Capital Expenditures and the Heterogeneous Response of Equity to Monetary Policy

Undergraduate Economics Honors Thesis

University of California, Berkeley

Michael Wei

May 5, 2017

supervised by

Professor Yuriy Gorodnichenko[1]

[1]**Acknowledgement**: I would like to thank my advisor Yuriy Gorodnichenko for his patience, guidance, and wisdom through which this project was made possible.

Abstract

This paper shows that on average, firms with higher levels of capital intensity experience larger responses in equity price to monetary policy action. This heterogeneous effect is economically large and statistically significant at high levels of confidence. These results suggest that monetary policy has important and substantial distributional implications on the US economy that can be measured by levels of firm capital intensity. This paper also shows that the marginal sensitivity of stock returns to capital intensity is not constant - the effect size varies across subgroups and is observed to be highest in median intensity firms

I. Introduction

Monetary policy is an important tool used by the Federal Reserve to influence real and nominal macroeconomic variables. One channel through which the monetary policy acts is the interest rate, set directly by the Federal Open Market Committee ("FOMC") by updates to the target federal funds rate. These interest rate changes generate reactions in the economy over time, affecting variables such as inflation and output. Additionally, and of specific interest to this paper, they cause an immediate reaction in asset markets as investors and firms react to the new information. These responses are observed to be rather non-uniform and non-random, and suggest that monetary policy action has an important and significant heterogeneous impact across firm stock returns.

This variable effect is especially important to policymakers concerned about the implications of such redistribution. To better understand how the heterogeneity looks, this paper isolates capital intensity as a firm characteristic with substantial explanatory power and quantifies the amount of heterogeneity that can be attributed to it across firms.

This analysis is one element of a greater effort to better understand the distributional effects of monetary policy. Which firms suffer the most? Which reap the highest reward? These questions are not new but remain important ones to be answered. The contribution of this paper to the field of economics is the identification of a firm characteristic that measures this response at significant and economically large magnitudes. The relationship is studied at aggregate and subgroup levels and we find that the association is strongest for median-intensity firms.

II. Literature Review

With the expanding economic role of the Federal Reserve in recent history[2], largely after the two major economic recessions since the early 2000's, the rising importance has been matched by an increased interest in studying the transmission mechanism of monetary policy. The state of research has advanced in both the study of the decision making process of the Federal Reserve - how policy responds to varying economic environments and time frames – and the resulting macroeconomic effects of the policy action. This section focuses on summarizing the advancement of understanding in the latter and discusses relevant studies. This paper builds on several models and methods used in prior analyses and contributes to the evolving study of monetary policy on asset prices.

Monetary Policy on Stock Prices

Bernanke and Kuttner

A seminal study in this field which evaluates the impact of unanticipated changes in the federal funds rate on equity prices was conducted by Bernanke and Kuttner (2005). They approach the question using an event-study style analysis, measuring the one-day-gain of CRSP's value weighted index on the day of an FOMC announcement. They compare these returns to the magnitude for the unexpected rate change, which they measure using a difference in the federal funds futures rate on the day of and the day prior. A very similar method for calculating the unanticipated portion of the rate change is presented in the following section.

Using a sample of announcements starting from June 1989 to December 2003, Bernanke and Kuttner find a negative relationship between the direction of the policy rate change and the direction of the stock market price response. When regressing on the entire rate change (the sum of expected and unexpected portions), the association is still negative, but statistically insignificant. Instead, narrowing their analysis to the unanticipated portion of the rate change, they estimate that a 25 basis point cut in the target federal funds rate is associated with a 1% increase in broad stock indexes, on average. This estimate was highly statistically significant. When large outliers are dropped, the estimate remains negative and significant but falls in magnitude by roughly half.

[2]Source: http://www.cfr.org/international-finance/role-us-federal-reserve/p21020

To asses the integrity of their results, Bernanke and Kuttner consider various internal threats to validity within their model. A concern arises regarding the orthogonality of the error term with the FOMC announcement size. One possible reason they propose that the condition may be violated is if the direction of causation were reversed: that a rise (fall) in the stock market was immediately responded to by a cut (hike) in the policy rate. They dismiss this possibility, however, on the basis that there were no clear examples of such a reaction taking placing. Another possible threat to orthogonality they consider is if both the stock market and monetary policy react to other new information. A plausible source of outside shock that Bernanke and Kuttner posit are announcements of slower than expected growth of the economy. In this example, an expansionary monetary policy action might be taken at the same time the markets fell from the negative news. They consider the concern trivial as any bias that did exist would only attenuate the estimate, not exaggerate it.

Another set of possible concerns Bernanke and Kuttner raise is in regards to asymmetries in their results. They consider the possibility that the magnitude of the monetary policy response depends on the direction of the action - that an equivalent sized monetary policy hike and cut might generate different magnitudes of equity response. To test, they include a dummy interacted with the surprise component of the policy change. The results indicated a slightly smaller absolute effects for policy hikes, but the estimate on the interaction term is statistically insignificant. Taking these above results together in addition to several omitted checks, Bernanke and Kuttner conclude that the results are consistent with a strong 1-day reaction of the stock market to unanticipated changes in the Federal funds rate. This paper will continually refer to the work of Bernanke and Kuttner (2005) for references of methodology and theory.

Heterogeneous Response to Monetary Policy

In addition to just quantifying the average effect of monetary policy on asset prices, there have also been numerous studies investigating the heterogeneity of response to monetary policy. We focus first on the literature that specifically studies stock return sensitivities and then proceed to summarize a series of papers that cover other macroeconomic responses.

Ehrmann and Fratzscher

Ehrmann and Fratzscher's study from 2004 had a portion of analysis parallel to that being conducted in this paper. The qualitative comparison is relatively similar, although the model being applied differs. They set out to test whether monetary policy had a different effect on the equity returns of firms based on financial constraints and/or the quality of investment opportunities. They use the following model:

$$r_{i,t} = \alpha + \beta_1 s_t + \beta_2 s_t x_{i,t} + \epsilon_{i,t} \tag{1}$$

Ehrmann and Fratzscher find strong empirical support for large differences in the effects of monetary policy across firms that have varying levels of financial constraints (represented by different proxies). In addition, they conclude that firms with low cash flows, small size, poor credit ratings, low debt to capital ratios, high price-earnings ratios or a high Tobin's q are affected significantly more by US monetary policy.

Gorodnichenko and Weber

Gorodnichenko and Weber (2016) study costs of price stickiness on firm response to monetary policy. They use a similar event-study approach as Bernanke and Kuttner (2005) with an event window of only 30 or 60 minutes in width for both stock returns and the Federal funds futures. The goal of their study was to show results consistent with the New Keynesian model: that sticky prices were costly to firms. They use the following specification to assess whether price stickiness led to heterogeneous response of equity returns:

$$R^2_{it} = b_0 + b_1 * v^2_t + b_2 * v^2_t * \lambda_i + b_3 * \lambda_i + FirmControls + FirmControls * v^2_t + error_{it} \tag{2}$$

The relationship studied was of the coefficient on the interaction term between the policy shock and price stickiness. Gorodnichenko and Weber show that the conditional volatility of stock market returns rises more for firms with stickier prices than firms with more flexible prices. These results were robust and help show why monetary policy heterogeneity persists within the stock market.

Peersman, Smets - Carlino, DeFina

Finally, we briefly summarize two papers that explore other channels through which monetary policy is suspected to create heterogeneity within the economy.

Peerman and Smets (2002) study the effects of a euro area monetary policy change on output growth. They find that on average, the contractionary effect of an interest rate hike is significantly greater in recessions than in booms. In addition, they find that the degree of asymmetry of policy effects are related to differences in financial structure i.e. maturity structure of debt, the coverage rate, financial leverage and firm size.

Carlino and Defina (1996) look into regional effects; they suggest that the diverse regions that are linked across the United States respond differently to changing economic circumstances. They provide as an example the fall in crude oil prices in the mid 1980's and how energy-producing regions were affected very differently than energy-consuming regions. From their analysis, they conclude that two regions in particular - the Southwest and Rocky Mountain - experienced smaller relative monetary policy effects within the local economy than the national economy. Carlino and Defina suggest that these non-uniform effects underscore the difficulty of conducting monetary policy across a region as diverse as the US.

III. Data Source

This section describes the data collection and data set construction process.

Monetary Policy Data

In order to measure firm response to monetary policy, we look at FOMC announcements from 1994 to 2009 and firm stock returns on each corresponding announcement date. Using an event-study method employed by Bernanke and Kuttner (2005), we compare asset prices before and after each announcement, and attribute any change in price to the release of the policy rate information. A difficulty in measuring this relationship, as noted in past literature, is the forward-looking nature of asset markets. Anticipated interest rate changes are assumed to be priced ex-ante into the stock market and should theoretically generate no real response at the time of announcement. To account for this, we compare firm asset response only to the unanticipated portion of the policy change. This unexpected portion should in effect "surprise" the market and generate a non-trivial response in asset value.

As described by Gorodnichenko and Weber (2016), who use a similar approach to Bernanke and Kuttner, one way to isolate this unanticipated change is by looking at the price of Federal funds futures contracts in a narrow time window surrounding the FOMC announcement. The price of these contracts, as shown by Kruger and Kuttner (1996), serve as a good approximation for market forecasts of the effective Federal funds rate; a change in price of this derivative is representative of a shift in expectation for the future Federal funds rate. Hence, taking the difference in contract price shortly before and shortly after an FOMC announcement provides an estimate for the unanticipated change in policy rate. The relationship can be described as

$$\Delta i^{unexpected} = \frac{D}{D-t}(ff^0_{t+\Delta t} - ff^0_{t-\Delta t}), \tag{3}$$

where $\Delta \mathrm{i}^{unexpected}$ is the unexpected rate change, D is the number of days in the month, t is the time when the FOMC issues an announcement, $ff^0_{t+\Delta t}$ is the Federal funds future rate just after t, and $ff^0_{t-\Delta t}$ is the Federal funds future rate just before t. The $\frac{D}{D-t}$ scalar takes into account the fact that the contract's settlement price is based on the average Federal funds rate and adjusts the price change accordingly.

We gather results of the above methodology from Gorodnichenko and Weber (2016, online appendix). They calculate the unanticipated policy rate change for every FOMC announcement that occurred between February 4, 1994 and December 16, 2009, less one. They choose to exclude the policy rate change of September 17, 2001 – the first day trading resumed after the September 11 terrorist attack. For each announcement, they find $\Delta \mathrm{i}^{unexpected}$ within a "tight window" (30 minutes in width, starting at t – 10 minutes) and "wide window" (60 minutes in width, starting at t − 15 minutes). We opt to run our analysis using the wide window shocks since the asset prices we compare are daily, not minutes apart. Table 1 and Figure 1 below summarize the 137 unanticipated shocks.

Table 1: Monetary Policy Shock (Percentage Points)

Mean	Median	SD	Min	Max	Observations
-0.015	0	0.091	-0.46	0.15	137

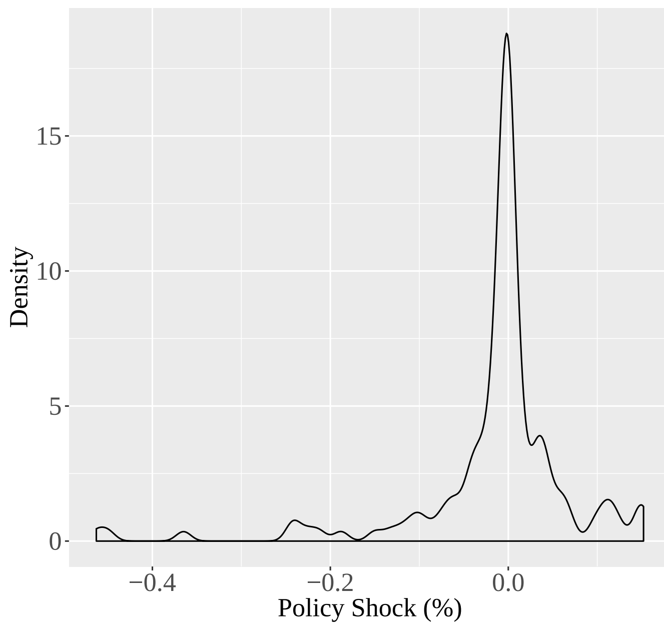

Figure 1: Density Plot of Monetary Policy Shocks

The average unanticipated monetary policy shock across the 1994 − 2009 period was approximately 0. Figure 1 shows that the distribution skews to the left and includes a few large, negative observations.

Firm Data

Database Description

To study firm response behavior, we limit our analysis to an unbalanced panel of large-cap companies across the same 1994 – 2009 period, pairing firm data to each FOMC announcement date. The sample of companies that we choose to study are the constituents of the S&P 500 as of January 1, 2010. These firms have high market capitalization, are incorporated in the US, and generally have robust information archived. Hence, the results are non-trivial in magnitude and are generalizable to other firms.

For each company, we obtain financial and fundamentals data from CRSP and Compustat databases, respectively. CRSP maintains a comprehensive collection of stock price, return, and volume data; Compustat is a database of fundamental, financial and price data for publicly traded companies.[3] Each company is tied to a searchable unique identifier which eliminates the possibility of misidentification. A Global Company Key (gvkey) is used in Compustat and a Permanent Company Identifier (PERMCO) is used in CRSP.[4] we obtain a primary list of 502 gvkey's from Compustat's Index Constituents database, which are then mapped (with a one to one relationship) into PERMCO's using the CRSP / Compustat linking table.

Variables of Interest

From CRSP, we find each company's 1-day stock return on each of the 137 announcement dates. Asset *return* is defined as the difference in the company's stock closing price (CRSP variable PRC) less the opening price (CRSP variable OPENPRC) as a proportion of the opening price. In the instances where a given company has more than one class of common stock at a time, we include only the security deemed the "primary issue" by CRSP and drop all "joiner secondary" issued securities.[5] If a branch of a company is spun off and retains the parent company's PERMCO

[3] Source: http://www.whartonwrds.com/datasets/crsp/, https://www.library.hbs.edu/Find/Databases/Compustat for more detailed descriptions.

[4] Source: https://www.alacra.com/alacra/outside/lei/info/gvkey.html, http://www.crsp.com/products/documentation/crsp-link

[5] Source: http://www.crsp.com/products/documentation/link-history-data

identifier (such that one PERMCO is used for two different entities), we drop observations of the newer entity. If a company only modifies its existing security, the returns on the prior and new security are treated equivalently.

From Compustat, we obtain quarterly firm-level data on capital expenditure levels (Compustat variable CAPXY), cash and liquid assets (Compustat variable CHEQ), and net sales (Compustat variable SALESQ). We define *capexratio* as capital expenditures over net sales and *cash* as the sum of cash and liquid assets over net sales.

Merging

To construct the desired dataset, we join the financial and fundamentals data by company identifier and date. In order to merge quarterly and daily data, however, we first transpose each of the 137 relevant dates into a year and calendar quarter. For example, we assign the policy rate change and asset returns on February 4, 1994 the value 1994Q1, the same date formatting produced by Compustat. Each observation in the resulting joint table captures a company's financial response, capital characteristic, and policy rate change for a given announcement date. The data collected, however, are subject to irregularities; issues such as missing values, suspicious outliers, and insufficient sample size are addressed in the next section.

Cleaning

The raw dataset of returns and fundamentals obtained contains 62,864 observations. Observations with missing values for *return*, *capexratio*, or *cash* are dropped. Additionally, observations with *capexratio* or *cash* values of 0 are deemed unrealistic for a large-cap company and are dropped.

There are three companies that have fewer than 10 observations each. The possible reason for so few data points is difficult to identify and raises concerns about their representativeness. We remove the three companies from our analysis. Additionally, we remove all companies that have a mean *capexratio* greater than 1. The statistic implies that a company's average quarterly capital expenditure exceeds their average quarterly income. This behavior is unsustainable and does not accurately represent the firm's true capital intensity.

The above criteria remove 8,976 observations from the dataset. The remaining single day returns consist of large outliers on both ends of the distribution. Across all observations, asset values rise as much 102.70% and fall as much as 38.49%. Given the narrow one-day time frame, these returns are anomalous and are likely the result of irregular events. Hence, the entirety of these large absolute returns should not be attributed to changes in the policy rate. To reconcile this fact without the complete loss of information, we cap absolute returns at 10%. Table 2 and Figure 2 below summarize the 1-day stock returns that occur on FOMC announcement dates after capping absolute returns.

Table 2: Asset Return (%)

	Mean	Median	SD	Minimum	Maximum	Observations
Raw Returns (%)	0.509	0.282	2.807	-38.49	102.7	53888
Capped Returns (%)	0.473	0.282	2.504	-10.00	10.0	53888

Note: This table summarizes 1-day stock returns on FOMC announcement days.

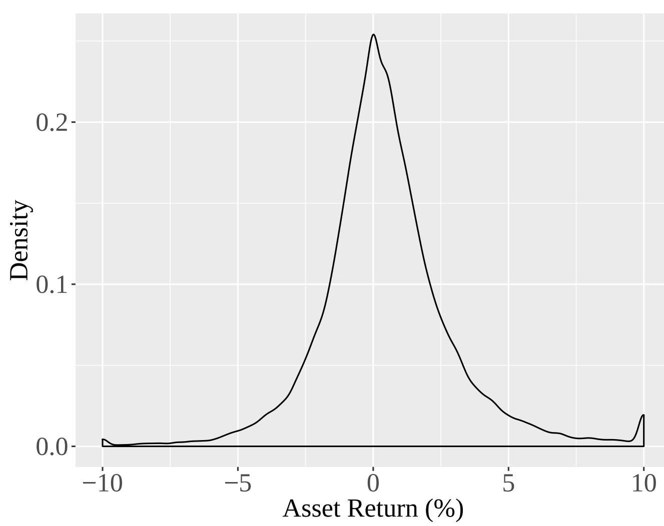

Figure 2: Density Plot of 1-Day Firm Stock Returns

The average daily return is positive and approximately 0.473% after applying the cap. The variability measured in standard deviations falls by 30.3 basis points from the cap while the sample median stays the same. This behavior is consistent with the presence of large, positive outliers and illustrates a slight right skew before capping returns. Since the the proportion of outliers in the dataset is low, the two distributions vary only marginally.

Table 3 below summarizes the *capex*, *sales*, *capexratio*, and *cash* variables obtained from Compustat for the same remaining sample of firms.

Table 3: Summary Statistics of Firm Fundamentals

	Mean	Median	SD	Minimum	Maximum	Observations
capex	193.273	56.404	446.180	-381.00	8564.00	53888
sales	3145.435	1274.600	6536.049	-25623.00	124238.00	53888
capexratio	0.085	0.045	0.159	-1.86	9.78	53888
cash	0.973	0.280	3.761	0.28	390.00	53888

Note: Capex, sales, and cash are measured in millions of USD

The mean capital intensity across the sample of high-cap companies is approximately 0.085. This indicates that on average, firms invest 8.5% of their total quarterly sales into capital outlays. Capital intensity, however, varies substantially across firms and across time periods. The standard deviation is 15.9%; the maximum ratio observed, an extreme outlier, is approximately 60 standard deviations above the mean.

IV. Model Specification

Theory

In order to ascertain the heterogeneous effect of monetary policy on equity prices, the following entity fixed effects regression is proposed:

$$\begin{aligned} return_{i,t} = constant + \beta_1 mp_t + \beta_2(mp_t * capexratio_{i,(t-1)}) + \beta_3 capexratio_{i,(t-1)} \\ + \beta_4(mp_t * cash_{i,(t-1)}) + \beta_5 cash_{i,(t-1)} + \alpha_i + error_{i,t} \end{aligned} \tag{4}$$

where $return_{i,t}$ is the one-day stock return of firm i on announcement date t, mp_t is the unanticipated change in the federal funds target rate on announcement date t, (mp_t * $capexratio_{i,(t-1)}$) is the interaction between the policy rate change and firm i's capital expenditure ratio in the quarter prior to announcement t, $capexratio_{i,(t-1)}$ is firm i's capital expenditure ratio isolated, (mp_t * $cash_{i,(t-1)}$) is the interaction between the policy rate change and firm i's short term access to funding, $cash_{i,(t-1)}$ is firm i's short term access to funding isolated, and α_i is firm i's fixed effect. Descriptive properties of these variables were described in the sections above.

β_1, the relationship between surprise policy change and equity price, has been studied previously and shown to be negative by Bernanke and Kuttner (2005). As the following analysis uses a similar event-study approach, we expect our results to be consistent. This paper instead centers its attention on the nature of β_2 - the coefficient of the *capexratio* interaction. β_2 reflects the "additional" effect that levels of capital intensity have on a firm's asset return, conditional on the presence of a monetary policy shock. A non-zero β_2 would suggest that company stock prices have varying levels of sensitivities to the federal funds rate, which can be at least partially explained by the firm's level of capital intensity. Quantifying this heterogeneous response across firms underlies the primary purpose of this paper, which contributes to a growing body of literature dedicated to studying the distributional effects of monetary policy.

An important property to acknowledge moving forward in this analysis is the sign of the individual capital expenditure observations. Variable *capexratio* is positive across almost all 53,888 observations with an exception to a few (0.6%) anomalous values. This fact has a convenient implication: the resulting directional sign of the interaction term will be the same as the sign on variable *mp* if β_2 is negative, and will be opposed if β_2 is positive. The relative direction of these two terms determines the direction of stock market heterogeneity. For example, a negative β_2 suggests that firms with a higher level of capital intensity are more responsive to changes in the federal funds rate: a one percentage point increase in the unanticipated target rate would change our estimate of asset return by (-β_2 - β_3 * $capexratio_{it}$). Alternatively, a positive β_2 changes the response to (-β_2 + β_3 * $capexratio_{it}$). The important relationship here is that $|(-\beta_2 - \beta_3 * capexratio_{it})| < \beta_2 < |(-\beta_2 - \beta_3 * capexratio_{it})|$.

Theory suggests that β_2 is negative – that firms with a higher capital intensity are more sensitive to changes in the federal funds rate. The justification for this depends largely on the method of transmission between monetary policy and equity price. Although the exact channel through which stock prices are affected is still being studied, Bernanke and Kuttner (2005) suggest that at least a partial effect is likely due to changes in the cost of capital. As the price of borrowing rises, firms that depend on external funding suffer relatively more, the present value of future earnings decrease, and the price of their equity falls as a result. A natural extension to this reasoning is that firms with a greater external dependency would feel the rise in interest rates more starkly. Hence, we anticipate that the equity value of firms with higher capital intensity respond with greater sensitivity to fluctuations in the Federal funds rate.[6]

Technical Notes

The construction of the model proposed above was careful and deliberate. This section describes particular characteristics and the economic reasoning that underlies them.

One preliminary concern about regressing on a firm's level of capital expenditure is unintentionally capturing the firm's size. Macroeconomic theory (and empirical results) suggests a correlation between firm size and response to monetary policy which would significantly bias the β_2 coefficient of interest (Ehrmann 2000). To account for this omitted variable, the fundamentals data for capital expenditure and liquid assets are normalized by the firm's level of quarterly net sales, a widely accepted method of normalization. The result of this approach is a desirable measure of a company's capital intensity, or the proportion of quarterly sales it reinvests back into capital.

To differentiate the conditional and unconditional relationships between *capexratio* and stock return, the model includes the $capexratio_{i,(t-1)}$ term by itself. Intuitively, any unconditional effect that a firm's capital intensity has on its stock returns will be differenced out of β_2 and captured instead by β_3. Economic theory, however, strongly suggests that the unconditional effect is trivial.

There is an additional concern around the model's implicit assumption that capital expenditures are paid for through borrowing – the connection that links the interest rate to the firm's future

[6]Note: We reference the federal funds rate synonymously to the cost of borrowing although there is a technical discrepancy. Companies borrow at the prime rate, which has been shown historically to closely track the federal funds rate. Source: http://www.businessinsider.com/federal-reserve-interest-rate-impact-consumers-businesses-2016-12/#prime-loan-rates-are-established-by-private-banks-as-a-baseline-rate-for-loans-to-businesses-and-consumers-2

earnings. If, instead, the company has enough cash or other liquid assets, the need to borrow is diminished and the stock price sensitivity would fall. To prevent this extra noise, the model includes two terms that control on the total cash and short term assets in a given quarter, as a proportion of sales. The first, ($mp_t * cash_{i,(t-1)}$) is included as a regressor to measure the monetary policy heterogeneity attributable to a firm's relative level of cash. In addition, $cash_{i,(t-1)}$ is added on its own to capture the unconditional effect cash may have on stock returns.

As indicated by the subscripts "*t-1*", the fundamentals data used to characterize each firm is lagged. For example, an FOMC announcement that occurs on February 4, 1994 will draw upon 1993Q4 data from Compustat. The goal is to use information that most accurately depicts firm characteristics at the time of each announcement. We assume the prior level of capital intensity and level of cash is recent and most representative of the firm. The concern with using data indexed t is the likelihood that observed levels capture the firm's response to the rate change, not the levels leading up to it.

Because the dataset used in this analysis is panel in nature, a firm fixed effect estimator α_i is included in the regression. α_i is intended to capture systematic time-invariant differences between companies that may, for example, affect the transmission between rate change to equity price.

V. Analysis

Consistent with the methods used by Gorodnichenko and Weber (2016), we conduct an initial pass through the data to assess the validity of monetary policy as a meaningful shock to equity price (shown on the following page). Column (1) shows that, in line with prior schools of thought, unanticipated increases in the policy rate generate equity value falls that are significant at the 1% level. A 25 basis-point increase lowers stock price by roughly 1.5%.

Regressions (2) and (3) are a first effort to isolate the value of β_2 - the amount of sensitivity associated with capital intensity. Regression (2) performs a standard OLS regression on the aggregate data while regression (3) includes entity fixed effects. Against our prior intuition,

Table 4:

	Dependent variable:		
	return		
	OLS		*panel linear*
	(1)	(2)	(3)
mp	-6.140^{***}	-6.218^{***}	-6.179^{***}
	(0.116)	(0.136)	(0.136)
mp * capexratio		2.494^{***}	2.428^{***}
		(0.787)	(0.788)
capexratio	-0.287^{***}	-0.236^{***}	-0.239^{**}
	(0.067)	(0.069)	(0.095)
mp * cash		-0.158^{***}	-0.172^{***}
		(0.033)	(0.033)
cash	0.013^{***}	0.009^{***}	-0.001
	(0.003)	(0.003)	(0.004)
Constant	0.390^{***}	0.390^{***}	
	(0.012)	(0.012)	
Fixed Effects	No	No	Yes
Observations	53,888	53,888	53,888
R^2	0.050	0.051	0.051
Note:	This table reports the results of regressing 1-day returns of the S&P 500 constituents on the unanticipated portion of the Federal funds rate policy action, the interaction (mp * capexratio), capexratio isolated, the (mp * cash) interaction, and cash isolated. The full sample ranges from February 1994 to December 2009, excluding the release of September 17, 2001, and captures data from 475 unique companies.		

both regressions suggest that the direction of β_2 is positive at highly significant levels. A net model estimate of +2.25 (the combined conditional and unconditional effect) can be roughly interpreted as a 35% reduction in equity sensitivity to monetary policy when a firm increases its capital expenditures by its net quarterly sales amount (assuming initial capital expenditures of 0).

In addition, Table 4 shows that the *capexratio* has a small but significant negative unconditional association with equity returns. On average, firms with 10% higher capital intensity experience returns that are lower by roughly 2 basis points. Cash and liquid assets are also shown to contribute a small, but statistically significant portion of the heterogeneous response; typically, firms with higher levels of cash intensity (cash as a proportion of sales) issue stock that are slightly more sensitive to monetary policy.

Outliers

The counterintuitive estimates for β_2, however, strongly suggest the presence of outlier companies: clusters of firms that behave systematically different than the rest. To identify these potential outliers, we classify individual firms into buckets and study their behavior at the local level. A natural classifier to use is the Global Industry Classification Standard (GICS) which was developed by MSCI and S&P as an industry classification system.[7] Each company is labeled at the sector, group, industry, and sub-industry level, which are recorded and available on Compustat. This analysis chooses to limit its scope to the sector-level, composed of the following eleven categories: energy, materials, industrials, consumer discretionary, consumer staples, health care, financials, information technology, telecommunication services, utilities, and real estate (most recently added on August 31, 2016).[8] With each company classified, the following OLS regression is run separately on individual companies,

$$r_{i,t} = \delta_0 + \delta_1 mp_t + \delta_2 capexratio_{i,(t-1)} + \delta_3 cash_{i,(t-1)} + error_{i,t} \tag{5}$$

for a total of 475 regressions. Each estimate of δ_1 represents an equity price sensitivity for that particular company. Figure 3 on the follow page plots the δ_1 estimate for each company on the company's mean *capexratio*, color coded by GIC sector. Mean *capexratio*'s are calculated using data from the 1994 – 2009 FOMC announcement quarters. Two extreme outliers are omitted from the plot to allow for reasonable scaling.

[7]Source: https://www.msci.com/gics

[8]Source: https://www.msci.com/documents/1296102/1339060/GICSSectorDefinitions.pdf/fd3a7bc2-c733-4308-8b27-9880dd0a766f for a description of each sector

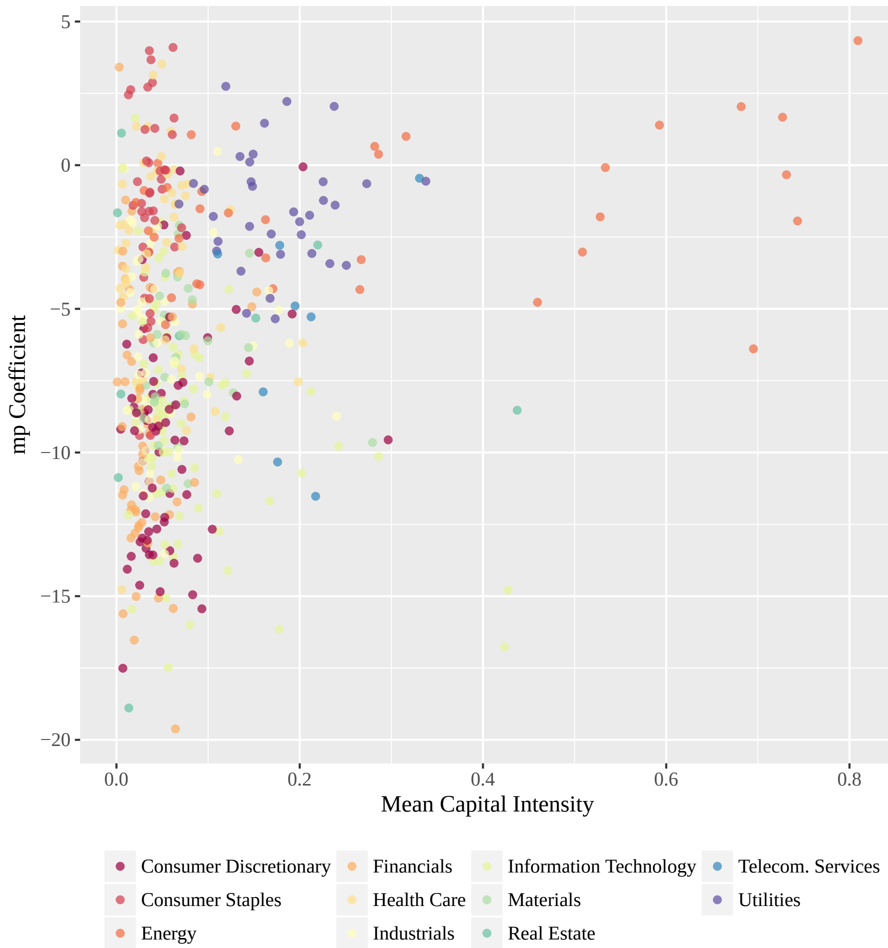

Figure 3: Equity Price Responsiveness and Mean Capital Intensity by Company

Consistent with economic theory, the sensitivities of most firms above are negative. We observe that estimates typically fall within the range $0 < |\delta_1| < 13$, with a few more extreme observations. Most notable in Figure 3, however, are the two prominent clusters of points located on the periphery of the cloud – points with high capital ratios but small absolute sensitivities. These companies fall very consistently into two GIC sectors: energy and utilities. Figure 3a and 3b on the following page highlight the two sectors.

Figure 3a

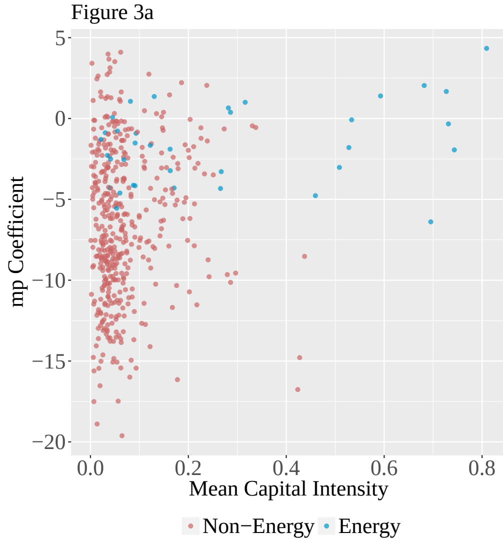

Figure 3b

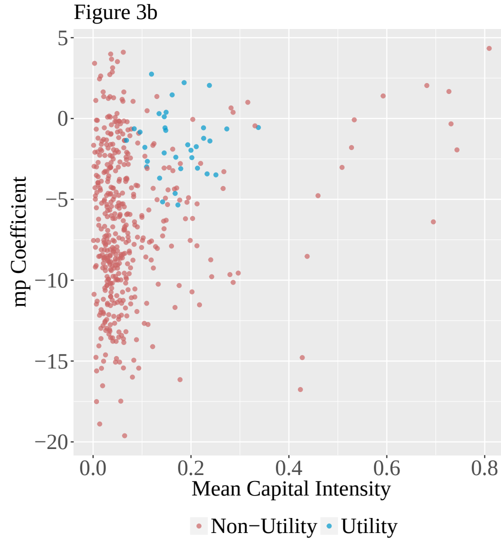

The plots strongly suggest that energy and utility companies behave systematically different than the remaining GIC sectors in terms of their equity response to monetary policy. This observation, however, is not novel. Cafariello (2015), in a piece that addressed the then recent anticipations of interest rate hikes by the Federal Reserve, argued that energy and utility companies wouldn't suffer to the extent to which investors might expect.[9] His analysis came at a time of growing concerns around the impact that higher interest rates would have on the extremely capital intensive energy and utility industries. In a similar strain to our own analysis, investors and speculators anticipated such firms to be most adversely affected. Cafariello's stance differed and offered three counter-reasons for why he believed the theory would not uphold behavior. The primary reason, or the one Cafariello most expanded on, was that energy and utility companies reacted more to the climate of the market than to the changes in the cost of borrowing. As evidence, he refers to the tracking of the utility sector as measured by the SPDR Utility Sector ETF (NYSE: XLU) to the federal funds rate across 2000 to 2008. Figures 4a and 4b illustrates the federal funds rate and SPDR ETF, respectively; the two charts are taken from the original article.

[9]Note: Although Cafariello does not explicitly refer to "energy" companies in his analysis, his use of Duke Energy Corporation and Eversource Energy (amongst others) as examples suggests the implicit inclusion.

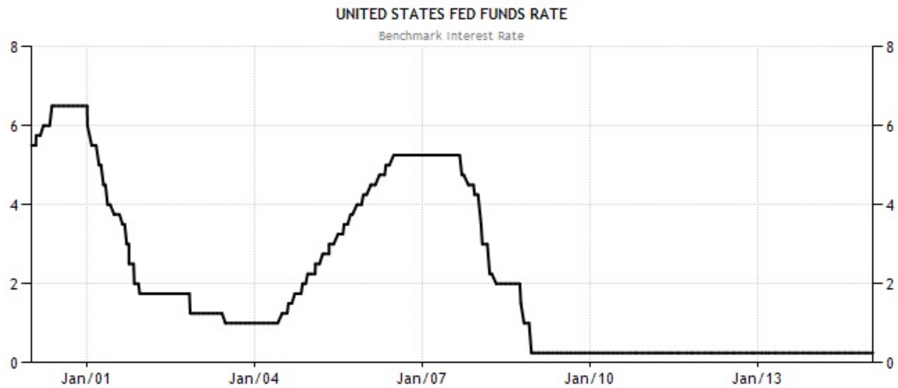

Figure 4a: Federal Funds Rate 2001 - 2015

Figure 4b: SPDR Utility Sector ETF 2001 - 2015

The parallel movements of the two markets suggest that increases in the Federal funds rate generate positive reactions across the utilities industry, and vice-versa. Cafariello attributes the counterintuitive reaction to a signaling effect, where policy rates increases indicate a strong economy, higher wages, and hence the ability to pass on higher costs of borrowing to consumers. These conclusions, in accordance with our own observations, suggest that energy and utility responses are anomalous. We consider the evidence sufficient to omit these industries from the analysis and proceed to for the remainder of this paper.

Results

We consider the 9 remaining GIC sectors and regress again on the aggregate. The regression before dropping values is included in column (1) for baseline comparison.

Table 5:

	Dependent variable:			
	return			
	(1)	(2)	(3)	(4)
mp	-6.179^{***}	-6.156^{***}	-6.346^{***}	-6.189^{***}
	(0.136)	(0.151)	(0.139)	(0.155)
mp * capexratio	2.428^{***}	-3.161^{**}	1.451^{*}	-7.929^{***}
	(0.788)	(1.266)	(0.820)	(1.404)
capexratio	-0.239^{**}	-0.926^{***}	-0.293^{***}	-1.122^{***}
	(0.095)	(0.138)	(0.098)	(0.146)
mp * cash	-0.172^{***}	-0.134^{***}	-0.157^{***}	-0.097^{***}
	(0.033)	(0.033)	(0.033)	(0.034)
cash	−0.001	0.003	−0.0004	0.005
	(0.004)	(0.004)	(0.004)	(0.004)
Fixed Effects	Yes	Yes	Yes	Yes
Energy Companies	X		X	
Utility Companies	X	X		
Observations	53,888	49,642	51,270	47,024
R^2	0.051	0.058	0.053	0.062

Note: This table reports the results of regressing 1-day returns of the S&P 500 constituents on the unanticipated portion of the Federal funds rate policy action, the interaction (mp * capexratio), capexratio isolated, the (mp * cash) interaction, and cash isolated. The full sample ranges from February 1994 to December 2009, excluding the release of September 17, 2001, and including a different number of companies in each regression. Depending on what sectors are included, the number of companies ranges from 405 to 475.

Fixed effects regression (2), which removes only energy companies, estimates β_2 (defined previously as the coefficient of the *capexratio* interaction) to be -3.16. The standard error is 1.27 and the estimate is significant at the 5% level. There were 4,246 energy company observations removed, or roughly 8% of the data. Notably, the omission of the energy industry had large, significant effects on the β_2 estimate. The direction of the coefficient swings from positive to negative and becomes even larger in magnitude.

Regression (3) conducts the same fixed effect regression, but instead removing the utility industry.

The estimate for β_2 is 1.45 with a standard error of 0.82. β_2 is significant at a 10% level. Comparing regressions (2) and (3), we see that the effect on β_2 from dropping utility companies is less pronounced than dropping energy companies. This is consistent with Figures 3a and 3b above, which show that utility companies are more similar to the remaining sectors. There are also roughly twice as many energy companies as there are utility companies in the original data set.

Regression (4) removes both energy and utility companies, dropping a total of 6,864 observations or roughly 12.7% of the data. The regression on the remaining 9 GIC sectors estimates a β_2 to be -7.93, with a standard error 1.40; the coefficient is significant at the 1% level. The result suggests the removal of the two industries had a significant effect on the relationship between the interest rate and equity price. The magnitude of the coefficient β_2 more than doubles relative to any of the other 3 regressions. The resulting value of -7.93 indicates that for a given 25 basis point rise in the policy rate, firms with capital expenditures of k + *sales* per quarter t will typically suffer a roughly 2% additional fall in asset price compared to firms with expenditures of k. The equity price of these firms are roughly twice as sensitive to unanticipated changes in the target rate.

The outcomes from above support our original hypothesis: that firms with higher capital intensity are more responsive to unanticipated rate changes. Through this analysis, we have shown not only that there is a heterogeneous response in the equity market to monetary policy, but have shown capital intensity to be a useful measure. The β_2 value is found to be both statistically significant and substantially large in macroeconomic terms. In addition to this primary conclusion, the model also measures the portion of the heterogeneity that is associated with the firm's cash intensity. The estimated effect is negligible in magnitude although statistically significant at the 1% level. A firm with cash and liquid assets of c + *sales* in quarter t will, on average, be one one-hundredth more sensitive to unanticipated rate changes.

Heterogeneity Across Subgroups

The assumption implicit in this model thus far is that firms are subject to a similar relationship between stock returns, capital intensity, and monetary policy - that β_2 is essentially constant across firms. This needs not be true, and likely does not hold across clusters that have inherently different levels of capital intensity. We suspect instead that the transmission mechanism is non-uniform and depends on several factors that are not accounted for within the model. These factors will change the relationship that capital intensity has on equity responsiveness and thus create variability in β_2 across clusters. An example of a plausible factor is the access to (or frictions in the access to) funding - firms that are more capital intensive will likely experience greater ease in borrowing.

This next series of tests look for the potential variation in β_2 across different subsets of firms. Companies are split into three different clusters based on their mean capital intensity level across the 1994 – 2009 period, measured as the mean $capexratio_{t-1}$. Group 1 is defined as the third of firms with the lowest mean $capexratio_{t-1}$'s (158 firms). Group 2 is defined as the firms with the middle third of $capexratio_{t-1}$'s (159 firms). Group 3 is defined as the third of firms with the highest mean $capexratio_{t-1}$'s (158 firms).

To understand the composition within each group, Figure 5 presents the distribution of companies across GIC sectors, clustered by Group. Note that the frequency of each sector is weighted by the number of observations, not by the number of companies; this is to account for the fact that companies have a variable number of observations. The distribution shows that Group 1 is disproportionately composed of the financial sector i.e. banks, while Group 2 is disproportionately composed of the consumer staples sector i.e. grocery and drug stores.

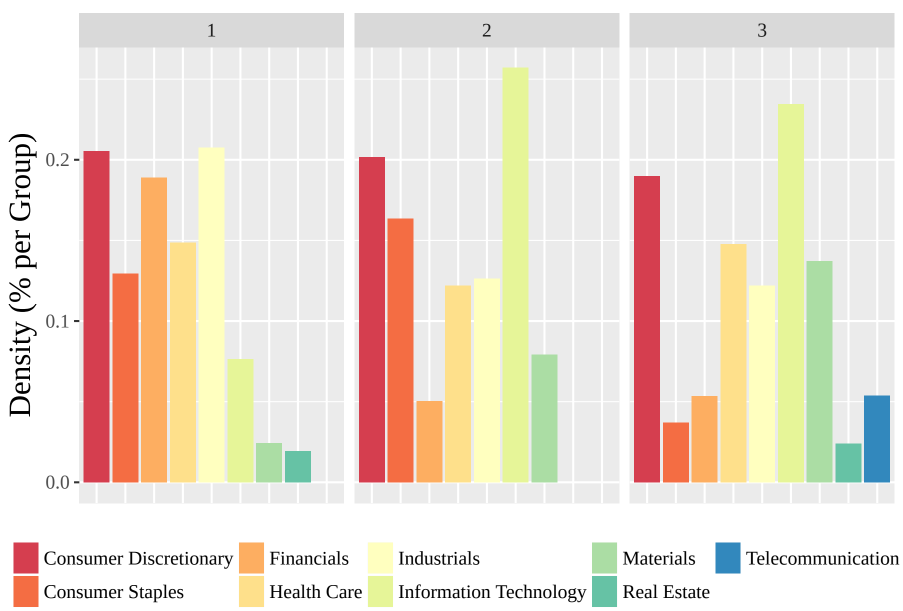

Figure 5: Bar Chart Illustrating the GIC Sector Composition by Group

In order to compare the behavior of firms across the three groups, we apply the model onto each group independently. There is a concern, however, about a direct application. Namely, the strong right skew of the *capexratio* distribution causes most values to fall within a narrow interval of roughly [0, 0.20], with fewer points composing the right tail. See Figure 6 for the distribution,

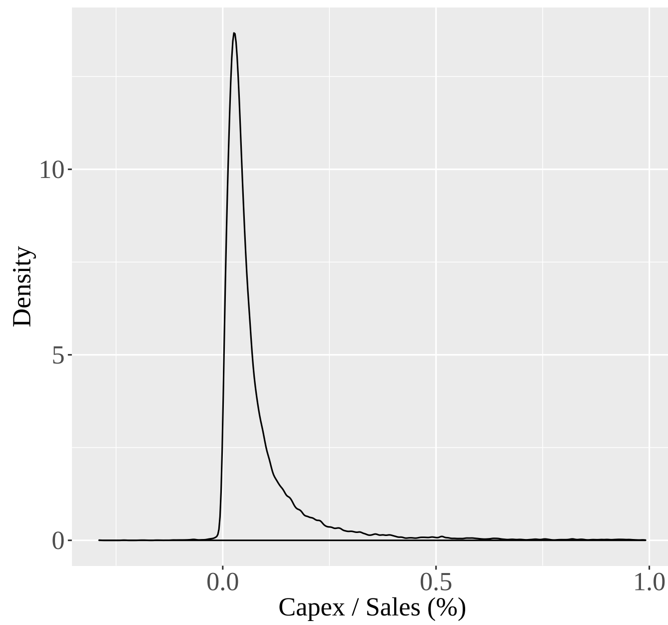

Figure 6: Density Plot of Capital Intensity Levels Across All Groups

This fact creates an undesired result: the spread of *capexratio* values in the first tertile is more narrow than the spread in the second tertile, which is additionally more narrow than the third. In other words, the variance of the *capexratio* regressor is not constant across the three Groups. The concern of uneven spread arises in the interpretation of the estimates: the coefficient β_2 will accurately describe the marginal association of *return* and *capexratio* but importantly misrepresents the true variation in sensitivity within each cluster. The magnitude of *return* that can be associated to the capital intensity of each firm is less than β_2 for any given cluster because each subset observes a net variation of *capexratio* that is less than one.

In order to make the estimates more intuitive, both in comparing firms intra-Group and across Groups, we make an accounting change to the measure of *capexratio*. We convert the variable from its original units into standard units, relative only to firms within its own Group "g." We call this converted variable $\widetilde{capexratio}$, which is defined as follows

$$\widetilde{capexratio}_{i,t-1} = \frac{capexratio_{i,t-1} - \overline{capexratio_g}}{\sigma_g} \tag{6}$$

and $\overline{capexratio_g}$, σ are defined as

$$\overline{capexratio_g} = \frac{1}{n_g T} \sum_{i=1}^{n_g} \sum_{t=1}^{T} capexratio_{i,t-1} \tag{7}$$

$$\sigma_g = \sqrt{\frac{1}{n_g T} \sum_{i=1}^{n_g} \sum_{t=1}^{T} (capexratio_{i,t-1} - \overline{capexratio_g})^2} \tag{8}$$

By converting the regressor into standard units, the variation in capital intensity captured from a 1 unit change is guaranteed to be observed within the cluster. The convenience of using this system of units is that β_2 now measures the difference in equity sensitivity between relative firms within a Group, and not for arbitrary 1-unit changes in capital intensity. Hence, the new estimate better represents the magnitude of sensitivity variation within a given Group and allows for more intuitive comparisons to be drawn. For similar comparability, we convert the *cash* regressor into standard units as well. As this method is only a change in units, the t-statistic and significance of a given estimate does not change. Table 6 on the following page describes the results of running the following regression on groups 1, 2, and 3 independently.

$$\begin{aligned} return_{i,t} = constant + \beta_1 mp_t + \beta_2(mp_t * \widetilde{capexratio}_{i,(t-1)}) + \beta_3 \widetilde{capexratio}_{i,(t-1)} \\ + \beta_4(mp_t * \widetilde{cash}_{i,(t-1)} + \beta_5 \widetilde{cash}_{i,(t-1)} + \alpha_i + error_{i,t} \end{aligned} \tag{9}$$

In addition, we include the results derived in original units in regressions (1), (2), and (3) as the marginal effects are still economically meaningful.

Table 6:

	Dependent variable:					
	return					
	Original Units			Standard Units		
	(1)	(2)	(3)	(4)	(5)	(6)
mp	−6.546***	−6.210***	−7.016***	−6.495***	−6.709***	−7.104***
	(0.224)	(0.235)	(0.226)	(0.218)	(0.213)	(0.218)
mp * capexratio	−22.351	−23.658***	−7.977***			
	(14.223)	(6.972)	(1.716)			
capexratio	0.078	−1.113*	−1.076***			
	(1.583)	(0.577)	(0.159)			
mp * cash	0.065	−0.520***	−0.062			
	(0.069)	(0.117)	(0.045)			
cash	0.037***	0.035**	−0.0003			
	(0.011)	(0.017)	(0.004)			
mp * capexratioSU				−0.345	−0.811***	−1.193***
				(0.219)	(0.239)	(0.257)
capexratioSU				0.001	−0.038*	−0.161***
				(0.024)	(0.020)	(0.024)
mp * cashSU				0.204	−1.164***	−0.353
				(0.216)	(0.263)	(0.256)
cashSU				0.116***	0.078**	−0.002
				(0.034)	(0.037)	(0.024)
Group 1	X			X		
Group 2		X			X	
Group 3			X			X
Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Standard Units	No	No	No	Yes	Yes	Yes
Observations	14,106	16,413	16,505	14,106	16,413	16,505
R^2	0.062	0.060	0.065	0.062	0.060	0.065

Note: This table reports the results of regressing 1-day returns of the S&P 500 constituents on the unanticipated portion of the Federal funds rate policy action, the interaction (mp * capexratio), capexratio isolated, the (mp * cash) interaction, and cash isolated. This regression is run separately onto each Group of firms, and then run again on the independent variables measured in standard units. The full sample ranges from February 1994 to December 2009, excluding the release of September 17, 2001, and including 405 unique companies.

In line with expectations from Bernanke and Kuttner (2005), the coefficient of variable *mp* is negative and significant at the 1% level for all three groups; increases in the policy rate unambiguously lower equity value. Additionally, the consistent magnitude of this coefficient across groups and across prior results suggests a strong estimation that isn't affected much by sampling variation.

Marginal Effects

Novel to this analysis, instead, is the heterogeneous relationship between conditional stock return sensitivity and capital intensity across the three Groups. Prior to analysis, we note the difference in estimates between the regressions on *capexratio* and $\widetilde{capexratio}$: in absolute terms, the size of each estimate for β_2 is substantially large when using *capexratio* than the latter. As mentioned prior, this magnitude of effect size may be misleading intuitively: regression (2)'s estimate of -23.66 is prorated for a 1-unit change in *capexratio*, or an increase in capital expenditures equal to one quarter's net sales. The total variation in *capexratio* in a given cluster of firms, however, never exceeds 0.4. Hence, these estimates provide poor intuition of the amount of variation in *return* actually explained by capital intensity within a given cluster. Regardless, these estimates do reflect the marginal sensitivity that conditional stock returns experience from very small changes in capital intensity. The strongest relationship exists within Group 2, firms with median levels: for a hypothetical 25bp hike by the FOMC, firms with capital intensity *k + 1% of quarterly sales* would typically experience a 23 basis point additional fall in returns compared to a firm with capital intensity *k*. These firm's stocks are roughly three times more sensitive than their less capital intense counterparts - the effect is economically large and significant at the 1% level.

This relationship is less pronounced at higher and lower levels of capital intensity. The marginal effect estimated by regression (3) for firms in Group 3 is smaller in magnitude; firms with the highest levels of capital intensity generate a β_2 estimate of only -7.09, a third of the marginal effect observed in Group 2. In a similar hypothetical as presented above, the more capital intense firm experiences an additional fall in returns of roughly 7bp, a difference that is significant at the 1% level. Across non-capital intense firms, varied capital expenditure levels have a trivial effect on their equity responsiveness to changes in the unanticipated Federal funds rate.

Average Intra-Group Effects

Regressions (4) - (6) provide for a different interpretation. When using standard units instead to compare firms within each cluster, the relative magnitude of the *capexratio* interaction coefficient falls. Despite the marginal effect being three times larger than that observed in Group 3, the net change in equity sensitivity from a one standard deviation increase in capital intensity (both from conditional and unconditional effects) is only around 0.85%, compared to the 1.35% change across Group 3. These estimates are statistically significant and better illustrate the true variation in sensitivity within clusters 2 and 3. The marginal and average changes in sensitivity differ due to the large differences in *capexratio* variation across groups. For example, a standard deviation of *capexratio* in Group 3 is roughly 0.4, ten times larger than that of Group 2; this effect is large, and illustrates why a strong marginal effect does not necessitate a greater difference in averages. Consistent with the conclusion above, the estimates for Group 1 are statistically insignificant.

Interpretation

The important relationship that results from this set of observations is that firm stock returns, conditional on monetary policy, are most responsive to variation in capital expenditures in median capital intensity firms, tapering off in either extreme direction. At low levels, additional expenditures are trivial and at high levels, the effect is dampened. We suspect that the reason for this phenomena is associated with the varying demand for, or access to funding in the capital markets. Capital expenditures make up too small a portion of spending amongst firms in Group 1 for the market to reflect small increases or decreases. Hence, we observe a negligible relationship between the *capexratio* and equity effect of FOMC announcements across these firms. On the reverse side of this curve, capital intense firms are likely less affected by changes in capital expenditures due to effects akin to economies of scale. When firms become so capital reliant, we expect them to develop efficient means of accessing capital funds and thus afford to increase capital expenditures with lesser exposure in the equity market. This cluster of firms does observe an increase in responsiveness, but at a third of the rate observed by median firms. The firms in the middle experience the brunt of the effect; an increase in capital expenditures must be traded off with significant increase in exposure, on average. The residual effects that stiffen stock reactions to monetary policy witnessed in high and low capital firms are not available to the middle tertile. Capital expenditures are large enough to generate investor reactions but too small to benefit from the efficiency of scale.

A few additional points worth noting about the results presented in Table 6: we see that the average equity price sensitivity to monetary policy is greatest in Group 3, then Group 2, and then Group 1. Converting units into standard deviations has a convenient implication: the monetary policy shock becomes the only regressor with a non-zero mean, and hence provides a comparable statistic across Groups. This ecological relationship is consistent with our primary conclusion that higher capital intensity is associated with greater stock return sensitivity.

From regressions (2) and (5), we also observe that firm levels of cash intensity have a small marginal contribution to the measured heterogeneity but a relatively large effect when compared to other firms within Group 2. This suggests that the high estimate in regression (5) is caused by a large variability in the level of cash intensity within the group, and that cash intensity otherwise plays a more minor role in determining stock return heterogeneity.

VI. Conclusion

The fact that monetary policy disproportionately affects firms is not novel; Ehrmann and Fratzscher (2004) showed heterogeneity stemming from financial constraints, size of cash flows, and other characteristics back in 2004. What this analysis has shown, which has not been studied before, is that the heterogeneity can be strongly linked to a firm's level of capital intensity and quantified with statistically significant results. The framework is simple and allows for straightforward interpretations – that higher capital intensity tends to generate higher monetary policy exposure. The objective of this paper is to build upon prior literature and better pinpoint what this heterogeneity actually looks like across firms so that we may potentially forecast the winner and losers of monetary policy.

These results have important policy implications. In a time where monetary policy has become an increasingly critical macroeconomic tool, the equity or inequality resulting from action should be scrutinized. Normative analysis suggests that policymakers should take into account what distributional effects their policy would likely generate in the economy; they should recognize that a 25 basis point hike will lead to substantially different results for individual firms. Quantifying the magnitude, direction, and heterogeneous behavior of this response is critical to answering the question of who wins and who loses in the game of monetary policy. The difference, as we have shown, is economically large and have important implications on the performance of firms across the US. We hope that moving forward, policy is considerate of this fact and designs action around our and prior contributions.

References

[1] Gorodnichenko, Yuriy, and Michael Weber. 2016. "Are Prices Costly? Evidence from the Stock Market." *American Economic Review* 106(1): 165–199.

[2] BERNANKE, B.S. and KUTTNER, K.N. (2005) "What Explains the Stock Market's Reaction to Federal Reserve Policy?" *The Journal of Finance* 60: 1221–1257.

[3] Gert Peersman and Frank Smets. 2002. "The Industry Effects of Monetary Policy in the Euro Area" *Working Paper,* No. 65. August. European Central Bank Working Paper Series.

[4] Michael Ehrmann. 2000. "Firm Size and Monetary Policy Transmission: Evidence from German Business Survey Data" *Working Paper,* No. 21. May. European Central Bank Working Paper Series.

[5] Gerald A. Carlino and Robert H. DeFina. 1996. "Does Monetary Policy Have Differential Regional Effects?" *Business Review,* Federal Reserve Bank of Philadelphia, pp. 17-25.

[6] Michael Ehrmann and Marcel Fratzscher. 2004 "Taking Stock: Monetary Policy Transmission to Equity Markets" *Working Paper,* No. 354. May. European Central Bank Working Paper Series.

[7] Joseph Cafariello. 2015. "How Utilities Respond to Interest Rate Changes" *Wealth Daily.*

[8] Krueger, Joel T., and Kenneth N. Kuttner. 1996 "The Fed funds futures rate as a predictor of Federal Reserve policy" *Journal of Futures Markets* 16, 865-879.