} \right)^{-\eta_k} Y_{t+s}}{E_0 \sum_{s=0}^{\infty} (\beta \theta_p)^s m_{t,t+s} \left(\frac{P_{t+s,k}}{P_{t+s}} \right)^{\eta_k} \left(\frac{P_{k,t+s}}{P_{t+s}} \right)^{-\eta} \left(\frac{P_t}{P_{t+s}} \right)^{1-\eta_k} Y_{t+s}} \quad (\text{II.22})$$

Economically, this means that a firm that is allowed to reset its price sets a price that is a markup $\eta_k/(\eta_k - 1)$ over future weighted marginal costs, where the weights are given by marginal utility of wealth times the probability of the firm having to stick with the price over the horizon.

Equilibrium Conditions and Taylor Rule

This section writes all the equilibrium conditions needed to solve the model. I start by normalizing all prices by the aggregate price level P_t . Therefore, I define: $p_{k,t} \equiv \frac{P_{k,t}}{P_t}$ for $k \in \{1, 2\}$, $w_t \equiv \frac{w_t}{P_t}$ and $r_t \equiv \log(R_t)$. Also, define H_t as aggregate labor supplied by households and L_t as aggregate labor demanded by firms. Define wd_t , $pd_{1,t}$, $pd_{2,t}$, pd_t as aggregate wage dispersion, sector 1 and 2 price dispersion and aggregate price dispersion respectively.

The variables in the model are: $p_{t,1}^*$, $F_{t,1}^p$, $K_{t,1}^p$, $p_{t,1}$, $pd_{t,1}$, $p_{t,2}^*$, $F_{t,2}^p$, $K_{t,2}^p$, $p_{t,2}$, $pd_{t,2}$, λ_t , a_t , Y_t , π_t , C_t , w_t , w_t^* , F_t^w , K_t^w , π_t^w , wd_t , pd_t , r_t , l_t , h_t , u_t , Ω_t .

Firms:

$$p_{t,1}^* = \frac{\eta_1}{\eta_1 - 1} \frac{F_{t,1}^p}{K_{t,1}^p} \quad (\text{II.1})$$

$$F_{t,1}^p = \lambda_t w_t \left(\frac{1}{A_t} \right) \left(p_{t,1} \right)^{\eta_1} \left(p_{t,1} \right)^{-\eta} Y_t + \beta \theta_1 \left(\pi_{t+1} \right)^{\eta_1} F_{t+1,1}^p \quad (\text{II.2})$$

$$K_{t,1}^p = \lambda_t \left(p_{t,1} \right)^{\eta_1} \left(p_{t,1} \right)^{-\eta} Y_t + \beta \theta_1 \left(\pi_{t+1} \right)^{\eta_1 - 1} K_{t+1,1}^p \quad (\text{II.3})$$

$$p_{t,1}^{1-\eta_1} = (1 - \theta_1) \left(p_t^* \right)^{1-\eta_1} + \theta_1 \left(\frac{p_{t-1,1}}{\pi_t} \right)^{1-\eta_1} \quad (\text{II.4})$$

$$pd_{t,1} = (1 - \theta_1) \left(p_t^* \right)^{-\eta_1} \left(p_1 \right)^{\eta_1} + \theta_1 \left(pd_{t-1,1} \right) \left(p_{t-1,1} \right)^{-\eta_1} \left(\pi_t^{\eta_1} p_1^{\eta_1} \right) \quad (\text{II.5})$$

$$p_{t,2}^* = \frac{\eta_2}{\eta_2 - 1} \frac{F_{t,2}^p}{K_{t,2}^p}. \quad (\text{II.6})$$

$$F_{t,2}^p = \lambda_t w_t \left(\frac{1}{A_t} \right) \left(p_{t,2} \right)^{\eta_2} \left(p_{t,2} \right)^{-\eta} Y_t + \beta \theta_2 \left(\pi_{t+1} \right)^{\eta_2} F_{t+1,2}^p \quad (\text{II.7})$$

$$K_{t,2}^p = \lambda_t \left(p_{t,2} \right)^{\eta_2} \left(p_{t,2} \right)^{-\eta} Y_t + \beta \theta_2 \left(\pi_{t+1} \right)^{\eta_2 - 1} K_{t+1,2}^p \quad (\text{II.8})$$

$$p_{t,2}^{1-\eta_2} = (1 - \theta_2) \left(p_t^* \right)^{1-\eta_2} + \theta_2 \left(\frac{p_{t-1,2}}{\pi_t} \right)^{1-\eta_2} \quad (\text{II.9})$$

$$pd_{t,2} = (1 - \theta_2) \left(p_t^* \right)^{-\eta_2} \left(p_2 \right)^{\eta_2} + \theta_2 \left(pd_{t-1,2} \right) \left(p_{t-1,2} \right)^{-\eta_2} \left(\pi_t^{\eta_2} p_2^{\eta_2} \right) \quad (\text{II.10})$$

Aggregate Stuff:

$$1 = \sum_{k=1}^2 \omega_k \left\{ (1 - \theta_k) \left(p_{t,k}^* \right)^{1-\eta} + \theta_k \left(p_{t-1,k}^{1-\eta} \right) \left(\pi_t^{\eta-1} \right) \right\} \quad (\text{II.11})$$

$$pd_t = \omega_1 p_1^{-\eta} pd_{t,1} + \omega_2 p_2^{-\eta} pd_{t,2} \quad (\text{II.12})$$

Household Stuff:

$$(w_t^*)^{1+\eta_w \sigma} = \phi_l \frac{\eta_w}{\eta_w - 1} \left(w_t \right)^{\eta_w \sigma} \frac{F_t^w}{K_t^w} \quad (\text{II.13})$$

$$F_t^w = H^{1+\sigma} + \beta_t \theta_w E_t \left(\pi_{t+1}^w \right)^{\eta_w (1+\sigma)} F_{t+1}^w \quad (\text{II.14})$$

$$K_t^w = \lambda_t H_t + \beta_t \theta_w E_t \frac{(\pi_{t+1}^w)^{\eta_w}}{\pi_{t+1}} K_{t+1}^w \quad (\text{II.15})$$

$$w_t^{1-\eta_w} = (1 - \theta_w) \left(w_t^* \right)^{1-\eta_w} + \theta_w \left(w_{t-1}^{1-\eta_w} \pi_t^{\eta_w-1} \right) \quad (\text{II.16})$$

$$\pi_t^w = \frac{w_t}{w_{t-1}} \pi_t \quad (\text{II.17})$$

$$wd_t = (1 - \theta_w) \left(w_t^* \right)^{-\eta_w} \left(w_t^{\eta_w} \right) + \theta_w \left((\pi_t^w)^{\eta_w} wd_{t-1} \right) \quad (\text{II.18})$$

$$\lambda_t = (C_t - bC_{t-1})^{-\gamma} \quad (\text{II.19})$$

$$1 = \beta R_t E_t \left[\frac{\lambda_{t+1}}{\pi_{t+1} \lambda_t} \right] \quad (\text{II.20})$$

Market Clearing conditions:

$$Y_t = \frac{AL_t}{pd_t wd_t} \quad (\text{II.21})$$

$$C_t = Y_t \quad (\text{II.22})$$

$$L_t = H_t \tag{II.23}$$

Stochastic processes and Taylor rule:

$$a_t = \rho_a a_{t-1} + \sigma_a \epsilon_t^a \tag{II.24}$$

$$r_t = \phi_\pi \pi_t + \phi_y \log \frac{Y_t}{Y_{t-1}} + \log\left(\frac{1}{\beta}\right) + \epsilon_t^r \tag{II.25}$$

$$u_{t+1} = \rho_u u_t + \sigma_u \epsilon_t^u \tag{II.26}$$

$$\Omega_{t+1} = \rho_\Omega \Omega_t + \sigma_\Omega \epsilon_t^\Omega; \tag{II.27}$$

Steady State

From equations (II.1)-(II.3), we have that:

$$p_{ss,1} = \frac{\eta_1}{\eta_1 - 1} w_{ss} \equiv \mu_{ss,1} w_{ss} \quad (\text{II.28})$$

From equations (II.6)-(II.8), we have that:

$$p_{ss,2} = \frac{\eta_2}{\eta_2 - 1} w_{ss} \equiv \mu_{ss,2} w_{ss} \quad (\text{II.29})$$

Combine the two equations above with equation (II.11) to get:

$$w_{ss} = \left(\frac{1}{\omega_1 \mu_{ss,1} + \omega_2 \mu_{ss,2}} \right)^{\frac{1}{1-\eta}} \quad (\text{II.30})$$

Now combine equations (II.13)-(II.15) and (II.19) to get:

$$\phi_l H_{ss}^\sigma \frac{\eta_w}{\eta_w - 1} - w_{ss} s (C_{ss} - b C_{ss})^{-\gamma} \quad (\text{II.31})$$

Use the market clearing condition in the steady state:

$$C_{ss} = \frac{H_{ss}}{p d_{ss}} \quad (\text{II.32})$$

The two equations above can be solved numerically for h_{ss} and C_{ss} . Then all the other steady state values follow directly.

AIII. Model Calibration and Aggregate Moments

In this section, I discuss the model fit as well as the calculation of the model’s empirical counterparts. Table A1 presents key macroeconomic statistics from the data and the model. The model is calibrated at a quarterly frequency and all statistics are annualized. The series for realized real consumption volatility is calculated using the non-durables and services consumption data series from the BEA and following the Beeler and Campbell (2009) computation method. I start by fitting an AR(1) process to the consumption growth series:

$$\Delta c_t = \gamma_0 + \gamma_1 \Delta c_{t-1} + u_t \quad (\text{III.1})$$

Then the annual realized volatility is computed as $V_{t,t+4} = \sum_{j=0}^{4-1} |u_{t+j}|$. The inflation time series is computed using the GDP implicit price deflator from the BEA (*series mnemonic: A191RD*).

To calculate per-capita hours in the data, I follow Smets and Wouters (2007) and Herbst and Schorfheide (2015). I take the index of average weekly nonfarm business hours (*FRED mnemonic / BLS series PRS85006023*) and call it $HOURS_t$. Then I take the number of employed civilians (*FRED mnemonic CE160N*) and call it EMP_t . Finally, as a measure of population, I use the quarterly average of the Civilian Non-institutional Population (*FRED mnemonic CNP160V*). Change in per-capita hours is then computed as:

$$\Delta l = \Delta \ln \left(\frac{HOURS_t EMP_t}{POP_t} \right) \quad (\text{III.2})$$

To compute frequency of price adjustment in the Amazon data, I follow the standard in the literature and identify regular prices using an algorithm discussed in Kehoe and Midrigan (2015) and Midrigan (2011). The algorithm is based on the idea that if a price change is regular then the store charges that same price frequently in a window around it. In practice, the algorithm works in 4 steps:

1. For each period t compute the modal price p_t^M in a window $[t - l, \dots, t, \dots, t + l]$
2. For t_1 set the regular price equal to the modal price
3. For each subsequent period:

- (a) if Amazon charges the modal price in that period and if 1/3 of prices in that window are equal to the modal price set the regular price equal to the modal price
 - (b) otherwise set the regular price equal to the last period regular price
4. Finally, one needs to eliminate regular price changes that occur in the absence of changes in Amazon's actual price. Replace the regular price with the current or previous regular price if they coincide.

Given the regular price series p^R , I then calculate frequency and duration of price adjustments. I first obtain the daily frequency per individual good, by calculating the number of daily price changes over the total number of valid observations for a particular good, then I compute implied duration using $-1/\ln(1 - frequency)$, and convert them to quarterly durations for calibration purposes. This is the same method used by Cavallo (2016).

The steady state markups on the high- a low-elasticity sectors are calibrated to match the first and last decile of the markup distribution using the same dataset as Epifani and Gancia (2011). Epifani and Gancia (2011) use a high-quality database on industry-level inputs and outputs, covering roughly 450 US manufacturing (4-digit SIC) industries for the period 1958 - 1996. I update their data until 2011 and compute the empirical distribution of price-cost margins. Price-cost margins are computed as the value of shipments (adjusted for inventory change) less the cost of labor, capital, materials and energy, divided by the value of shipments. I then take the first and last decile of this markup distribution and set the elasticity parameters in my model to match those values.

Finally, I estimate the equity premium using S&P 500 index returns during the post-war sample, using data from CRSP. Stock returns are continuously compounded including dividends. The risk-free rate is the treasury-bill rate that I take from the Welch and Goyal (2007) dataset. To compute the dividend price ratio, I use 12-month moving sums of dividends paid on the S&P 500.

Table A1: Macroeconomic Moments

This table presents statistics for key macroeconomic variables from the data and the model. The model is calibrated at a quarterly frequency and all the statistics are annualized.

	Model	Data
$\sigma(\Delta c)$	0.011	0.014
$\sigma(\pi)$	0.028	0.023
AC1(Δc)	0.242	0.239
AC1(π)	0.706	0.760
$\sigma(\Delta l)/\sigma(\Delta c)$	2.182	2.630
Frequency of Price Adjustment (quarters)	4.49	4.93
SS Markup low-elasticity sector	0.08	0.08
SS Markup high-elasticity sector	0.50	0.46
Equity Premium	0.05	0.08
Sharpe Ratio	0.31	0.52
E(P/D)	33.40	36.21
$\sigma(P/D)$	0.16	0.28

AIV. Keepa Data

The micro product data on Amazon products is from **Keepa**. **Keepa GmbH** is a German based firm that maintains price histories for almost all products on Amazon, the worlds largest and most trusted online merchant. Users of **Keepa** can individually track the price development of products that interests them, and **Keepa** will notify them when the price has reached a predetermined threshold. As of 2018, **Keepa** tracks more than 500 million products from Amazon US, UK, Germany, Japan, France, Canada, China, Italy, Spain, India, Mexico, Brazil and Australia. More than half of products tracked are from Amazon US. **Keepa** started operating in 2011, the date when my sample starts.

Keepa allows any user to track the price of any product through a graphic interface and charges a fee to individuals and businesses who need to download the data. The **Keepa** database is accessed through a proprietary Application Programming Interface (API). Access is controlled and charged according to the number of products per minute the user wants to download, with higher costs to download data for a larger set of products per minute. Therefore, unlike most standard financial databases, one cannot download the full database at once. **Keepa** requires the user to provide the API with a specific Amazon Standard Identification Number (ASIN), which uniquely identifies an Amazon product in order to download its data. Figure A2 shows the syntax of an API request to obtain data on a given product. To request a product, the user needs to specify three things: (i) the API key that grants access to the database, (ii) the market where the product is (US, Canada, France, etc.) and (iii) the ASIN of the product.

The difficult part is to setup a database of ASINs that belong to US public firms. Take for example the company *Fortune Brands Home Security* (ticker symbol: FBHS), a US public firm traded on the NYSE. Searching for the company name on Amazon yields almost no results. The reason for this is that Amazon does not store the name of the parent company that produces the goods it sells. Instead, Amazon just keeps track of the brand of the good.

Fortune Brands owns more than ten different brands and all of its brands are sold on Amazon. It is possible therefore to search on Amazon for a specific brand and from that search query get the ASINs that are needed to feed the **Keepa** API.

Thus, in order to obtain data on products sold by public firms on Amazon I need to catalogue all brands sold on Amazon and check whether they belong to a public firm. I proceed in three steps.

First I use the same list of companies as Huang (2016) as a starting point. Huang (2016)

identified several companies that sold their products on Amazon and researched the impact of positive/negative product reviews on asset returns.

Second, I web-crawl Amazon to get a list of all the brands it sells. Under each product subcategory, Amazon stores a list of brands. Figure A3 shows a list of brands beginning with the letter “A” that Amazon sells under its *Baby & Child Care* product subcategory. Amazon has more than five-hundred product categories and under each category, it sells thousands of brands.²⁵ To compile a list of all such brands, I develop a Python algorithm that slowly web-crawls Amazon webpages and builds a database with all Amazon brands. I remove from that list all brands that sell less than 5 products on Amazon, as they are unlikely to be relevant and unlikely to belong to a public traded firm. Using the list of brands, I manually search them on the Global Brands Database from the World Intellectual Property Organization and on Google, to assign a parent company to each brand. This allows me to get a list of public companies and all the brands they sell on Amazon.

Third, I use the above list of public companies and brands and again employ a web-crawling algorithm to retrieve all the ASINs of products that Amazon sells under a specific brand/firm pair.

The final step is to download the data from *Keepa* using the list of ASINs. Each product request to the *Keepa* database yields a JSON (JavaScript Object Notation) response with the data for the requested product. Figure A4 shows the format of the JSON objects.²⁶ The most important fields from the JSON object for this study are the fields `csv(0)` and `csv3`. These two fields store the Amazon price history and Sales Rank history of the product. Each of these fields is a two-dimensional object where the first dimension has the time/minute of the observation and the second dimension the observation value. Importantly, *Keepa* only appends a new entry to the history array if the price/value changes and not every time they update the product. Given that Sales Rank changes much more frequently than prices, the Sales Rank vector always has a much larger dimension than the price dimension.

The overall sample has around 280 thousand products, an average (median) of 1,100 (237) products sold per firm (see Table 3 for details).

Close Substitutes

²⁵See <https://www.browsenodes.com/> for a full list of Amazon subcategories/nodes.

²⁶The full list of output fields from a *Keepa* product API Request can be found here: <https://keepa.com/#!discuss/t/product-object/116>.

In Section IV.C, I identify products that are close substitutes for those in my original *Keepa* sample. To do so, I pick each product from my sample and store its narrowest Amazon product category or node (i.e. the narrowest product category to which the product belongs). Amazon’s narrowest nodes are very precise and therefore all products in such a node are very similar (substitutes). *Keepa* allows the user to make a “best sellers query” for a given product category, i.e. given an Amazon product category, an API best seller query returns up to 500,000 ASINs of products within that category.²⁷ Given that I am looking into very narrow categories, this limit is non-binding, as for any given product, the deepest Amazon category level always has less than 500,000 products. Given the ASINs from the best seller query, I can then use the standard product query to get the prices and sales ranks of products that are close substitutes.

The Keepa Canada Sample

The *Keepa* database is extremely powerful for identifying the price of the same product in different markets. Given that Amazon ASINs uniquely identify a product and are therefore the same across different markets, in order to obtain the price and sales rank of the products in my sample but on the Amazon Canada market, I simply need to modify the API request to the Canadian market instead of the US market. If a given product in my US sample is not sold on Amazon Canada, the API request will return an empty value.

²⁷Details of this query can be found here: <https://keepa.com/#!discuss/t/request-best-sellers/1298>

Figure A2: Keepa API Request

This figure shows how to make a Keepa API request. Requests can be made using a browser or programming languages suited for back-end web development, such as Python or Java.



Keepa 

Product Request

Token costs: 1 per ASIN

Retrieves the [product object](#) 2.2k for the specified ASIN and domain. If our last update is older than 1 hour it will be automatically refreshed before delivered to you to ensure you get near to real-time pricing data.

You can request products via either their ASIN (preferred) or via UPC and EAN codes. You can not use both parameters, `asin` and `code`, in the same request. Keepa can not track Amazon Fresh and eBooks.

Query:

```
/product?key=<yourAccessKey>&domain=<domainId>&asin=<ASIN> [or] &code=<productCode>
```

- `<yourAccessKey>`
Your private API key.
- `<domainId>`
Integer value for the Amazon locale you want to access. Valid values:
[1: com | 2: co.uk 12 | 3: de | 4: fr | 5: co.jp | 6: ca | 7: cn | 8: it | 9: es | 10: in | 11: com.mx | 12: com.br | 13: com.au]
- `<ASIN>`
The ASIN of the product you want to request.
- `<productCode>`
The product code of the product you want to request. We currently allow UPC, EAN and ISBN-13 codes. For batch requests a comma separated list of codes (up to 100). Multiple ASINs can have the same product code, so requesting a product code can return multiple products.

Figure A3: Amazon Brands

This figure shows the list of brands sold on Amazon under the product category *Health & Personal Care*, subcategory *Baby & Child Care* that begin with the letter “A”.

Health & Personal Care : Baby & Child Care : A

Viewing: Top Brands # A B C D E F G H I J K L M N O P Q R S T U V W X Y Z

A (1)	All Funny Animals (1)	Apert (1)
A - Z (1)	All Funny Animals (8)	APC (2)
A Little Lovely Company (3)	AE Naturals (1)	Apoc (3)
A Real Cool World (1)	ABJunior (2)	Apex Tool Group (1)
A to Z Bath Toys (2)	ALLCAMP (4)	Aphogee (1)
AB&D (1)	Allcura (1)	API (1)
ABD Medical (1)	Alayne Puel (4)	APIEU (1)
ADomo (2)	aller-mates (1)	APIG (2)
A+ (1)	Allergika GmbH (1)	Apivita (7)
A-Derma (1)	Allergy Kids (2)	APOLINE (2)
A-goo (2)	Allerhand (4)	Aposan (1)
A-logic (3)	alles-mehle.de GmbH (2)	Apostle (1)
A.C. Milan (2)	Allga San (1)	Apothecary 87 (1)
A.D. Sutton (1)	ALLIANCE HEALTHCARE IT.DIG.SpA (2)	Apotheke (2)
A.Mistry Ltd (4)	Allis (2)	Apple & Bears (3)
AZZ 4 Kids (7)	Allice (1)	Apple Park (8)
AZZ Home Solutions (1)	Allkindathings (4)	AppleCheeks (7)
A4DB (1)	ALLPHARM Vertriebs GmbH (1)	Applesauce (1)
A99 Golf (1)	AllRight (2)	APRAMO (2)
A&S Skincare (1)	AlTopBargains (1)	Aprica (1)
AAA-A (2)	Allure (1)	Aprilaire (1)
AABuild (1)	Allegre (1)	Apron (2)
AAE CAVALIER INC (2)	Almirall (4)	Aprons (1)
AARDVARK (2)	Almon (2)	APROPOS (2)
AI (1)	Almidea (2)	Aptamil (27)
AI Baby (2)	Alo (1)	Aqua Leisure (2)
Ababalaya (1)	Aloe Dent (4)	Aqua Sphere (1)
Ababy (2)	Aloe Farms (1)	AquaTresh (8)
abbyfrank (1)	Aloe Plus (5)	AquaSlide (1)
ABC (2)	Aloe Plus Lanzarote (4)	Aqualion (1)
ABC Design (9)	Aloe Pura (1)	Aquamax (1)
ABEDOE (1)	Aloe Vesta (1)	Aquadet (4)
Abena (2)	Alendia (2)	AquaShor (4)
ABI Kids (1)	Alpaca Annie (1)	AquaPlay (1)
Ability Superstore (1)	Alpen (1)	Aquasafe-Coronatio... (3)
AbleGrow (2)	Alphanova (17)	Aquascale (2)
ABM (1)	AlphaPeach (1)	AQUARTICA (1)

Figure A4: Output of Keepa API Product Request

This figure shows the format of the JSON objects that are returned by Keepa when making a product request for a specific ASIN.



Keepa

About:

The product object contains all of our price history data and basic product information.

Returned by:

The product object is returned by the following requests:

[Product Searches](#) 249 and [Product Request](#).

Important:

Always evaluate the *productType* field first. This field determines what data for the product is available.

Format:

```
{
  "productType": Integer,
  "asin": String,
  "domainId": Integer,
  "title": String,
  "trackingSince": Integer,
  "listedSince": Integer,
  "lastUpdate": Integer,
  "lastRatingUpdate": Integer,
  "lastPriceChange": Integer,
  "lastEbayUpdate": Integer,
  "imagesCSV": String,
  "rootCategory": Integer,
  "categories": Long array,
  "categoryTree": Object array,
  "parentAsin": String,
  "variationCSV": String,
  "frequentlyBoughtTogether": String array,
  "eanList": String array,
  "upcList": String array,
  "mpn": String,
  "hasReviews": Boolean,
  "type": String,
  "manufacturer": String,
  "brand": String,
  "label": String,
  "department": String.
```

AV. Additional Statistics and Tests

Table A2: Financial Properties of Elasticity Sorted Portfolios

This table reports several financial variables of portfolios sorted on demand elasticity. Tobin's Q is the inverse of the ratio the market value of equity plus book value of debt and preferred equity divided by book value of assets. Investment rate are the net capex expenditures divided by assets. Size is the market value of equity. R&D and SGA to assets are the ratio of research and expenditures and selling general and administrative expenses divided by total assets. Data are from 2011 to 2017.

	Less Elastic	2	3	4	More elastic	Difference (5-1)
Tobin's Q	0.434	0.400	0.472	0.417	0.391	-0.044**
Size	40.967	21.775	18.175	28.226	30.500	-10.467**
Investment Rate	0.039	0.036	0.038	0.035	0.040	0.001*
R&D to Assets	0.027	0.027	0.064	0.031	0.043	0.016
SGA to Assets	0.352	0.319	0.364	0.244	0.344	-0.008
Leverage	0.876	0.944	0.887	0.871	0.887	0.011

Table A3: Fama-French 5-Factors

This table reports results for Fama-French 5-factor model time-series regressions. I regress returns of the 5 elasticity-sorted portfolios and the return on the High-Low Elasticity Portfolio on the market, size, book-to-market, profitability and investment factors. The regressions are done with daily data, and the factors data is taken from the Kenneth French data library. The first row of Panel A reports the intercept of the regression in annual terms. The last three rows of this table compare the R^2 's of the CAPM, Fama-French 3-Factor model and Fama-French 5 factor model. The standard errors are reported in brackets. The sample ranges from January 2011 to April 2017.

	Less Elastic	2	3	4	More elastic
Annualized Alpha	-0.009	0.011	0.021	0.035	0.037
<i>Std. Error</i>	[0.011]	[0.009]	[0.010]	[0.009]*	[0.011]*
Rm-Rf	0.948	0.983	1.017	1.006	1.077
<i>Std. Error</i>	[0.011]***	[0.009]***	[0.010]***	[0.009]***	[0.011]***
Book-to-Market	0.453	0.249	0.438	0.323	0.381
<i>Std. Error</i>	[0.019]***	[0.016]***	[0.018]***	[0.015]***	[0.019]***
Size	-0.150	-0.116	-0.125	-0.104	-0.081
<i>Std. Error</i>	[0.024]***	[0.021]***	[0.023]***	[0.020]***	[0.025]***
Profitability	0.167	0.314	0.134	0.259	0.158
<i>Std. Error</i>	[0.031]***	[0.027]***	[0.030]***	[0.026]***	[0.032]***
Investment	0.123	0.287	0.227	0.137	0.203
<i>Std. Error</i>	[0.038]**	[0.034]*	[0.037]**	[0.031]**	[0.039]***
R^2 Fama-French 5 Factor model	0.886	0.903	0.905	0.923	0.902
R^2 Fama-French 3 Factor model	0.883	0.891	0.901	0.917	0.898
R^2 CAPM	0.843	0.881	0.869	0.899	0.877

Table A4: Double sorted portfolios on degree of price stickiness and demand elasticity

This table reports time-series averages for annual excess returns of double sorted portfolios on demand elasticity and the degree of price adjustment. Firms are assigned to one of five portfolios based on the average elasticity of demand they face for their products (columns). Portfolio low (high) has the stocks with the lowest (highest) demand elasticity. The sample is also split into portfolios with low degree of price adjustment (first row) and high degree of price adjustment (second row). The last column (row), reports the difference in returns between the portfolio with low demand-elasticity (degree of price adjustment) and the one with high demand-elasticity (degree of price adjustment). The sample ranges from January 2011 to March 2017.

	Less Elastic	2	3	4	More elastic	(5)-(1)
Low FPA	0.146**	0.154**	0.186**	0.192**	0.205**	0.048*
High FPA	0.121**	0.158**	0.153**	0.165**	0.189**	0.068*
(1)-(2)	0.026	0.004	-0.033*	-0.028	0.016	

Table A5: Panel Regressions of Demand Elasticity on Firm Net Income and Revenues

This table reports the results estimating equation (19). The regression is done using annual data. Standard errors are clustered at the firm level and t – stats are reported in parentheses.

	Net Income	Net Income	Net Income	Revenues	Revenues	Revenues
Elasticity Quintiles	-0.150	-0.148	-0.093	0.062	0.062	0.034
(t-stat)	(-3.62) ^{***}	(-3.57) ^{***}	(-2.19) ^{**}	(2.80) ^{**}	(2.80) ^{**}	(1.80) [*]
Year FE	N	Y	N	N	Y	N
Year x Industry FE	N	N	Y	N	N	Y
R^2	0.010	0.012	0.050	0.002	0.003	0.031
# Obs	1358	1358	1358	1597	1597	1597

AVI. Financial Data and Markups

Variable Construction

In this section, I discuss in more detail the construction of financial variables. The data on equity prices comes from the Center for Research in Security Prices (CRSP) monthly stock file. I manually merge the companies in my *Keepa* sample with the CRSP database. Balance sheet data comes from Compustat. I use the standard in the literature to compute measures of size and book-to-market (e.g. Lettau and Ludvigson (2001)). Market equity is the natural logarithm of firms' total market capitalization (price times number of shares outstanding) at the end of June. The ratio of book to market is calculated by dividing book equity at the last fiscal year end in the prior calendar year by the market equity at the end of December in the prior year. Market betas and Fama and French (1993) betas are the regression coefficient on the market factor and Fama and French (1993) factors excess returns in times-series regressions. I make use of the Kenneth French data library to get the returns on the SMB (Small Minus Big) and HML (High Minus Low) factors. The Herfindahl-Hirschman index (HHI) is computed as the sum of the squares of the market shares of the firms within a given industry, where I use net turnover of a firm as a proxy for its market size and define industry at the 4-digit Standard Industry Classification Code (SIC) of a firm.

As a proxy for markups of firms in my *Keepa* sample, I use gross margin as in Bustamante and Donangelo (2017) and Corhay, Kung, and Schmid (2017). Let the markup of firm i at time t be defined as:

$$Markup_{t,i} = \frac{\text{Value of Sales}_{t,i} - \text{Value of Costs of Goods Sold}_{t,i}}{\text{Value of Sales}_{t,i}} \quad (\text{VI.1})$$

where value of sales is the total value of revenue of firm i (Compustat annual data item REVT) and the value of the cost of goods sold is the Compustat annual data item COGS. I also define measures of markups at the industry level j , where the definition of markups used is the same as above but $\text{Value of Sales}_{t,j}$ is the sum of firms sales in industry j , and $\text{Value of Costs of Goods Sold}_{t,j}$ is the sum of costs of goods sold in industry j .

Markups and the Cross-Section of Returns

In the main text, I analyzed the relation between demand elasticity and stock returns, as well as the relation between demand elasticity and firm markups. In Section V.B of the paper, I show empirically that firms with higher demand elasticities have higher returns, and in Section V.C, I show that firms with higher demand elasticity have lower markups. One corollary of these two findings is that, in my sample, returns and markups have a negative correlation.

There is mixed evidence in the literature regarding the relation between markups and returns. Hou and Robinson (2006) document that firms in highly concentrated industries (higher markups) have lower returns even after controlling for size, book-to-market, momentum and other known return predictors. On the other hand, Kung (2015) finds, at the industry level, that high-markup industries are associated with higher expected returns. Finally, Bustamante and Donangelo (2017) study two different channels that affect the relation between markups and the cross-section of returns in opposite directions. On the one hand, they argue that firms with lower markups have a lower margin to buffer against negative shocks and therefore have a higher exposure to risk. On the other hand, incumbents in high-margin industries face a higher threat from entry, which makes them riskier and thus reduces their expected returns. They find support for both channels in the data.

Although it is beyond the scope of this paper to fully understand all the economic drivers of this relationship, Table A6 below provides some evidence on the time-series and cross-sectional relation between markups and stock returns. In this analysis I make use of the full Compustat universe of firms excluding financial industries (SIC codes from 6000 to 6999) and regulated industries (SIC codes from 4900 to 4999). For each year and firm (industry) I compute their markups using equation (VI.1) above. I then sort firms (industries) into terciles according to their markup and compute their expected returns the following year. The first (second) columns of Table A6 show the expected returns of firms (industries) with low and high markups using the entire Compustat time series. At the industry level, firms with higher markups have higher expected returns. The economic significance is the same as the one in Corhay, Kung, and Schmid (2017). At the firm level in the full 1964-2017 sample, the pattern is the opposite. Firms with lower markups have marginally higher expected returns. The last two columns of the table repeat the same exercise, but using only the latter part of the sample (2011-2017). On this latter sample, both at the industry level and at the firm level, firms with lower markups

have higher returns (the effect is more pronounced at the industry level). The direction of this correlation is the same as the one I find when I analyze how demand elasticity impacts on both returns and markups. The possible structural break on the relation between markups and returns is an empirically interesting question for future research.

Table A6: Markups and the Cross-section of Returns

This table presents time-series and cross-sectional relations between returns and markups. For each Compustat firm and industry, I compute a measure of markups using equation (VI.1). I sort stocks (industries) into terciles according to their markup and compute the returns the following year. The first two columns show the relation between markups and returns for the overall Compustat sample (1964-2017). The first column shows the results at industry level and the second column the results at firm level. The last two columns of the table report the same thing, but using the same time series as the *Keepa* sample (2011-2017).

	Excess Returns			
	Full Sample (1964-17)		Keepa Sample (2011-17)	
	Industry level	Firm level	Industry level	Firm level
Low Markup	0.125	0.137	0.147	0.128
High Markup	0.136	0.132	0.099	0.108