

Demand Elasticities, Nominal Rigidities and Asset Prices

Nuno Clara*

February 26, 2018

Abstract

This paper examines the interactions between demand elasticity and nominal rigidities and their implication to firm fundamentals and asset prices. In a multi-sector new-keynesian model that firms facing more elastic demands bear higher risk due to the presence of nominal frictions. I develop a novel method to estimate demand elasticities at the firm level by using high frequency Amazon product data. Consistent with the model I find that firms facing more elastic demands have lower markups and earn a return premium of 6.2% compared to firms facing more inelastic demands.

*Department of Finance, London Business School.

I. Introduction

Marshall's (1890) seminal work developed the concept of demand elasticity: how much does demand vary in response to a price change. Since then the concept has been central in economics (e.g. Gali (1994), Raith (2003)). However, the recent new-keynesian asset pricing pricing is silent on the impacts of different elasticities of demand as it mainly assumes that firms face the same elasticity. As a consequence, we typically do not look into the relation between demand elasticity, nominal rigidities and asset prices. Nominal rigidities create operational leverage in firms and therefore create a role for demand elasticity to matter for firm fundamentals and cross-sectional asset pricing. 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. My contribution is to consider the relation between demand elasticities and frequency of price of price adjustment at the firm level and study their joint implications for firm fundamentals and asset prices.

I start by laying out a standard new-keynesian model where firms are heterogeneous in terms of the demand elasticity they face. I allow for a correlation between demand elasticities and the degree of nominal rigidities that firms have. The basic structure of the model is similar to Carvalho (2006), Nakamura and Steinsson (2010) and Weber (2015) in which firms face nominal rigidities in adjusting the prices. In my model firms have both heterogeneous degrees of price stickiness and face heterogeneous demand elasticities. I calibrate the benchmark model using standard parameters in the literature and calibrate the correlation between sticky prices and demand elasticity to a value close to zero which I estimate in the data. This parameter is important as firms with higher demand elasticities should in principle adjust prices more often so that their degree of operational leverage is lower. My benchmark calibration generates a 7% returns spread between firms facing high elasticities of demand versus firms facing more inelastic demands. The mechanism is fairly simple: when faced with a shock, firms optimal prices change and, all else equal, optimal prices 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 expected returns.

The model also predicts that firms in the high elasticity sector to have lower equilibrium markups. This is an important prediction as it allows me to distinguish my model mechanism from the mechanism in a model where heterogeneity is only in the degree of nominal rigidities such as Weber (2015). In fact, 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 these type of firms fear more selling at a loss and therefore charge slightly higher markups. This prediction is at odds with what a model with heterogeneous demand elasticities would deliver.

To test the model predictions I firm-level need estimates of demand elasticity. There are challenges in estimating demand elasticities. First it is difficult to obtain them due to the standard endogeneity problem. In general, firms' decisions to change prices are endogenous and therefore empiricists only observe equilibrium prices and quantities, which makes it difficult to estimate the slopes that generated the equilibrium outcomes. This challenge can be overcome by either using instruments to trace out demand or to by relying on parametric assumptions regarding the shape of the demand (e.g. Berry, Levinsohn, and Pakes (1995) and Feenstra (1994)). Second it is not straightforward to estimate demand for a large cross-section of firms. It is difficult to find good instruments for a large cross-section of firms and therefore this challenge has no immediate solution.

One of the main empirical contributions of this paper is the estimation of demand elasticities for a large cross-section of firms. I use publicly available high-frequency micro-level product data (prices and quantities) provided by *Keepa*, one of the largest Amazon product trackers, for a very large number of products and firms. I address the identification challenges by looking into how quantity moves in a very narrow window around a price change. In particular, I measure quantity demanded right before the price change and see how quantity evolves within a short-time frame (12 hours) after the price has changed.

The identification strategy could fail if either aggregate demand is moving or if firms are quick to react to competitors price changes, which would trigger a shift in the demand curve. To ensure that this is not the case in these windows I ensure that there are no demand shifts nor demand shifts are expected and that therefore only the price shift is affecting quantity demanded. First, I exclude price changes that occur when there are predictably demand shifts (such as holidays and sales periods). Second, I verify that in such narrow windows there are

no competitors moving their prices. In fact, firms are indeed slow to adjust their prices when competitors change their prices. The degree of price synchronization across products that are close substitutes is around 5%, an order of magnitude lower than what the Calvo (1983) model would imply. In a Calvo model price changes are purely random and even in such a case the degree of price synchronization is higher than what is observable in the data.

The estimated elasticities have several reasonable economic properties. First, elasticities are mostly negative; as standard microeconomic theory would predict, an increase in the price of a good leads to a reduction in the demand. Second, consumer goods industries such as clothing and non-durables have a more inelastic demand than durable and manufacturing industries. 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 fully 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 paper is dedicated to showing that the results are robust. To avoid potential biases arising from increases in demand for specific products during holidays and sales periods (such as Christmas and Thanksgiving), I exclude from my sample the week surrounding these days and replicate the main result. The result is robust to the exclusion of these periods

and holds in a large out-of-sample period. Also, I show that the results are not driven by industry specific characteristics. Even within industries there is a 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. Further, I show using panel regressions that the relation between demand elasticity and the frequency of price adjustment is weakly positive but statistical insignificant, i.e. empirically there seems to be no relation between elasticity of demand and nominal rigidities. This is an important and puzzling finding that might help future researchers to explain heterogeneity in the degree of price adjustment.

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, sections V and VI show the empirical results and section VII 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 Consumer Price Index (CPI) dataset and scanner datasets (such as the Nielsen dataset) suffer from several biases. On one hand, price imputation and substitutions temporarily missing products in the CPI dataset does 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, biasing average price and quantity observed at week's end.

Monthly and weekly data collection that is 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* 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 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)) to trace out demand. Due to the difficulty in 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 estimate demand elasticities that makes use of high-frequency data.

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), who look at how price stickiness relates to asset prices. My study tries to take a step forward, by connecting very granular product level data to asset prices. In particular, I try to identify if different demand elasticities have implications for firm fundamentals and asset prices and show evidence regarding the relation between the degree of price stickiness and demand elasticities.

III. Framework

I solve a 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. To provide the basic intuition I start by qualitatively describing the mechanism in a static one-period model, but then develop a dynamic stochastic sectoral equilibrium model to evaluate the mechanism quantitatively.

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

A. *Static Model*

Consider a one period (two-dates) 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:

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

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.¹ It follows from the above demand 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:

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

Now consider a shock to either marginal costs c_i or aggregate demand Y . In an equilibrium model, the former could be motivated by a shift in aggregate productivity and the latter by a monetary policy shock. What would happen to the firm's profit if it cannot adjust its price? Figure 1 plots the profit loss due to price stickiness if a shock to c_i (Panel A) or a shock to Y (Panel B) 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:

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

The measure is plotted for several values of η_i . If there is no change in costs or demand there is no profit loss. This is the middle point in the graphs. If c moves then the profit losses start to increase (Panel A). The larger the shock, the further away firm i is from the optimal

¹This demand function can be micro-founded through a Dixit and Stiglitz (1977) aggregator over different varieties.

price and the bigger are its losses. The loss is more pronounced for cost increases than for cost decreases. The effect is an order of magnitude larger if the elasticity of demand is larger, meaning that for the same shock firms with larger 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 larger elasticity of demand are 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 such correlation.

B. Dynamic Model

The previous section showed in reduced form why nominal rigidities might lead to larger riskiness of firms with higher elasticities of demand. In this section I develop a quantitative neoclassical multi-sector equilibrium model in which sectors face different elasticities of demand and quantitatively evaluate differences in risk among firms in the different sectors. The model builds on the new-keynesian multisector models of Carvalho (2006), Nakamura and Steinsson (2010) and Weber (2015); therefore I lay down the main model equations here and leave details and derivations for the appendix. 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 both face different demand elasticities and different degrees of nominal rigidity.

I make use of the standard Dixit and Stiglitz (1977) aggregator to combine different types of consumption goods. This framework imposes a link between the elasticity of demand and the markup of prices over marginal costs, as in equilibrium the optimal firm markup is a function of the demand price elasticity perceived by each firm.

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

alternative mechanisms.

B.1. Households

There is a continuum of households indexed by $i \in [0, 1]$. Each household i has an utility function separable in a consumption bundle, C_t , and differentiated labour $n_{i,t}$. The representative agent utility function is given by:

$$E_t \sum_{s=0}^{\infty} \beta^s \left[u(C_{t+s} - \nu C_{t+s-1}) - \int_0^1 v(n_{t+s,i}) di \right] \quad (4)$$

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 $k \in \{1, 2\}$.² 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:

$$C_{k,t} = \left[\int_0^1 C_{t,k,j}^{\frac{\eta_k-1}{\eta_k}} dj \right]^{\frac{\eta_k}{\eta_k-1}} \quad k \in \{1, 2\} \quad (5)$$

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

$$\frac{\partial C_{t,k,j}}{\partial P_{t,k,j}} \frac{P_{t,k,j}}{C_{t,k,j}} = -\eta_k \quad (6)$$

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 equation is given by:

$$1 = \beta R_t E_t \left[\frac{1}{\pi_{t+1}} \frac{(C_{t+1} - bC_t)^{-\gamma}}{(C_t - bC_{t-1})^{-\gamma}} \right] \quad (7)$$

²The model can be easily generalizable to have an arbitrarily larger number of sectors

B.2. Wage Setting

I follow Erceg, Henderson, and Levin (2000) and Woodford (2013, chapter 4.1) and 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, labor used to produce each good is a CES aggregate of the continuum of types of labor supplied by the representative household:³

$$N_t \equiv \left[\int_0^1 n_{i,t}^{\frac{\eta_w-1}{\eta_w}} di \right]^{\frac{\eta_w}{\eta_w-1}} \quad (8)$$

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:

$$E_t \sum_{s=t}^{\infty} (\beta \theta_w)^{s-t} [\lambda_s w_{t,i} n_{s,i} - v(n_{s,i})] \quad (9)$$

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 labour 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.

B.3. Firms

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

$$Y_{t,k,j} = A_t n_{t,k,j} \quad k \in \{1, 2\} \quad (10)$$

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

³It 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.

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.⁴ 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:

$$\max_{P_{t,k,j}^*} E_t \left[\sum_{t=0}^{\infty} (\beta \theta_k)^t m_{0,t} [P_{t,k,j}^* Y_{t,k,j} - w_t n_{t,k,j}] \right] \quad (11)$$

subject to:

$$Y_{t,k,j} = \left(\frac{P_{t,k,j}^*}{P_{k,t}} \right)^{-\epsilon_k} Y_{t,k} (\omega_k)^{-1} \quad (12)$$

$$Y_{t,k,j} = A_t n_{t,k,j} \quad (13)$$

I close the model 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:

$$r_t = \phi_{\pi} \pi_t + \phi_y \log \frac{Y_t}{Y_{t-1}} + \log \left(\frac{1}{\beta} \right) + \epsilon_t^r \quad (14)$$

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 this model. In the context of the zero lower bound, the monetary policy shock should be interpreted as forward guidance shocks.

B.4. Calibration

I calibrate the model at quarterly frequency by taking parameter estimates from the literature.

Household Parameters

I set the households impatience level to 0.99, which implies a quarterly risk-free rate of 1%. The habit adjustment parameter, v , is 0.66, as estimated by Galí, Smets, and Wouters

⁴Up 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.

(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 duration of labour contracts 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).⁵ 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) there seems to be no relation between the degree of price stickiness and demand elasticity both at the firm and at the industry level, (ii) portfolios sorted on demand elasticity do not have any difference in terms of frequency of price adjustment. Consequently 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

⁵To back out these numbers I assume that markups are uniformly distributed.

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 , 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. I equally-weight the returns within each sector to estimate returns at the sector level.

B.5. Implications and Mechanism

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.

Mechanism

To understand the main forces at work, figure 2 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 price of goods in the sector with higher elasticity versus sectors with lower elasticity have now decreased and therefore demand is relatively higher.

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 is important to explain the heterogeneity in the risk of firms. Firms in the high elasticity sector pay more when marginal utility of wealth is lower and therefore are riskier. Panel (e) 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:

$$pd_{t,k} = \int_0^1 \left(\frac{P_{t,k,j}}{P_{t,k}} \right)^{-\eta_k} dj \quad (15)$$

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 away from the optimal price. This makes this sector riskier than the sector with a more inelastic demand. To see 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:

$$E_t[r_{t+1}^H - r_{t+1}^L] = -cov_t(m_{t,t+1}, d_{H,t+1} - d_{L,t+1}) \quad (16)$$

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 (16) can be written as:

$$\begin{aligned} E_t[r_{t+1}^H - r_{t+1}^L] &= -[(1 - \eta)cov_t(m_{t,t+1}, p_{H,t+1} - p_{L,t+1}) \\ &\quad + (\eta_H - 1)cov_t(m_{t,t+1}, \log(\mu_{H,t+1})) \\ &\quad - (\eta_L - 1)cov_t(m_{t,t+1}, \log(\mu_{L,t+1}))] \end{aligned} \quad (17)$$

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 2). But 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 now move to describe its implications, which can be directly tested in the data. The baseline model generates reasonable macroeconomic moments - the annual volatility of consumption growth is 1.1%, the annual volatility of inflation is 1.9%, the average price-dividend ratio and the volatility of the price-dividend ratio are 30.1 and 16%, both values in line with the estimates of Lettau and Wachter (2007).

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 perfectly. 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.

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 Wecken (2010). In

particular De Paoli, Scott, and Weenen (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 or an increase in the substitutability of goods across the two sectors increases the spread in returns among the sectors.

Firms with a higher degree of monopoly power have higher markups. 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 the results. The top panel reports the returns and markups of my model: in equilibrium the more elastic sector earns higher returns and has lower markups. The bottom panel reports the results for the two-sector version of the heterogeneity in sticky prices model. In line with the results of Weber (2015), sectors with a higher degree of price stickiness earn higher expected returns. In this 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 are taken 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.

The `Keepa` database stores several fixed attributes of the products it tracks: its name, category node⁶, Universal Product Codes (UPC), International Article Numbers (EAN), product brand, label, model, color and size, among other. `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.⁷ In contrast, the Nielsen Dataset and the BLS dataset have a weekly and monthly update frequency (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 Indicator, 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 exception are the Nielsen scanner datasets, which do have information on quantities. The `Keepa` dataset fills this void in the micro-level product datasets, having both quantity information and high frequency of update.

`Keepa` is one of the largest Amazon online price trackers available. It provides free access to interactive product price and sales-rank charts and 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,

⁶e.g. Home Care & Cleaning

⁷The update frequency of each product varies, depending on general interest in the product in question. If a product is on someone's wishlist, it will be updated at least once every hour. Most others are updated several times daily.

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.

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 use a two-tier approach to identify such firms. 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 and therefore are 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 use to into the *Keepa* database to download the data.

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 both in 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.

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

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 products per firm is 298, which reflects the fact that a few firms sell many more products than others (e.g., clothing firms sell many products whereas manufacturing firms sell less products). The average market capitalization of the firms in the sample is \$26 billion, as most of the firms that have their products sold by Amazon are large. They also have high average revenues of \$3 billion and high average EBITDA of \$700 million. This reflects that fact that most of these firms are not only large, but also mature and 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 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.

B. Amazon Prices and Sales

I use some examples to explain the nature of the data. Figure 3 plots screenshots taken from *Amazon.com* of 6 products on the sample. As the figure illustrates, Amazon sells a vast range of products: drills, sportswear, food, sprays, copying devices and sports watches among others. At each 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 3, the Fitbit Flex 2, a wristband produced by Fitbit, Inc., a US public firm traded on the New York Stock Exchange (NYSE). Panel A of Figure 4 shows the time-series of the original price p_t of 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 our specific example, the price is usually either

80.50 or 60 US dollars. At a higher frequency, prices experience a great deal more movement. 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 used 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, as 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, 25th November 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 the paper results are robust to the exclusion from the sample of weeks surrounding event days such as Black Friday, Christmas, New Years, 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 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:

$$\log(\text{SalesRank}) = c - \theta \log(\text{QuantitySold}) \tag{18}$$

where c and θ are a function of parameters of the power law distribution. They estimate equation (18) 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 (18)).

I rely on their result, 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 only refer to the difference between the concepts when appropriate.

The sticky prices literature places great emphasis on the distinction between high frequency and low frequency price movements, and most frequency 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 are referred to Midrigan's (2011) online appendix for further details.

The orange line in panel C of Figure 4 plots the Amazon Price for Fitbit and the blue dashed line plots the regular price. 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 and then 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: $-1/\ln(1 - \text{frequency})$. The second line of Panel A of Table 3 reports the frequency of price adjustment at the firm level. On average, firms adjust

their prices every 4.93 months. 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 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 with firms with more flexible prices. I test this hypothesis in the appendix of the paper and find that in the more recent sample this result does not hold. This is consistent with Weber (2015), who claims the pattern to be true only before 2007. Furthermore, I try to establish a relation between the elasticity of demand of a firm and the frequency with which firms alter their prices and fail to find such an association.

C. 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 such hypothesis, I need to rank firms according to the elasticity of demand they face. Having both individual product quantity and price data is not enough to identify elasticities due to the standard endogeneity problem: empiricists only observe equilibrium prices and quantities and therefore the slopes that generate the equilibrium outcomes are unknown.

I make use of the higher frequency *Keepa* data on prices and quantities to estimate elasticities. Given the nature of my data, I can measure the level of demand right before the price change and 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, supply shocks, or could be simply changing markups (as often occurs in sales). Shifts in the demand curve are problematic for this identification strategy. Thus for the identification strategy to be valid I need to ensure that in such a narrow window around the price change demand is not moving. There are two main reasons for demand of 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, electronics on the 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 and then 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. I remove from my sample weeks surrounding holidays and festivities. These are weeks where demand is likely to shift and therefore I remove from my sample weeks surrounding Christmas, New Year, Valentine's, Thanksgiving, Black Friday, Father's day and Mother's.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 Amazon product category *Electronics* and subcategories *Office Electronics* > *Presentation Products* > *Presentation Remotes*, therefore its narrowest product category is *Presentation Remotes*. Figure 5 illustrates a few products in this subcategory. All the products are close substitutes, have similar price tags and therefore in a frictionless market a change in the price of a product in such a narrow category should trigger a change in competitor's price.

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 6 plots an histogram with the probability of a change in price of a competitor when product i changes its price. I find that 50% of the time only 2.5% of competitors change their price in such a narrow window and 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. Formally, I define the synchronization

⁹The main results of the paper are not sensitive to using either the full sample or the sample with festivities included

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^c is the number of products in \mathbb{S} that changed their prices at any point in time t , then one can define the synchronization rate as:

$$sync_i = \frac{s_i^c - 1}{\#\mathbb{S}_i - 1} \quad (19)$$

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

The synchronization rate defined above ranges from one if all products change price at the same time, 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 and therefore firms cannot synchronize price changes. However, given that some price changes coincide in time the synchronization rate is not zero.

Figure 7 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 adjust its prices is only 4%.

This implies that firms are slow to react to demand shocks such as price changes and therefore this should not be an issue for the identification strategy outlined above. The Calvo synchronization rate is depicted in a blue line.¹² The Calvo synchronization rate always lies above the synchronization rate of the sellers in my *Keepa* sample, which is inline 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 panel of the figure plots the same synchronization rates but excluding the tempo-

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

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

¹²Given the evidence from Weber (2015) that different sectors might have different frequency of price adjustments, I compute the degree of stickiness at the sector level and average them.

rary 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 prices of competitors. Remember that the products are almost perfect substitutes and therefore should have a high degree of coordination in prices. 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. In February 2nd, 2017 at 8:15 a.m. UTC -4, Amazon changed the price of this product from 99 to 79.5 US dollars. Economic theory would predict that quantity sold of this product should increase immediately: a movement *along* the demand curve. Panel A of Figure 8 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 (many potential buyers of the product have price alerts and automated purchase decisions if prices go below a certain threshold). During the next 12 hours, the sales rank kept increasing and becoming closer to 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 8 plots the Amazon price and Sales Rank 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.

I now make use of the result of Chevalier and Goolsbee (2003) to argue that a change in sales rank is a good proxy for a change in quantity sold and compute elasticities of demand.

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

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:

$$\epsilon_{t \rightarrow t+w, i, f} = -\frac{\Delta \log SalesRank_{t \rightarrow t+w, i, f}}{\Delta \log p_{t \rightarrow t+w, i, f}} \equiv \frac{\Delta \log Q_{t \rightarrow t+w, i, f}}{\Delta \log p_{t \rightarrow t+w, i, f}} \quad (20)$$

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. I choose a window of 12 hours after the price change to estimate elasticities. The estimated elasticities and main results are also robust to the choice of a 8 and 16 hour window.

I average elasticities resulting from equation (20) at the product level and then at the firm level to get an estimate of the average demand elasticity per product and the average elasticity a firm faces.

V. Empirical Results

A. Demand Elasticity

Figure 9 plots a histogram of the estimated elasticities at the firm level. The average elasticity of demand faced by a firm is -0.1243. In line with what economic theory predicts, most firms in our sample face a negative elasticity, with significant heterogeneity across firms: the standard deviation of firms' elasticities is equal to 0.1852. There are a few firms in our sample with a positive estimate for elasticity. This is likely due to noise in the data, as only 0.8% of firms have positive estimates that are statistically different from zero at 5% significance level. Therefore, I interpret these estimates as being close to zero.

I categorize each firm in the sample according to the Fama-French 12-Industries classification 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 related to 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)).

All elasticities are statistically negative except for the more 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 10 plots the average elasticity per firm on the y -axis and the industry on the x -axis. Despite the large heterogeneity in elasticities across industries, there is still a large variation in elasticities within each industry and therefore the effects reported in this paper are robust to industry level controls. In the section below I highlight the potential concerns with the identification strategy and the potential bias of the estimated elasticities.

VI. Asset Prices

A. 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 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 demand a firm faces is quite stable over time (more on this later). Therefore, I do not rebalance portfolios but only sort them once to minimize concerns about measurement error in firm elasticity estimates, an approach similar to Weber (2015). In the robustness section, I show that the elasticity estimates are quite 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 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.

In my theoretical framework different exposures to the elasticity of demand are fully explained by systematic risk, i.e. CAPM betas line up perfectly with stock returns. To test this prediction, I perform standard time-series tests and regress the returns of the elasticity sorted portfolios on the 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 , $R_{m,t}^e$ the excess market return, and 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:

$$R_{i,t}^e = \alpha_i + \beta_i R_{m,t}^e + \Gamma_i X_t + u_{i,t} \quad (21)$$

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 portfolios excess returns. This implies that the intercept α_i of the time-series regression (21) should be zero. Panel A of table 6 shows the results. 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 the Fama-French model fully explains the returns.

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 spread in market betas implies that the risk underlying demand elasticity is fully captured by the CAPM. 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 (21) using only the market as a factor. The results are similar to the ones 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 festivities. These are days where there may be predictable changes in the prices of products. For example, retailers often discount their products on Black Friday and therefore if agents anticipate significant price changes during that day that might bias the elasticity estimates as demand would not be 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.¹⁴ 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% spread when holidays and festivities are not excluded).

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 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 if there is any relation between the two.

B. Panel Regressions and Implications for Markups

The model from section III.B predicts that firms with higher demand elasticities should 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:

¹⁴For Christmas and New Year I exclude the full set of days between those dates as well.

$$Y_{f,t} = \alpha + \beta_{\epsilon_f} \times \epsilon_f + \gamma_t + u_{t,i} \quad (22)$$

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.¹⁵ 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 5.¹⁶ 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 we use EBITDA margin. In the model the only costs of production are labor costs, which could potentially encompass not only wages but costs of intermediary goods. EBITDA margin encompasses all such costs (costs of goods sold, wages, and selling, general and administrative expenses). As we move from a firm with a low elasticity to a higher elasticity quintile, markups decrease by 5 percentage points. This yields a margin differential of around 25 percentage points between the lowest and highest quintile.

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. In the next section I try to establish the relation between the degree of price stickiness and elasticity of demand and make a few other robustness tests, such whether the results hold out-of-sample and whether elasticities are stable through time.

¹⁵The panel regression results are robust to the use of terciles, quintiles or deciles.

¹⁶Moving 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 $1.1\% \times 5$.

C. Robustness

C.1. Degree of Price Stickiness and Demand Elasticities

I study the relation between the degree of price stickiness and demand elasticities. I have shown that firms face heterogeneous price demand elasticities. Weber (2015) shows that firms also face heterogeneous degrees of price adjustment. If firms facing higher demand elasticities are riskier, then it would be plausible that these firms would change their prices more frequently. To test this relation formally I estimate the following regression:

$$\epsilon_i = \alpha + \beta_{FPA}FPA_i + \sum_n \beta_n X_{t,i,n} + \mu_t + \epsilon_{i,t} \quad (23)$$

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 9 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 the relation between price stickiness and elasticity. The relation is weakly positive. Panel B from 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. Once a set of standard controls from the literature are added to the panel regression above, the coefficient becomes insignificant. The last column of Table 9 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 VI.A: heterogeneity in demand elasticities is fully captured by differences in systematic risk.

C.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. The estimated portfolios were not rebalanced and therefore it is important

that this is indeed the case. I test this assumption using two different approaches. First, I set a cutoff date that exactly splits the sample in two identically sized time-series and estimate elasticities at the firm level on both samples. If elasticities are stable over time, there should be no differences between the two samples. Second, I make an out-of-sample exercise, by repeating the portfolio sorts done in section IV-A. using a different sample from the one with which I estimated the elasticities.

Figure 12 plots the results for the sample splits. The blue dots on the figure are 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 has 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).

If elasticities are indeed stable over time, then the return spreads from table 5 should be similar in an 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 still 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.

VII. Conclusion

I have developed a novel way of estimating demand elasticities at the firm-level by exploiting on high frequency data from 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 fully 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 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 on the micro-foundations of sticky prices. Standard Calvo 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.

References

- BANSAL, R., D. KIKU, AND A. YARON (2010): “Long run risks, the macroeconomy, and asset prices,” *The American Economic Review*, 100(2), 542–546.
- 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,” .
- CALVO, G. A. (1983): “Staggered prices in a utility-maximizing framework,” *Journal of monetary Economics*, 12(3), 383–398.
- 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.

- 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.
- 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.
- 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.
- 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.
- 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.
- 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 J. A. WACHTER (2007): “Why is long-horizon equity less risky? a duration-based explanation of the value premium,” *The Journal of Finance*, 62(1), 55–92.
- 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.
- OLIVEI, G., AND S. TENREYRO (2007): “The timing of monetary policy shocks,” *The American Economic Review*, 97(3), 636–663.
- 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.
- TAYLOR, J. B. (1993): “Discretion versus policy rules in practice,” in *Carnegie-Rochester conference series on public policy*, vol. 39, pp. 195–214. Elsevier.
- 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.

Figure 1: 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 in regards to the elasticity of demand parameter η . The parameters underlying the figure are: $c = 5$, $Y = 5$ and $\eta \in \{4, 5, 6\}$.




Figure 2: Impulse response function to a one-standard deviation positive productivity shock

y-axis are % deviations from steady state

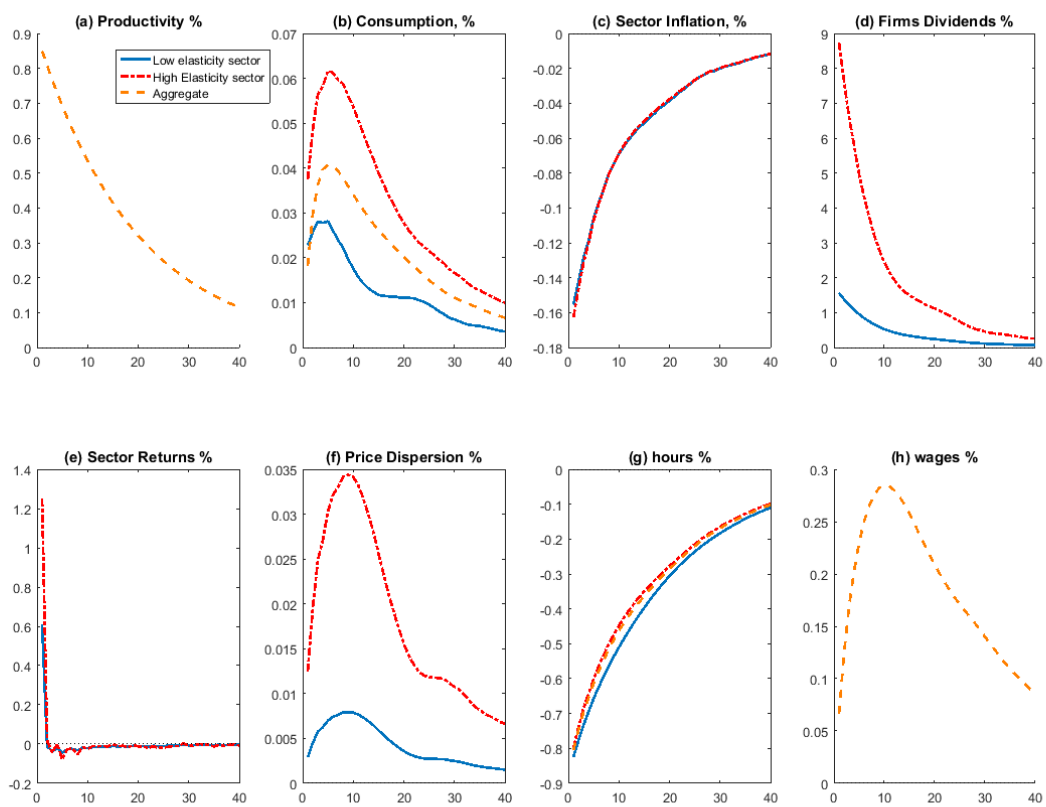


Figure 3: Examples of products in the dataset

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



BLACK+DECKER
BLACK+DECKER LDX120C 20-Volt MAX Lithium-Ion Cordless Drill/Driver
★★★★☆ * 3,679 customer reviews | 656 answered questions

Price: **\$47.89**

In Stock.

Packaging may reveal contents. Choose **Conceal Package** at checkout.



Columbia
Columbia Sportswear Rogue River Pack
★★★★☆ * 1 customer review | 3 answered questions

Price: **\$39.99** & **FREE Shipping**

Color: **Sea Sale**



Campbell's Homestyle Soup, Chicken Noodle, 15.4 Ounce (Pack of 8)
by Campbell's Homestyle

\$12.84 (\$0.10/Ounce)
Subscribe & Save

★★★★☆ * 414

Product Features
Classic home-style soup you can conveniently savor on the go

More options available:
\$13.52 **prime**
FREE Shipping on eligible orders



WD-40
WD-40 490057 Multi-Use Lubricant Smart Straw Spray 12 OZ (Pack of 12)
★★★★☆ * 116 customer reviews

List Price: ~~\$404.88~~
Price: **\$48.48** (\$4.04 / Can of WD-40)
You Save: **\$55.40** (52%)



Xerox
Xerox WorkCentre 3215/NI Monochrome Multifunction Printer
★★★★☆ * 98 customer reviews | 60 answered questions

Was: ~~\$445.97~~
Price: **\$109.99**
You Save: **\$35.98** (25%)

In Stock.

Packaging may reveal contents. Choose **Conceal Package** at checkout.



Fitbit Flex 2, Lavender (US Version)
by Fitbit

\$59.95 ~~\$99.95~~

prime
FREE Shipping on eligible orders

★★★★☆ * 3,382

Product Features
... pendants and bangles, The *Fitbit Flex 2* wri
made of a ...

More options available:
\$52.99 **Other Sellers**

Figure 4: Amazon Price, Sales Rank and Permanent Price

This figure plots price and sales rank time series for 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.




Figure 5: Close substitutes - example

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




Figure 6: Probability of a competitors' price change in the 12 hour windows

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




Figure 7: 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 (19). 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 panel plots the synchronization rates without filtering for temporary price changes and the right panel plots the synchronization rate excluding temporary price changes. The permanent price changes are computed using the algorithm of Midrigan (2011).

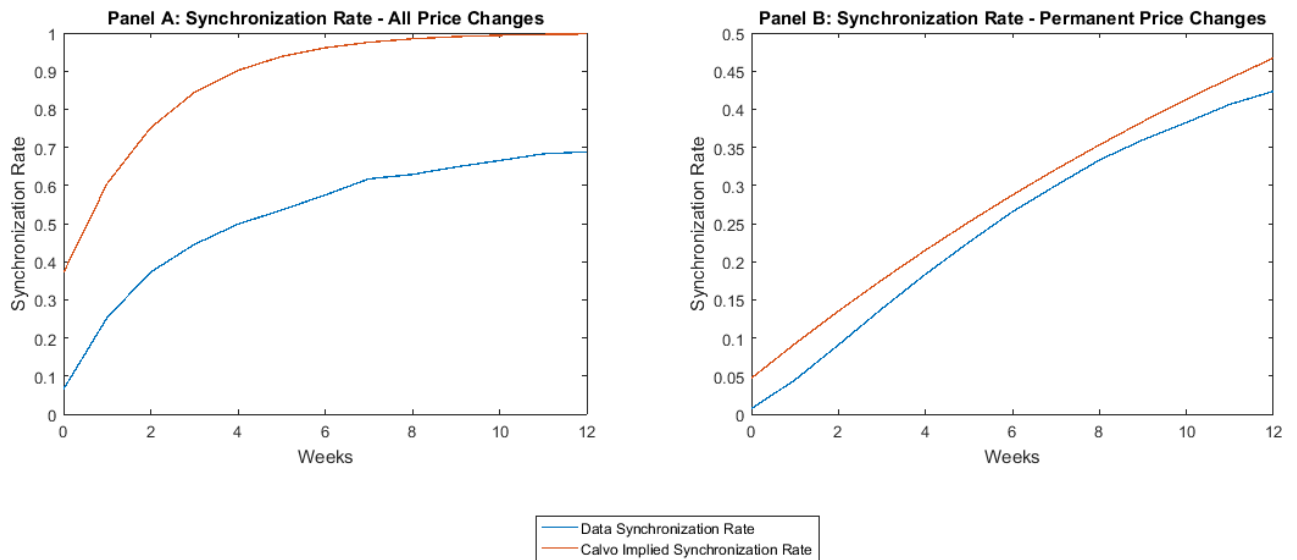


Figure 8: High Frequency Amazon Price and Sales Rank

This figure plots high frequency price and sales rank for 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.

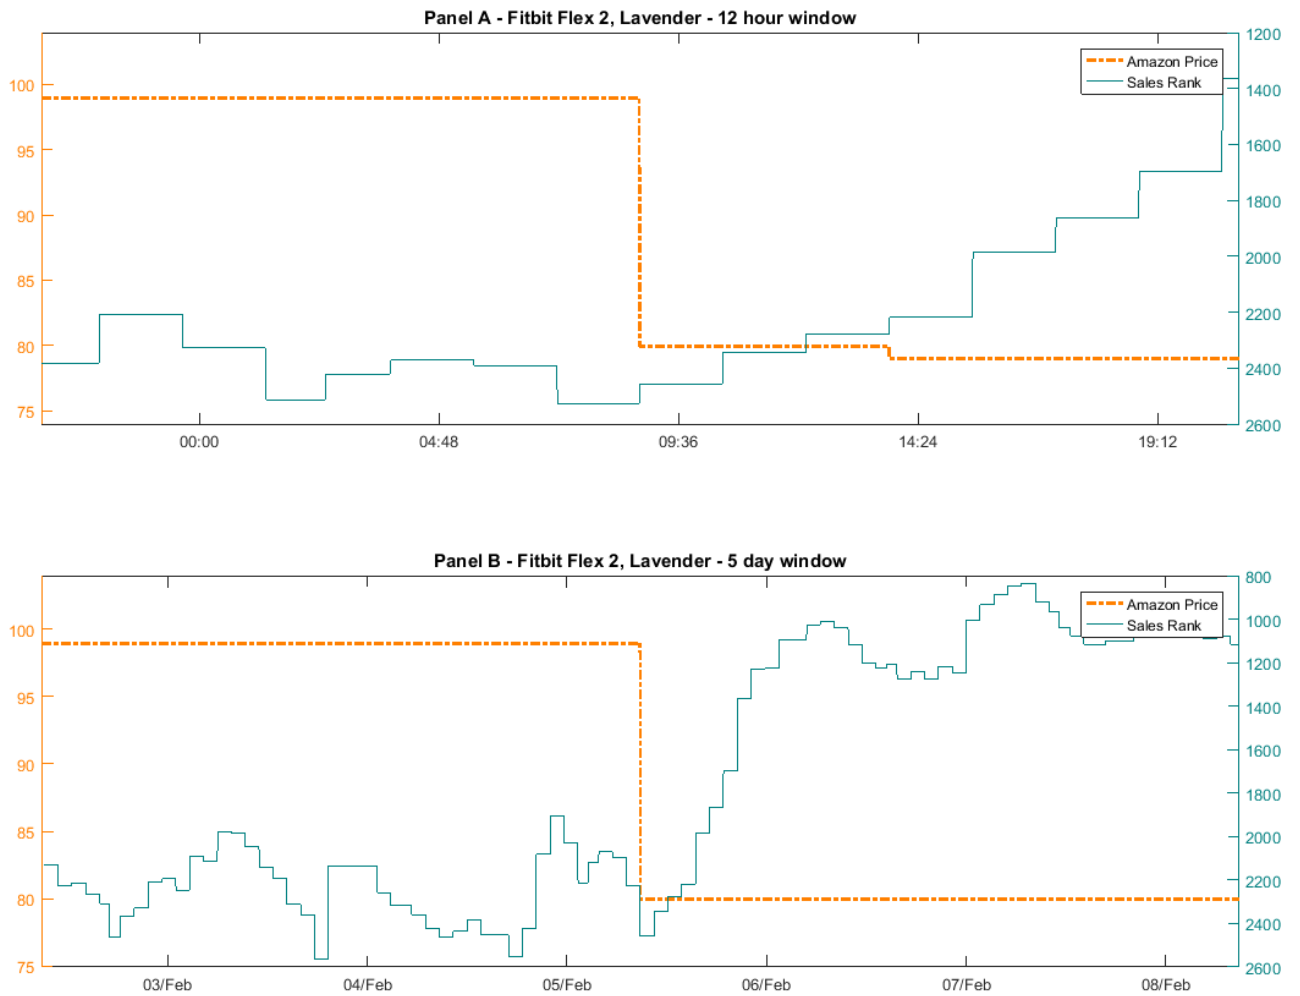


Figure 9: Average Elasticity Per Firm

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 20 and then average across products and firms to get an average elasticity per firm. The sample with which we compute elasticities ranges from February 2015 to April 2017.




Figure 10: Average Elasticities by firm - industry level

This scatter plot shows the average elasticities by firm (y -axis) conditional on the industry the firm is in (x -axis). Firms are classified into industries according to the Fama-French 13-industry classification. The elasticities are computed using equation 20 and the *Keepa* data.

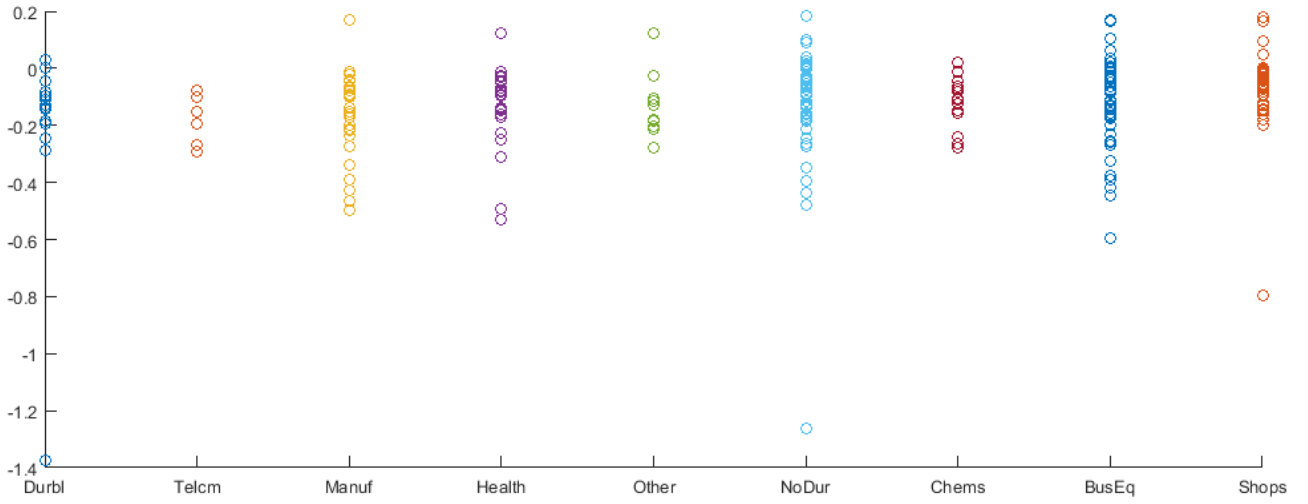


Figure 11: Cross-sectional relation between price stickiness and demand elasticity

This figure shows illustrating 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.

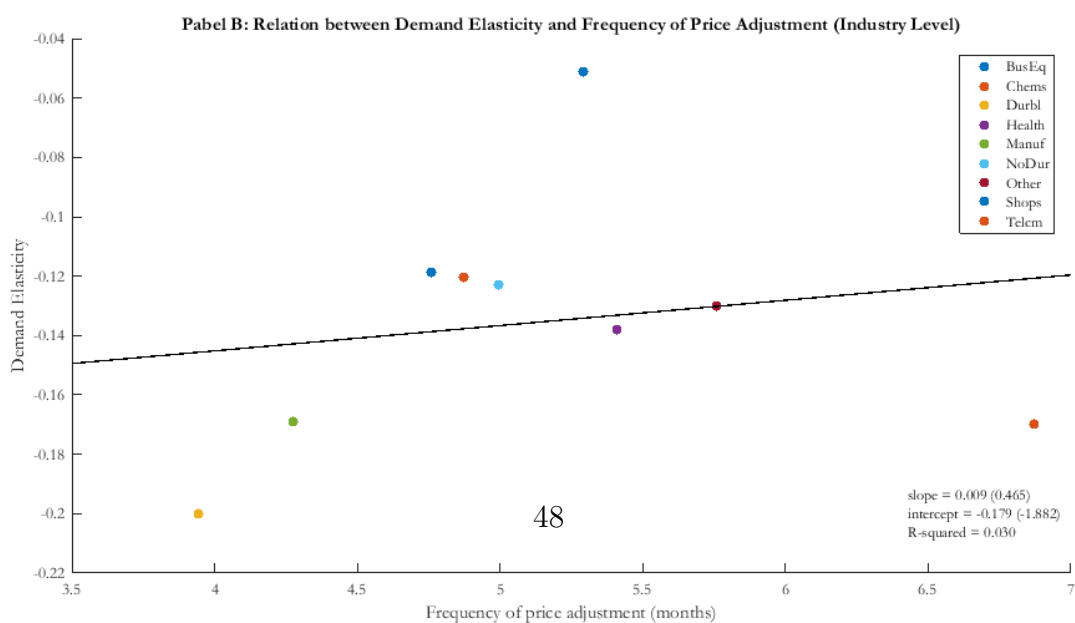
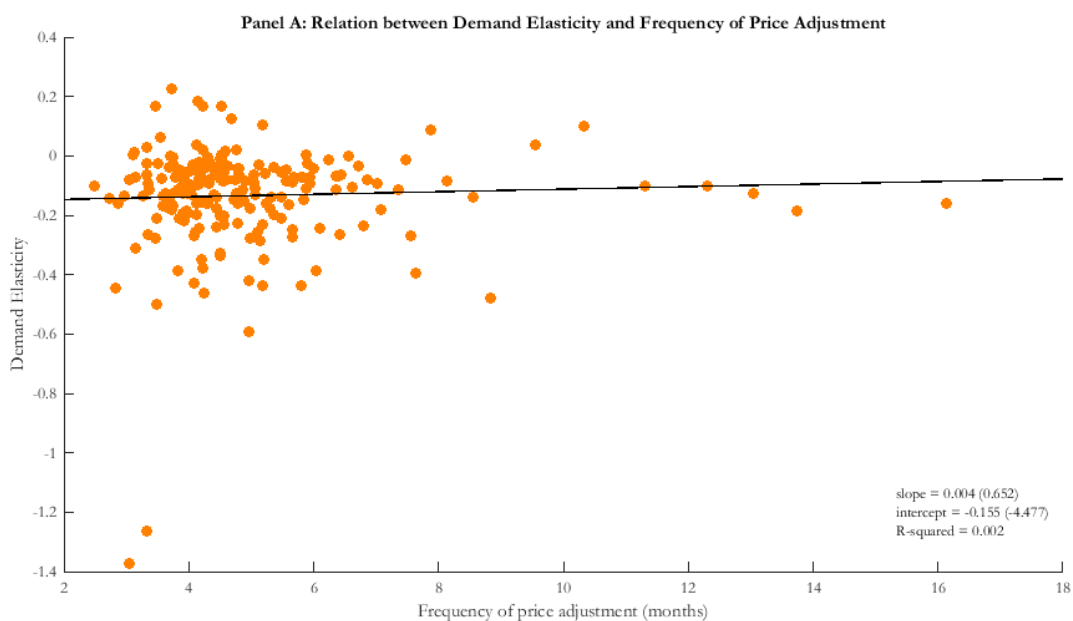


Figure 12: Subsample Splits

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




Table 1: Model implied stock returns

This table reports annualized asset pricing moments for the model described in section B. The second and third columns report average excess returns for stocks in the low elasticity and high elasticity sector 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

Table 2: Heterogeneity in sticky prices versus heterogeneity in elasticities of demand in the model

This table reports returns and markups for the baseline model described in section B (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)).

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

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

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.

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

Table 4: Average Elasticity per Industry

This table reports the average equally-weighted elasticities per industry. I use the Fama-French 12-industry classification. The elasticities are computed using equation (20) and the *Keepa* data. The sample period is February 2015 to April 2017. The second row of the table reports the t-statistics.

	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

Table 5: Sorted portfolios on demand elasticity

This table reports time-series averages of 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.

Average Annual Excess Returns						
	Low Elasticity	2	3	4	High Elasticity	H-L
<u>Panel A: Equally Weighted Returns</u>						
$R_i - R_f$	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]***
<u>Panel B: Value-Weighted Returns</u>						
$R_i - R_f$	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]**

Table 6: CAPM Regressions

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 row 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.

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]***

Table 7: Elasticity Sorted Portfolios - Robustness

This figure reports time-series averages of annual excess returns of elasticity sorted portfolios. Panel A from 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 festivity is excluded (except for Christmas and New Year, where I exclude the full period between those two dates as well). I exclude the days around Christmas, New Year, Valentine's, Thanksgiving, Black Friday, Father's day and Mother's day.

Panel A: Portfolio Sorts Excluding Holiday Days and Festivities

	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
$t\text{-stat}$	[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
$t\text{-stat}$	[2.01]***	[2.57]***	[2.48]***	[2.85]***	[2.71]***	[2.15]***
FPA (months)	4.89	4.72	5.2	4.92	4.86	

Table 8: Panel Regressions of demand elasticity on firm returns and markups

This table reports the results estimating equation (22). The regression is done using annual data. Standard errors are clustered at the firm level and t – stats are reported in parenthesis.

	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

Table 9: Demand Elasticity and Frequency of Price Adjustment

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 parenthesis.

	(1)	(2)	(3)	(4)	(5)	(6)
freqpradj	0.00431*					0.00457
	(1.729)					(1.648)
hhi		0.0259				0.0617**
		(0.901)				(2.047)
beta			-0.0810***			-0.0974***
			(-4.754)			(-4.523)
size				-0.00245		-0.00693**
				(-1.192)		(-2.196)
b2m					0.0108**	0.00785
					(1.995)	(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

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

Table 10: Portfolio Sorts out-of-sample

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.

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]**