

THE UNIVERSITY OF CHICAGO

ESSAYS IN FINANCIAL ECONOMICS

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE UNIVERSITY OF CHICAGO
BOOTH SCHOOL OF BUSINESS
AND
THE FACULTY OF THE DIVISION OF THE SOCIAL SCIENCES
KENNETH C. GRIFFIN DEPARTMENT OF ECONOMICS
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

BY
BENJAMIN MARROW

CHICAGO, ILLINOIS

JUNE 2025

To my grandmother, Joanna S. Rose

TABLE OF CONTENTS

LIST OF FIGURES	vi
LIST OF TABLES	vii
ACKNOWLEDGMENTS	viii
ABSTRACT	x
1 THE PRICING OF EARNINGS NEWS	1
1.1 Introduction	1
1.2 A Model of Investor Demand	10
1.2.1 Investor Demand	10
1.2.2 Market Clearing	13
1.2.3 Estimation	17
1.3 Data and Summary Statistics	21
1.4 Results	23
1.4.1 Market-Level Estimates	23
1.4.2 Estimates by Sector	26
1.4.3 Decomposition of the Earnings Passthrough	30
1.5 Asset Pricing Implications	36
1.5.1 Heterogeneous Earnings Passthrough	37
1.5.2 Earnings Surprises and Misreaction	41
1.6 Extending the Approach to Other Variables	52
1.7 Conclusion	54
1.8 Appendix	56
1.8.1 Specialized Model with Homogeneous Earnings Passthrough	56
1.8.2 Bootstrap Details	58
1.8.3 Dynamic Model and Earnings Momentum	59
1.8.4 Additional Figures	63
2 INFLATION EXPECTATIONS AND STOCK RETURNS	65
2.1 Introduction	65
2.2 Data	72
2.3 Stocks prices rise hand-in-hand with inflation expectations	74
2.4 Causal Effect of Changes in Inflation Expectations	79
2.4.1 Identification	80
2.4.2 CPI release day summary statistics	81
2.4.3 Causal estimates	82
2.4.4 Asymmetric Responses: Disinflation Risk	87
2.4.5 Inflation Surprises vs. Inflation Expectation Surprises	89
2.4.6 Changes in beliefs about future expectations or changes in term premia?	95
2.5 Disentangling the Channels	95

2.5.1	Market factor implied β_π	96
2.5.2	VAR decomposition of β_π	97
2.5.3	Mediation Regression Decomposition of β_π	101
2.5.4	Expected dividend channel and cross-sectional β_π	106
2.6	Towards an Economic Interpretation	107
2.6.1	Inflation Expectations and Future Output	109
2.7	Conclusion	111
2.8	Appendix	114
2.8.1	Decomposition Result	114
2.8.2	Robustness to Changes in Term Premium	116
2.8.3	Robustness to Changes in Probability of Hitting Zero Lower Bound	120
2.8.4	Robustness of Volatility Response to Inflation Expectations	122
2.8.5	Survey Expectation Changes and Inflation	122
REFERENCES		126

LIST OF FIGURES

1.1	R^2 of Quarterly Log Returns on Characteristics	5
1.2	Aggregate Price and Earnings Elasticities	25
1.3	Earnings Elasticities by Sector	27
1.4	Price vs. Earnings Elasticities by Sector	29
1.5	Reduced-form vs. Structural Earnings Elasticities	31
1.6	Time Series of Raw Earnings Passthrough	32
1.7	Earnings Passthrough Decomposition	35
1.8	Improvements to R^2 from Elasticity Measures	42
1.9	Portfolio Timing	46
1.10	Future Returns by SUE and Beta Quintile	48
1.11	Earnings-Sensitivity Portfolio Returns	49
1.12	Bias and Precision of Standardized Unexpected Earnings	63
1.13	Discrepancy Between $\tilde{\beta}_{St}$ and β_{St}	64
2.1	Inflation Betas Across 30 Fama-French Industries	78
2.2	Rolling Inflation Expectations Beta	79
2.3	CPI day regression DAG	81
2.4	Inflation Betas Across 30 Fama-French Industries	85
2.5	Reversal	86
2.6	Realized Inflation vs Inflation Expectations	93
2.7	Market beta implied estimates vs. direct estimates of β^π	98
2.8	Inflation Betas By Fama-French Industries with and without Interest Rate Controls	105
2.9	Inflation Betas by Total Payout Ratio and Leverage	108
2.10	Impulse Responses to an Increase in Expected Inflation	112
2.11	Breakeven inflation Fama-Bliss regressions	119
2.12	Impulse Responses to an Increase in Expected Inflation	123

LIST OF TABLES

1.1	Earnings Passthrough and Stock Level Sensitivities	40
1.2	Earnings Sensitivity Portfolio Returns	50
1.3	Price Sensitivity Portfolio Returns	52
2.1	Summary statistics of main measures	74
2.2	The exposure of stocks to changes in inflation expectations	76
2.3	Inflation expectation measures CPI vs non-CPI days	82
2.4	CPI day effect of changes in inflation expectations on returns	83
2.5	Inflation Beta by 2x2 Sort	88
2.6	Baseline Regression for November 2021 to October 2023 Sample	92
2.7	Inflation Surprises vs. Inflation Expectation Surprises	94
2.8	VAR decomposition of excess return response	100
2.9	Decomposition of Inflation Premium by Channel	104
2.10	Inflation Expectations Premium After Controlling for Term Premium	120
2.11	Expected Inflation Beta by Change in ZLB Probability (2018-2019)	121
2.12	Response of SPF Forecast Revisions to Innovations in Inflation Expectations . .	125

ACKNOWLEDGMENTS

Alfred Tennyson’s *Ulysses* observes that “I am part of all that I have met.” Reflecting on these words, I feel profoundly grateful for the many outstanding individuals I have had the fortune to meet and who have supported me throughout this long and winding PhD journey.

First, I want to express my deepest gratitude to my dissertation committee: Stefan Nagel, Ralph Koijen, Niels Gormsen, Lubos Pastor, and Lars Hansen. I could not have asked for a more intelligent committee, and their advice, support, and criticism were invaluable during my time in the program. Beyond their immediate impact on my academic research, each has structured my way of thinking: I thank Lars for forcing me to think about the real economy; Lubos for forcing me to think about the real world; Ralph for pushing me to ask ambitious questions; Niels for encouraging me to think deeply about what we can learn from the answers to those questions; and above all, Stefan, for reminding me that economics is ultimately a study of human behavior, and all of the expectations, learning, preferences, and imperfections that accompany it. I am particularly grateful that my advisors always treated me as a colleague and peer; I still have much to learn from each of them.

My committee is one subset of a faculty that has been extraordinarily generous with its mentorship. I want to also thank Sam Hartzmark, Constantine Yannelis, Amir Sufi, Steve Kaplan, Anil Kashyap, and Zhiguo He, who were supportive of my journey in various ways. Not all mentorship is scholarly in nature; I’m grateful to those professors who offered words of encouragement or simply checked in on me when they did not have to. I thank the PhD office staff, especially Malaina, Cynthia, and Kim, for helping me navigate the ins and outs of the program. And I remain indebted to the teachers and mentors I had before graduate school, including Blake Eldridge, Doug McKee, and Fernando Duarte, who taught me how to think and helped spark the passion for learning that ultimately led me here.

It is sometimes joked that, in a variant of the Hobbesian quip, the PhD is “nasty brutish and *long*.” Fortunately, it is made less so by the constant support—intellectual and otherwise—

of my classmates and colleagues. Beyond the valuable companionship they provided, Aditya, Federico, Filippo, and Zhiyu were remarkably patient in explaining economics to me time and again and answering my many questions. I learned a great deal from my collaborators Manav and Sangmin, not only from what they taught me directly but also from the example they continue to set about communication, framing, and clear thinking. And the denizens of the Room 384 ecosphere, including Ching-Tse, David, Jingoo, Jingtao, Jessica, Cong, Yueyang, and Pengyu were a constant source of camaraderie, learning, and discussion.

I certainly would not have survived the PhD without the support of my peers outside of my program. I could always count on Brandon and Jordan to hear me out, and Patrick Adams was a voice of sanity, support, and humor when I frequently needed one. I was especially lucky to have David Crane, Joshua Altman, Philip Groenwegen, and William Ezekowitz, who provided friendship in the truest sense of the word. There are of course many more individuals to whom I owe thanks for their counsel and friendship, and I hope they know the crucial role they played in this chapter of my life.

Finally, I want to extend my unconditional gratitude to my family. Perhaps the one challenge of having such a large and loving family is that it is difficult to acknowledge everyone individually, but I am deeply grateful to “the whole mishpochah”. I want to give special thanks to Wendi Rose, whose unwavering support of my academic, professional, and personal life during these years meant more than she knows, and to Gideon Rose, whose curiosity and encouragement sustained me. I cannot ever thank my parents enough, James Marrow and Emily Rose, for the unconditional support and love they have given me my entire life. I would not be where I am or who I am without them. I would like to conclude by recognizing my late grandmother, Joanna S. Rose, who was an undying source of love and a model of intellectual curiosity. I think of her daily, but especially when reflecting on my own experience in the PhD, I think of the words inscribed on her epitaph: “she was never happier than when she was learning.” This dissertation is dedicated to her memory.

ABSTRACT

This dissertation consists of two chapters studying how news is priced in equity markets.

Chapter 1 studies how earnings news becomes priced into stock returns. I use a demand system approach to show that this passthrough depends on investor responses to both earnings and prices and that these sensitivities are heterogeneous across investors. A key identification challenge is that earnings news is rapidly incorporated into prices; as a result, it is difficult to distinguish whether investors react to the earnings news itself or the concurrent price change. Using a two-step procedure to isolate price from earnings responses, I identify an average asset-weighted earnings elasticity of 3, i.e. for a stock that beats earnings expectations by 1%, the average investor would increase his number of shares held by 3% if prices were held fixed. These estimates vary across sectors, with most institutional investors more earnings elastic and price inelastic compared to the residual ("household") sector. The stock-level sensitivities implied by their ownership account for heterogeneous earnings passthroughs, as stocks with higher earnings sensitivity and lower price sensitivity see larger return responses from the same earnings surprise. Extremes of price and earnings elasticities are also closely related to misreaction: a strategy that bets on subsequent reversal (momentum) insensitive (insensitive) stocks in response to earnings news generates significant outperformance and alpha. These findings suggest that the pricing of earnings news is closely related to the ownership structure of stocks.

Chapter 2 studies how inflation expectations affect stock returns, and what accounts for this relationship. I directly measure investors' expectations using traded inflation-indexed contracts and show that, post-2000, stocks offer positive returns in response to higher expected inflation: unconditionally, a 10 basis point increase in 10-year breakeven inflation is associated with a 1.1% increase in the value-weighted stock index. Using a wide range of approaches, I show that this positive relationship is almost entirely due to aggregate variations in expected excess returns rather than changes in firm cash flows (e.g., due to

higher mark-ups) or fluctuations in risk-free rates (e.g., due to expected monetary policy response). Overall, a risk premium “proxy” mechanism appears to explain this dominant role of expected excess returns: higher long-term inflation expectations signal stronger future economic growth and reduced volatility.

CHAPTER 1

THE PRICING OF EARNINGS NEWS

1.1 Introduction

Models of asset pricing typically specify a close link between stock fundamentals (dividends, earnings, etc.) and valuations. Because these models imply that shocks to, or news about fundamentals translate into changes in prices, they raise a set of empirical questions about how precisely fundamental news becomes incorporated into stock returns. In this paper, I adopt a demand-system approach to study the pricing of earnings news – an organizing source of fundamental news about firms – in the stock market. I show that investor-level responses to earnings news depend on elasticities to prices and fundamentals separately, and that these elasticities are heterogeneous over time and across investors.

Moreover, because stocks vary in their ownership structure, stocks inherit the sensitivities of their investors: in the cross-section, stocks with investors who are more earnings sensitive, or less price sensitive, see a larger price impact from an earnings surprise of a given magnitude. This differential pricing is not a result of investors sorting across stocks that differ along other dimensions that explain these cross-sectional responses, but rather a feature of market mispricing: long-short directional strategies that bet on *reversal* in stocks expected to overreact to fundamental news (i.e., earnings sensitive or price insensitive stocks), and on *momentum* in stocks expected to underreact (i.e., earnings insensitive or price sensitive stocks) generate a four-factor alpha on the order of 1.0-1.7% per quarter. The stock-level analysis implies that the ownership structure of stocks plays a key role in the pricing of earnings news, and that variation in both investor earnings and price sensitivities closely reflects misreaction to fundamental news.

The challenge in identifying investor-level responses to earnings is that such news is contemporaneous with returns induced by *other* investors' responses to those surprises. In

short, an investor’s general-equilibrium response to earnings news embeds both his response to the news itself (the “fundamental” response) as well as his response to the aggregate price impact (the “price” response). To understand the challenge this poses for identification, consider the case of a price momentum trader who pays no direct attention to earnings news and instead trades on the past performance of various stocks. Such an investor will appear to buy those stocks that beat earnings expectations (and sell those that underperform), but only because stocks that beat earnings tend to rise in price as well. A naive regression of this investor’s portfolio weights on earnings surprises will deliver a positive coefficient, even as the investor in question does not, by assumption, care about earnings news. A potential alternative solution – controlling for returns – is inadequate, since returns are endogenous to the quantity responses of *all* investors, who may share the demand function of the investor analyzed. In other words, controlling for returns in this case introduces the “bad control” problem that arises when the control variable (stock returns) is itself a consequence of the treatment (earnings surprises) (Angrist and Pischke (2009)).

I instead address the identification problem through a two-step procedure that first adjusts an investor’s trades to account for how that investor responds to prices. Intuitively, consider two investors who are equally sensitive to an earnings surprise, i.e. believe it has the same implications for future earnings and dividend growth, while the rest of the market believes it has a larger impact. The investors’ total trading response will then depend not only on their elasticity to the earnings news, but also on how they trade against the perceived overvaluation induced by the pricing wedge between their own perception of the earnings news and the market’s. An investor who trades actively against perceived overvaluations – a price elastic investor – will reduce his net purchases of the asset by more than an equally earnings-sensitive investor who is less active in his trading against prices. Hence these two investors, who share the same earnings elasticity, may have different general equilibrium responses to the earnings news due to differences in their respective price elasticities.

More formally, I start with a simple model in which an investor’s portfolio choice is expressed as a log-linearized function of earnings surprises and price movements (i.e., returns). Through the logic of market clearing, I show that the coefficient from a simple regression of quantities on earnings surprises equals the investor’s true earnings elasticity, plus a price response: the investor’s price elasticity, times the market-level price impact of the earnings surprise. To recover the true fundamental elasticity, it is necessary to net out the investor’s response to the aggregate price change. While the formulation is specific to the linear model I write down, the necessity of adjusting trading responses for the associated market-induced price response applies more generally, both to alternative specifications of demand and a broader set of news and characteristics.

Estimating an investor’s price response requires an estimate of price elasticities, i.e., how an investor’s portfolio holdings respond to a pure price change unassociated with other changes in fundamentals. I identify price elasticities using the market equity instrument from Kojien and Yogo (2019). This instrument generates cross-sectional variation in prices by exploiting variation in the investment universes of large investors: stocks that are in the investment universes of more (or wealthier) investors have higher prices *ceteris paribus*, and the relationship between investors’ portfolio weights and this implied market equity provides an estimate of their sensitivity (elasticity) to prices. Investors that tilt their holdings further away from stocks that have exogenously higher market equity due to their larger presence in the investment universes are more price elastic. This price elasticity provides an estimate of how an investor would respond if the price moved without any corresponding change in fundamentals. Any systematic residual change in holdings beyond this predicted amount can then be attributed to an investor’s sensitivity to earnings news itself.

I estimate an earnings-augmented demand system on the full panel of 13-F investors 1980-2017. The average asset-weighted earnings elasticity is approximately 3, but ranges from slightly greater than 0 (in the late 1980s) to nearly 7 (in 2014). In other words, a


stock that beats its expected earnings per share by 1% sees an asset-weighted increase in demand of around 3%. However, these responses vary significantly in the cross-section. Traditional institutional investors, including mutual funds and pension funds, tend to have the strongest sensitivity to earnings but are fairly price inelastic, while the 13-F residual “household” sector – the most price elastic of all investor types – is the least responsive to earnings. Meanwhile, investment advisers – which include large hedge funds – are price and earnings elastic. These investor-type estimates suggest that price sensitivity and earnings sensitivity are not well described as pure complements or substitutes: within institutional investors, they are positively correlated, but are negatively correlated unconditionally.¹

In principle, the procedure I introduce for estimating earnings elasticities is not limited to earning news; it can be applied to any non-price characteristic. An investor’s response to ESG-related news, for example, can be decomposed into their true sensitivity to ESG news plus an adjustment for how the market as a whole prices the ESG news. I focus on earnings news specifically not only for the prominent role that news about fundamentals plays in asset pricing models, behavioral theories, and earnings-related anomalies, but also for the simple empirical fact that earnings news accounts for a large component of variation in returns. To illustrate this point, figure (1.1) plots the time series of R^2 from quarterly cross-sectional regressions of log returns on various characteristics. In most periods, the variation in returns that can be attributed to characteristics associated with the 5 Fama-French factors (shown in the dotted line) is dominated by that from a single non-price characteristic, the standardized unexpected earnings (SUE) (shown in the light blue). Yet another measure of earnings news, the three day return around earning announcements days, is striking not only for its relative dominance, but also for its ability, in absolute terms, to explain low frequency returns: close

1. To clarify terminology, price elasticity by convention is defined as the *negative* partial derivative of quantities with respect to prices $-\partial \log Q / \partial \log P$, while earnings elasticity is defined as the *positive* partial derivative of quantities with respect to earnings news: $\partial \log Q / \partial SUE$. In other words, a price elastic investor shifts his holdings *away* from stocks that have exogenous increases in price, while an earnings elastic investor shifts his holding *toward* stocks with exogenous increases in earnings news.

to 30% of variation in quarterly log returns is explained by the returns that span one day on either side of earnings announcements. A clear implication of this result is that earnings news must play an important role in investor demand for it to feature so saliently in aggregate returns.

Figure 1.1: R^2 of Quarterly Log Returns on Characteristics



This figure shows the cross-sectional within-adjusted R^2 from three different models. The dotted line shows the in-sample R^2 from 5-quarter panel regressions of $R_{nt} = \alpha_t + \beta'_t \mathbf{X}_{nt} + \epsilon_{nt}$, where $\mathbf{X}_{nt}$ includes CAPM beta, asset growth, profitability, log book equity, log dividends-to-book-equity, and where R_{nt} is the return over the quarter. The light blue line shows the in-sample R^2 from 5-quarter panel regressions of $R_{nt} = \alpha_t + \beta_t SUE_{nt} + \epsilon_{nt}$. Finally, the dark blue line shows the in-sample R^2 from 5-quarter panel regressions of $R_{nt} = \alpha_t + \beta_t Ret_{nt}^{(-1,1)} + \epsilon_{nt}$, where $Ret_{nt}^{(-1,1)}$ is the total cumulative return of stock n from the open price the trading day prior to the earnings announcement through the close the trading day following the earnings announcement.

An advantage of using the demand system approach is that it gives interpretation to the earning passthrough, which has been a long object of study in asset pricing and accounting literature. In the linearized model I use, the earnings passthrough, i.e., the cross-sectional coefficient from a regression of log returns on standardized earnings surprises, equals an average ratio of two share-weighted elasticities: the asset-weighted earnings elasticity divided

by the asset-weighted price elasticity. Intuitively, the earnings passthrough is larger if investors in aggregate become more responsive to earnings news (higher earnings elasticity) *ceteris paribus* or less responsive to price changes (lower price elasticity). I show that the ratio of these share-weighted elasticities, estimated from holdings data, well approximates the magnitudes and trends in the aggregate earnings passthrough estimated from simple pricing regressions. A longstanding question in asset pricing asks why the earnings passthrough varies over time; my approach allows me to decompose the time-series variation into the respective contributions of asset-weighted earning elasticities and price elasticities. The vast majority of the time-series variation in this coefficient is driven by investors' changing response to fundamentals rather than changing price elasticities.

In the second half of the paper, I connect measures of investor-level sensitivities to stock-level sensitivities. These measures are constructed as the weighted average sensitivities of a stock's investors, with weights equal to the ownership shares of the stock. I first show that these measures explain heterogeneous earnings passthroughs in the cross-section. Consistent with the implications of the model, in which the passthrough equals the ratio of stock-specific earnings elasticities to price elasticities, stocks with higher earnings sensitivity, or lower price sensitivity, see larger returns in the same quarter from a given earnings surprise. These effects are not spanned by many characteristics associated with heterogeneous passthroughs (such as size, value, leverage and market beta), and generate large relative improvements in R^2 compared to models that only use these stock characteristics to explain the passthroughs. The asset pricing results suggest that the demand-system measures capture important characteristics about the stock that are important for explaining the cross-sectional pricing of earnings news.

A natural question is whether these heterogeneous responses to earnings surprises are caused by variation in ownership, or are merely associated with it. For example, the observed cross-sectional results could occur if there are unobserved characteristics that are

simultaneously associated with larger passthroughs and with the types of investors that hold the stock. Under this account, heterogeneous responses to earnings news could reflect rational market pricing insofar as they capture these omitted variables. If earnings news is more predictive of future earnings growth for tech stocks, for example, and if earnings elastic or growth-oriented investors (such as hedge funds) hold a comparatively larger share of tech stocks, then a larger response to earnings news for these stocks could be consistent with rational pricing.

I instead show that stock-level differences in ownership are more consistent with mispricing through over- and under-reaction. If stocks with the most earnings sensitive investors overreact to earnings news, the prices of these stocks may be expected to revert in the subsequent period. Likewise if stocks with earnings insensitive investors underreact, these stocks may be expected to exhibit momentum. To test this mechanism, I form two portfolios to capture earnings misreaction. These portfolios are constructed by first taking the intersection of 5 cross-sectional sorts on the earnings surprise (SUE) and 5 cross-sectional sorts on stock sensitivity – either for earnings or price. The “earnings sensitivity” portfolio itself consists of two long-short portfolios to capture misreaction: one part of the earnings portfolio bets on momentum in stocks expected to underreact due to earnings insensitive investors. That is, it goes long in stocks with good earnings news with earnings insensitive investors – those expected to increase further – and short stocks with bad news and earnings insensitive investors – those expected to decrease further. The second earnings portfolio bets on subsequent reversal in stocks expected to overreact due to earnings over-sensitive investors. It goes long in stocks with bad earnings news and earnings sensitive investors (stocks presumed to have overreacted to the bad news) and short stocks with good news and sensitive investors (those presumed to have overreacted to the good news). The combined “earnings sensitivity” portfolio, constructed as the average return on these two portfolios, generates quarterly alpha of 1.7% equal weighted and 1.0% value weighted.

I also form a corresponding “price sensitivity” portfolio, constructed by intersecting earnings news sorts with stock *price*-sensitivity sorts. My model suggests that very price sensitive stocks should underreact to earnings news: price-elastic investors trade against any earnings-induced price movement they disagree with, dampening the contemporaneous response, but potentially generating a future “correction”. Meanwhile, for price-inelastic stocks, there may not be enough capital to trade against earnings-induced price movements, leading to a contemporaneous overreaction corrected through a reversal pattern the following quarter. This logic calls for one portfolio that is long price-elastic stocks that have received good news and short price-inelastic stocks that have received good news, and a second portfolio that is long price-inelastic stocks that have received bad news and short price-elastic stocks that have received bad news. To ensure that both the earnings and price sensitivity portfolios are not driven by cross-sectional correlations of stock-level price sensitivity with stock-level earnings sensitivities, I first residualize each elasticity measure on the other before forming the portfolio. The combined “price sensitivity” portfolio – which does not use data on investor responses to earnings news, except in the orthogonalization step – generates quarterly alpha of 1.2% equal-weighted and 1.0% value-weighted.

This paper is connected with several strands of literature, most closely to the literature on Demand-Systems Asset Pricing (DSAP) that endeavors to jointly explain asset prices alongside investor holdings (demand). As with my paper, Koijen and Yogo (2019), Koijen, Richmond, and Yogo (2023), Huebner (2023), and Van Der Beck (2022) among others, estimate demand functions for 13-F institutions, including the estimation of price elasticities over time and across investors. My contribution to this literature is two-fold: first, I introduce a new variable to the asset demand system – earnings surprises – that explains a greater proportion of returns than the typically modeled characteristics. Additionally, using the logic of demand systems, I closely illustrate the identification problem from reduced-form pricing regressions that do not account for price-inelastic investors.

My paper also addresses the behavioral finance literature that studies the existence, magnitude, and identity of “price” vs. “fundamental” investors. On the theoretical side, behavioral models typically model agents as *either* responding to fundamental (Barberis, Shleifer, and Vishny (1998), Fuster, Hebert, and Laibson (2012), Choi and Mertens (2019), Hirshleifer, Li, and Yu (2015)) *or* as responding to prices (Cutler, Poterba, and Summers (1990), Hong and Stein (1999), Barberis, Greenwood, et al. (2015)). On the empirical side, researchers use survey data (Dahlquist and Ibert (2023), Greenwood and Shleifer (2014)) to elicit beliefs and expectations, or study the behavior of specific subsets of investor like retail traders (Laarits and Sammon (2022)). Relative to this literature, I simultaneously estimate price vs. fundamental responses; I use trades rather than reported beliefs to study the passthrough of market-relevant beliefs into prices; and I study heterogeneity across a wider cross-section of investors (Banks, Mutual Funds, Investment Advisors, Pension Funds, Insurance Companies, and the residual “household” sector).² I also improve upon earlier estimation methods that either do not take into account the endogeneity of prices, or that naively control for returns to partial out the price response. I show that these approaches can misclassify the direction of the response to earnings news of entire sector types, such as the household sector.

My paper also relates to a long accounting and finance literature studying asset price responses to earnings news, and investor-level responses (e.g., Livnat and Mendenhall (2006); Bernard and Thomas (1989); Skinner and Sloan (2002); Barber et al. (2013); McClure and Nikolaev (2023)). Relative to papers that study asset prices at high frequency around earnings news, or even medium frequency (e.g. post-earnings announcement drift), I focus on the very low frequency component of earnings news that shows up in returns and holdings quarter to quarter. The method I use for constructing stock-level sensitivities is closely related

2. The elasticities I measure reflect contemporaneous responses that occur over the same quarter, whereas behavioral finance literature often focuses on “extrapolation,” i.e. responses to past innovations. In appendix section (1.8.3), I develop a dynamic model that could be used to estimate the full term structure of elasticities.

to the approach in Blank, Kwon, and Tang (2023), which constructs a “holding gap” based on investor responses to earnings news. Since this approach does not distinguish whether investors’ responses are due to a response to prices or to fundamentals, one contribution of the asset pricing section is to separately – and independently – estimate and form portfolios on these two different elasticities (indeed, the price sensitivity portfolio is, by construction, independent of investor responses to earnings news). Moreover, while their approach focuses on long run price run-ups and drawdowns, I focus purely on the impact of (non-price) earnings surprises. Finally, a contribution relative to this literature is to interpret earnings responses through a demand system framework, which enables both proper identification of elasticities as well as a decomposition of the passthrough.

The remainder of the paper is as follows. Section II introduces the model of demand and the identification challenge associated with identifying price elasticities. Section III introduces the data and the construction of earnings surprises. Section IV estimates the model on a panel of 13-F investors, decomposes the (log) earnings passthrough, and discusses the heterogeneity by sector. Section V analyzes the implications of the investor-level responses for stock-level asset pricing in the cross-section and time-series. Section VI discusses how to extend the identification approach to a more general set of characteristics and asset classes. Section VII concludes.

1.2 A Model of Investor Demand

1.2.1 *Investor Demand*

In this section, I introduce a model of investor demand in order to illustrate the identification problem in traditional approaches to measuring earnings sensitivity.

I begin with a log-linearized demand curve for investor i over stock n in quarter t that expresses his trades as a function of earnings surprises and returns:

$$\Delta q_{int} = \alpha_{it} + \beta_{it} SUE_{nt} - \zeta_{it} \Delta p_{nt} + u_{int} \quad (1.1)$$

Here, Δq_{int} is the change in log shares demanded by investor i in stock n at quarter t ; SUE_{nt} is the standardized unexpected earnings of stock n in quarter t (i.e., the earnings news), and Δp_{nt} is the quarterly log return. As a linearized approximation to the investor's true demand curve, this model does not impose that earnings and price news are the only determinants of an investor's demand: u_{int} captures all the other variables that drive changes in holdings.

In this model, the coefficients on earnings news and log returns, β_{it} and ζ_{it} , represent the investor's fundamental and price elasticity respectively. Crucially, they are partial derivatives of holdings with respect to characteristics: β_{it} is the percent change in demand associated with a stock that beats expected earning by 1%, holding all else – including prices – constant. ζ_{it} is the percent decrease in shares demanded associated with a stock that sees an exogenous price increase of 1%.

In standard models of portfolio choice, β_{it} and ζ_{it} are typically both nonnegative. To see this, consider the demand curve implied by an investor with CARA utility over one-period future returns. The quantity of shares demanded by investor i , Q_{it} , is given by:

$$Q_{it} = \frac{E_{it} [(P_{t+1} + D_{t+1}) / P_t]}{P_t A_{it} \sigma_{t+1}^2} := \frac{\pi_{it+1}}{P_t A_{it} \sigma_{t+1}^2} \quad (1.2)$$

where A_{it} is the investor's risk aversion and σ_{t+1}^2 is the variance of returns. Taking logs of (1.2), the log number of shares demanded is

$$q_{it} = \log(\pi_{it}) - p_t - \log(A_{it} \sigma_{t+1}^2) \quad (1.3)$$

Assume that investors form expectations over future prices and dividends using conditional expectation functions f_i and g_i :

$$E_{it} [P_{t+1}] =: f_i (P_t, SUE_t)$$

$$E_{it} [D_{t+1}] =: g_i (P_t, SUE_t)$$

In other words, investors forecast future prices and future dividends as a function of current price levels and current earnings news. Investor i 's price elasticity, ζ_{it} , is then given by:

$$\zeta_{it} := -\frac{\partial q_{it}}{\partial p_t} = \frac{1}{\pi_{it}} [(f_P + g_P) - (f + g) / \exp p_t]$$

where f_P and g_P are partial derivatives of f_i and g_i with respect to P_t . If prices do not affect expectations of future prices and future dividends, i.e. if $f_P = g_P = 0$, then the price elasticity simplifies to

$$\zeta_{it} = -\frac{1}{\pi_{it}} [-(f + g) / \exp p_t] > 0$$

Intuitively, if a stock has an exogenous increase in price, and investors do not change their views on *future* prices, dividends, or risk, then the expected returns for this stock are lower, and investors demand less, i.e. a positive price elasticity.

A similar logic applies to the earnings elasticity, defined as $\beta_{it} \equiv \partial q_{it} / \partial SUE_t$:

$$\frac{\partial q_{it}}{\partial SUE_t} = \frac{1}{\pi_{it}} (f_{SUE} + g_{SUE}) \frac{1}{\exp p_t}$$

where f_{SUE} and g_{SUE} are partial derivatives of f_i and g_i with respect to SUE_t . In many models, we may expect f_{SUE} and g_{SUE} to be greater than 0. Assume, for example, that all unexpected earnings are paid out as a dividend next quarter, then $g = E [D_{t+1} | SUE] = SUE \times P_t > 0$. Under many belief models (e.g. diagnostic expectations), agents may also change their expectation of future earnings growth, in which case they may also expect the

price to increase ($f_{SUE} > 0$). Indeed if the estimated persistence is stronger, then f_{SUE} is larger. As with price elasticities, earnings elasticities are not guaranteed to be positive for all investors, but will depend on how earnings news affect expectations of future prices and fundamentals. Inattentive investors, for example, who do not pay attention to fundamental news, may exhibit a small passthrough of news to beliefs, in which case β_{it} will be close to 0. Alternatively, investors may decrease their expectations of future fundamentals if they associate the positive earnings surprise with discretionary earnings management, or if they expect the market to overreact to the earnings news.

1.2.2 Market Clearing

I next apply market clearing across individual investor demand curves. If stocks are in fixed net supply, market clearing implies that the share-weighted changes in quantities must sum to zero. Put differently, prices must adjust such that every investor's demand is offset by the residual supply provided by other investors in the market. Let S_{int} denote the fraction of shares outstanding of stock n held by investor i at time t :

$$S_{int} := Q_{int} / \sum_i Q_{int} \quad (1.4)$$

Market clearing imposes that $\forall n, t$,

$$\sum_i S_{int} \Delta q_{int} = 0 \quad (1.5)$$

Aggregating over all investor demands from (1.1), and solving for prices, we obtain

$$\Delta p_{nt} = \frac{\alpha_{Snt}}{\zeta_{Snt}} + \frac{\beta_{Snt}}{\zeta_{Snt}} SUE_{nt} + \frac{u_{Snt}}{\zeta_{Snt}} \quad (1.6)$$

where the S subscript denotes S_{int} -weighted sums: $X_S := \sum_i S_{int} X_i$. Equation (1.6) expresses that returns are related to earnings surprises through the ratio of ownership-weighted

earnings elasticities (β_{Snt}) to similarly weighted price elasticities (ζ_{Snt}). If the market becomes more earnings elastic (higher β_{Snt}), the passthrough of earnings surprises to returns is larger. If the market becomes more price elastic (higher ζ_{Snt}), there is more capital to “undo” others investors’ increased demand from a given earnings surprise, and the earnings passthrough is lower. Because β_{Snt} and ζ_{Snt} are share-weighted elasticities, time-series variation in these coefficients need not be driven by changing elasticities of each investor; a reallocation of capital from from less elastic to more elastic investors will increase the market-level elasticities.

Substituting the expression for prices in (1.6) back into investor demand provides the reduced-form investor demand function:

$$\begin{aligned}
\Delta q_{int} &= \alpha_{it} + \beta_{it} SUE_{nt} - \underbrace{\zeta_{it} \left(\frac{\alpha_{Snt}}{\zeta_{Snt}} + \frac{\beta_{Snt}}{\zeta_{Snt}} SUE_{nt} + \frac{u_{Snt}}{\zeta_{Snt}} \right)}_{\Delta p_{nt}} + u_{int} \\
&= \underbrace{\left(\alpha_{it} - \zeta_{it} \frac{\alpha_{Snt}}{\zeta_{Snt}} \right)}_{\tilde{\alpha}_{int}} + \underbrace{\left(\beta_{it} - \zeta_{it} \frac{\beta_{Snt}}{\zeta_{Snt}} \right)}_{\tilde{\beta}_{int}} SUE_{nt} + \underbrace{\left(u_{int} - \zeta_{it} \frac{u_{Snt}}{\zeta_{Snt}} \right)}_{\tilde{u}_{int}} \\
&= \tilde{\alpha}_{int} + \tilde{\beta}_{int} SUE_{nt} + \tilde{u}_{int}
\end{aligned} \tag{1.7}$$

Equation (1.7) illustrates the principle identification challenge: the passthrough $\tilde{\beta}_{int}$ relating earning surprises to trades is not the structural coefficient β_{it} , but rather the structural coefficient plus a stock-specific price adjustment.

$$\tilde{\beta}_{int} = \beta_{it} - \zeta_{it} \frac{\beta_{Snt}}{\zeta_{Snt}} \tag{1.8}$$

The price adjustment is the investor’s response (ζ_{it}) to the price change arising from the market’s asset-weighted response (β_{Snt}/ζ_{Snt}). While (1.7) is not properly a regression (due to the presence of investor $\times$ stock $\times$ quarter coefficients), the regression analog of (1.7) suffers

from the same identification challenge. Suppose one were to estimate a regression of investor trades on earnings surprises:

$$\Delta q_{int} = \alpha_{it} + \tilde{\beta}_{it} SUE_{nt} + \epsilon_{int} \quad (1.9)$$

The coefficient $\tilde{\beta}_{it}$ from this regression would equal a weighted average of β_{int} across stocks. Taking weighted averages of (1.9), the estimated reduced-form coefficient would equal:

$$\begin{aligned} \tilde{\beta}_{it} &= \beta_{it} - \zeta_{it} \gamma_t, \\ \gamma_t &:= \dot{E}_t [\beta_{Snt} / \zeta_{Snt}] \end{aligned}$$

where $\dot{E}[\cdot]$ denotes the weighted expectation across stocks, with weights proportional to $SUE_{nt} (SUE_{nt} - \overline{SUE_{nt}})$.

In effect, two things occur simultaneously in a period when a firm realizes an earnings surprise. First, the earnings surprise, isolated from any additional changes in price, affects the investor's demand for the asset. Second, the market as a whole responds to the earnings news and drives a price change in proportion to the ratio of aggregate elasticities. If the market response is much larger than the investor's individual earnings sensitivity, the investor should decrease the shares they hold due to the perceived overvaluation (assuming positive price elasticities). But how much this investor decreases their holdings in response to overvaluation depends on her price elasticity: more elastic investors will have a greater response to the overvaluation than less elastic investors. Investors with negative price elasticities, i.e. with upward sloping demand may even *increase* the shares they hold if the market response to the earnings news is larger than their individual sensitivities. This could occur for investors who learn from prices (about fundamentals) or for momentum traders, for example.

Rearranging, we can express an investor's true structural elasticity, β_{it} , as a function of

his reduced-form coefficient, price elasticity, and market response:

$$\beta_{it} = \tilde{\beta}_{int} + \zeta_{it} \frac{\beta_{Snt}}{\zeta_{Snt}} \quad (1.10)$$

Without observing β_{it} , if the remaining terms $(\tilde{\beta}_{int}, \zeta_{it}, \beta_{Snt}, \zeta_{Snt})$ can be well-estimated, then one can recover the structural elasticity by adjusting the biased coefficient for the magnitude of the bias. Intuitively, this approach undoes the price response embedded in the reduced-form coefficient $\tilde{\beta}_{int}$.

$$\underbrace{\beta_{it}}_{\text{Earnings Elasticity}} = \underbrace{\tilde{\beta}_{int}}_{\text{Reduced-Form Response}} + \underbrace{\zeta_{it}}_{\text{Price Elasticity}} \times \underbrace{\gamma_{nt}}_{\text{Earnings Passthrough}} \quad (1.11)$$

Equation (1.11) clarifies how different types of investors may be expected to respond. For example, consider an index fund that holds the market portfolio. Since index funds hold the market portfolio, they do not trade ($\Delta q_{int} \equiv 0$), so that the holdings regression implied by (1.9) will deliver a reduced-form $\tilde{\beta}_{it} = 0$. By the same token, these funds perfectly track market equity, i.e. are perfectly price inelastic, and so $\zeta_{it} = 0$. For any aggregate earnings passthrough γ_{nt} , an index fund's estimated structural earnings elasticity is $\beta_{it} = 0$. As expected, since index fund weights are only functions of market equity, there is zero sensitivity to earnings news when holding fixed prices.

Another insight of the identifying equation is that trade is not needed to identify an investor's demand for earnings. This is clearly true for the case of the index fund, which has no earnings elasticity, but can also hold for investors with positive earnings demand. An investor who shares the market elasticities for a given stock $\beta_{it} = \beta_{St}$, $\zeta_{it} = \zeta_{St}$ will have a reduced-form response of zero:

$$\tilde{\beta}_{it} = \underbrace{\beta_{it}}_{=\beta_{St}} - \underbrace{\zeta_{it}}_{=\zeta_{St}} \times \gamma_t = \beta_{St} - \zeta_{St} \dot{E}_t \left[\frac{\beta_{Snt}}{\zeta_{Snt}} \right] \approx 0$$

A reduced-form response close to zero – i.e., in which an investor does not trade at all (or trades very little) in response to earnings news – does not imply that an investor has zero sensitivity to earnings. To the contrary, this investor by assumption shares the market’s average response to earnings news, which must be positive since the earnings passthrough is, on average, positive. Precisely because the investor agrees with the market’s elasticities, he does not need to trade at all; the price response perfectly tracks his views and make him indifferent to purchasing additional shares.

An additional insight reflected in equations (1.11) and (1.6) is that earnings passthroughs can be heterogeneous, even when each investor, by assumption, has elasticities that are common across all stocks. In the model, the earnings passthrough $\gamma_{nt} := \beta_{Snt}/\zeta_{Snt}$ depends on the stock (n), even as investor elasticities (ζ_{it}, β_{it}) do not. The heterogeneity arises because of the cross-variation in stock-ownership, which appears due to the stock-specific share-weighting of elasticities. If the ownership structure of each stock was identical across stocks, the model suggests there would be no cross-sectional variation in the earnings passthrough. If this does not hold, then earnings passthroughs are larger for stocks whose ownership consists of investors that are more earnings elastic or less price elastic. To be sure, there may be additional sources of heterogeneity in passthroughs, if, for example, investor elasticities depend on stock characteristics (e.g. industry). The log-linearized demand curve in (1.1) generalizes away from these stock-specific loadings, allowing heterogeneity only through variation in ownership structure. In section 1.5.1, I verify empirically that this ownership channel accounts for cross-sectional variation in the earnings passthrough.

1.2.3 Estimation

Equation (1.10) is suggestive of a method to recover the earnings elasticity. If $\tilde{\beta}_{it}$ differs from the structural earnings parameter by a magnitude equal to the earnings passthrough times the investor’s price elasticity, then one could estimate the $\tilde{\beta}_{it}$ from a reduced-form

pricing of trades on earnings surprises, and correct it for the magnitude of the bias. The challenge with this approach is that the earnings passthrough in the correction, γ_{nt} , is stock- and time- specific, so cannot be estimated in a cross-sectional or time-series regression.³ The alternative approach taken in this paper is to estimate a demand system that adjusts each investor’s trading for her price response before estimating her response to earnings. In words, this approach regresses the component of trades net of the price-induced response on the earnings surprises and controlling for changes in other characteristics that affect investor demand.

The first stage in the estimation procedure is to estimate investor-by-quarter level price elasticities, ζ_{it} . Identifying price elasticities in asset pricing setting is typically challenging due to the need for orthogonal drivers of changes in price. It is not sufficient to regress change in log quantities on prices, since prices themselves are endogenous to the trades (quantities) of investors. Nor is it possible to use asset-relevant news – such as earnings news – to identify price elasticities, since the investor’s response to this news depends on both his price and earnings elasticity (as the demand system framework, and equation (1.6), make clear).

Instead, I use the “investment universe” approach from Kojen and Yogo (2019) to estimate price elasticities. This approach involves a two-stage estimation procedure that generates cross-sectional variation in the market equity of stocks by exploiting variation in the investment universes of asset managers. Formally, it constructs an instrument for market equity as the counterfactual market equity that would obtain at the market clearing price if every asset manager held an equal-weighted portfolio of all stocks in their investment universe, where a manager’s investment universe is defined as the set of stocks held in the preceding or subsequent 11 quarters. This instrument is defined as:

3. For estimation of a restricted model in which earnings passthroughs are assumed to be constant across stocks, i.e. $\gamma_{nt} = \gamma_t$, see Appendix section 1.8.1.

$$\hat{p}_{int} = \log \left(\sum_{j \neq i} A_{jt} \frac{\mathbb{I}_{jnt}}{1 + \sum_{m=1}^N \mathbb{I}_{jmt}} \right) \quad (1.12)$$

where $\hat{p}_{int}$ is the log market equity of stock n at time t , implied by the investment universes of investors other than i . An advantage of using the price elasticity estimated from the investment universe of holdings is that it provides investor-by-quarter measures of elasticity, where alternative approaches typically only have sufficient statistical power for pooled estimation over the full sample. In addition to the standard exclusion restrictions and instrument relevance, the underlying assumption required to use this instrument in this demand system is that the price elasticities estimated from cross-sectional variation in price levels closely reflect the elasticity identified in the *time series* from a shift in price. In other words, it assumes that how an investor shifts his holdings of two stocks that differ in (exogenous) price levels at a point in time tracks how an investor would respond to an exogenous increase in the price level of a stock in a quarter.⁴ An additional assumption is that price elasticities on earnings announcement days are identical to those on non-earnings announcement days. It is worth noting that for the purposes of recovering β_{it} , any measure of price elasticity that satisfies these assumptions will work; one could alternatively use measures estimated from other methods (e.g. Granular Instrumental variables, flow induced trading, benchmarking intensity, etc.). However, the evidence in section 1.5 demonstrates that the investment-universe instruments are able to account for several important asset pricing patterns consistent with theory.

Following Koijen and Yogo (2019) the procedure to estimate price elasticities ζ_{it} involves investor-by-quarter GMM estimation on moment condition:

$$\mathbb{E}[\epsilon_{int} \mid \hat{p}_{nt}, X_{nt}] = 1 \quad (1.13)$$

4. One reason that these two elasticities could differ is if investors separately use price *levels* and price *changes* in forming expectations of future returns.

where $\hat{p}_{nt}$ is defined in (1.12), and ϵ_{int} is the residual of investor demand:

$$w_{int} = \exp \left\{ (1 - \zeta_{it}) \hat{p}_{nt} + \mathbf{\Gamma}'_{it} \mathbf{X}_{nt} \right\} \epsilon_{int} \quad (1.14)$$

In estimation, ζ_{it} is bounded to be greater than 0 (i.e. downward sloping demand). In exercises in which ζ_{it} is unbounded, I estimate a 2SLS “log” version of (1.13)-(1.14) in which the first stage projects log market equity on the implied market equity in (1.12), and the second stage regresses log portfolio weights, $\log w_{int}$, on instrumented log market equity and the characteristics in $\mathbf{X}_{nt}$.

Once price elasticities are estimated, (1.1) can be rearranged to obtain the second estimating equation:

$$\Delta q_{int} + \zeta_{it} \Delta p_{nt} = \alpha_{it} + \beta_{it} SUE_{nt} + u_{int} \quad (1.15)$$

The left-hand side reflects a compensated measure, Δq_{int}^* that adjusts each investor’s trades for the price response:

$$\Delta q_{int}^* := \Delta q_{it} + \zeta_{it} \Delta p_{nt} \quad (1.16)$$

Since Δq_{int} represents the general equilibrium (i.e., total response) of trades, it includes the response to both prices and other earnings news. The compensated measure Δq_{int}^* adjusts this quantity by subtracting the investor’s trading due to changes in market equity, $-\zeta_{it} \Delta p_{nt}$, where ζ_{it} is estimated from (1.13)-(1.14) and where Δp_{nt} is the actual – not instrumented – log return over the quarter. In the final estimating equation, I regress this “price-adjusted” response on the earnings news and control for innovations in the characteristics from (1.14) that determine the level of portfolio weights:

$$\underbrace{\Delta q_{it} + \zeta_{it} \Delta p_{nt}}_{\Delta q_{int}^*} = \alpha_{it} + \beta_{it} SUE_{nt} + \mathbf{\Psi}'_{it} \Delta \mathbf{X}_{nt} + \epsilon_{int} \quad (1.17)$$

In the linearized demand curve in (1.1), $SUE_{nt} \perp u_{int}$, where u_{int} is latent demand. An identifying assumption to estimate (1.17) is thus that $E[\varepsilon_{int} | SUE_{nt}, X_{nt}] = 0$, i.e. mean independence of latent demand given the earnings surprise and the characteristics. It is known that earnings announcements are associated with a large amount of information beyond the earnings beat. They may contain, for example, forward guidance, information on firm activities, or other numbers (such as the revenue beat) relevant to future expectations. The identifying assumption needed to estimate (1.17) is not that investors ignore this other information, but rather that latent demand is mean zero conditional on a given realization of the earning surprise. For example, if investors associate the earnings surprise with future growth in fundamentals, or subsequent investment in the firm, the identification is still valid: investors change their quantities only due to the earning surprise itself, and indeed the elasticity is designed to capture the broad and heterogeneous ways in which investors interpret fundamental news. What is important for identification is that there are not other factors correlated and concurrent with the release of earnings that affect investor demand. For example, suppose investors have a “taste” for stocks associated with artificial intelligence, and in a given quarter, stocks that beat earnings surprises are systematically more likely to mention AI. Then the measured effect on trades is contaminated by the demand for AI, and the estimated coefficient β_{it} is biased. I assume that earnings surprises are close to randomly assigned, and so the assumption of mean independence is likely to hold.

1.3 Data and Summary Statistics

Earnings News: I use the standardized unexpected earnings (“SUE”) based on I/B/E/S unadjusted reported analyst forecasts and actuals. The earnings surprise is constructed as

$$SUE_{nt} = \frac{(EPS^{Actual} - EPS^{Expected})}{\text{Price per Share}}$$

where $EPS^{Expected}$ is the median analyst expectation among forecasts made in the 90 days prior to earnings. The sample begins in 1984Q1, but coverage in the early period is sparse: there are fewer than 1000 stocks until 1986Q1, and fewer than 2000 until 1994Q1. In most analyses, I filter to quarters where there at least 500 stocks in the cross-section ($\geq 1985Q1$), or – for the asset pricing analysis in section 1.5 – post 1990, and I cross-sectionally winsorize the earnings variable at percentiles 3 and 97.⁵ For stocks with multiple earnings announcements in a quarter, which occurs when there are annual releases and quarterly releases, I aggregate to a single stock-by-quarter SUE_{nt} within the quarter by taking a simple average across the earnings announcements. Summary statistics on the sample are below:

Variable	P25	P50	P75	Mean	St. Dev
SUE_{nt}	-0.001	0.0003	0.002	-0.0005	0.010
$ SUE_{nt} $	0.001	0.002	0.005	0.004	0.009

Holdings: I use holdings data from 13-F disclosures available from Thomson Reuters. All asset managers with $>\$100M$ AUM are required to disclose quarterly positions of registered 13-F securities, which include common stocks. Common stock holdings by 13-F reporting institution account for 70% of total market capitalization at the end of the sample. Following Koijen and Yogo (2019), I classify 13-F managers into one of 6 types: Banks, Insurance Companies, Investment Advisers, Mutual Funds, Pension Funds, and Other. The residual sector – which includes non-13-F reporting institutions, and the aggregate short sector – are classified as the “household” sector. For sector- and aggregate-level measures (section 1.4), I construct Δq_{int} as the change in log quantities quarter to quarter. However, since

5. I focus on post-1990 because the precision ($E[|SUE_{nt}|]$) and bias ($E[SUE_{nt}]$) of SUE stabilizes after this period, as seen in appendix figure 1.12.

holdings are sparse at the manager-level, the change in log quantities is not well defined when a stock position moves to or from zero portfolio weight. Therefore, for the construction of stock-level sensitivities, which require estimates for each investor in a stock, I use $\Delta q_{int} = (Q_{int} - Q_{int-1}) / (2(Q_{int} + Q_{int-1}))$ in the estimation of the demand system, which approximates the percentage change in shares by normalizing by the average number of shares held across the two quarters. This measure closely tracks the log-constructed measure on the matched sample ($\rho = 0.87$).

Additional Stock Characteristics: Stock characteristics, such as returns and market equity, are from CRSP, while accounting data (including dividends, assets, profits, and book equity) are from Compustat. Most regressions use the set of (non-price) control variables from Kojien and Yogo (2019) that are designed to span the FF5 factors: profitability, dividend payout ratio, asset growth, beta, and log book equity.⁶ For the heterogeneous earnings passthrough analysis, I also add assets-to-book equity, and log book-to-market equity. Section 1.5 examines the alpha of portfolios constructed on earnings sensitivity; I use data on factor returns available from Ken French’s website. The matched sample of stocks that appear in 13-F reported securities and have a constructed measure of earnings surprises includes 276,680 stock-quarter observations from 1984Q1-2017Q4, including 10,652 distinct permnos.

1.4 Results

1.4.1 Market-Level Estimates

I estimate the earnings-augmented demand system from 1985Q1-2017Q4. While estimation at the investor-by-quarter level is, in principle, possible, trades are sparse at the manager

6. These 5 characteristics, along with (instrumented) market equity, explain a significant component of investor demand. Reduced-form IV regressions of log portfolio weights on the market equity instrument and the 5 characteristics deliver R^2 of $\geq 30\%$ for the majority of 13-F investors.

level. I therefore conduct two types of pooling: first, following Koijen and Yogo (2019), I pool managers with below 1000 holdings into a size-by-type group (a form a group shrinkage). Second, I estimate the set of regressions over rolling four-quarter windows. I construct standard errors by jointly bootstrapping a common set of stock-quarter observations in the regressions, as described in appendix section 1.8.2.


Figure (1.2) plots the asset-weighted earnings and price elasticities over this period.

Panel (a) shows the asset-weighted earnings elasticity. The market-level elasticity averages around 3: in other words, for a stock that beats consensus expectations by 1%, the average investor would demand a 3% increase in the number of shares he holds if shares of that stock were available at the pre-news price. As noted, this elasticity is a partial equilibrium parameter that asks how quantities demanded would change if prices were held fixed. It does not itself imply that investors must (or do) change their holdings in equilibrium. Indeed, it is perfectly possible for no-trade to occur in equilibrium even as investors have high demand for earnings. The reason is that prices adjust such that investors are indifferent to keeping their current holdings at the new price and realization of earnings news.

While earnings elasticities average near 3, there is significant low frequency variation over the time series. Toward the early part of the sample, pre 1990, earnings elasticities are small and close to zero (potentially a consequence of the smaller information content of consensus expectations in the early days of I/B/E/S forecasts). The highest earnings sensitivity observed over a sustained period is in the late 1990's and mid to late 2000's, when the average earnings elasticity was closer to 5.

Panel (b) shows the asset-weighted price elasticity. The market-level price elasticity is small and close to 0.26 on average. The interpretation is that a price increase of 1% translates into a 0.26% decrease in the quantity of shares demanded. Equivalently, a 1% flow into stocks leads to a $1/0.26 \approx 3.85$ times multiplier in returns. As noted in Gabaix and Koijen (2022), these are well below the estimates implied by standard asset pricing models. Two features

Figure 1.2: Aggregate Price and Earnings Elasticities



(a) Asset-Weighted Earnings Elasticity



(b) Asset-Weighted Price Elasticity

Panel (a) shows the asset-weighted earnings elasticity, computed by aggregating β_{it} (estimated from equation (1.17)) using AUM-weights. Standard errors are bootstrapped using the method in appendix section 1.8.2. Panel (b) shows the asset-weighted price elasticity, computed by aggregating ζ_{it} (estimated from equation (1.13)) using AUM-weights. Following Kojen and Yogo (2019), price elasticity estimates are bounded below by zero and are estimated by pooling estimation for investors with fewer than 1,000 holdings a quarter at an investor by size-quantile level. Estimation uses rolling four quarter windows.

of Figure (1.2) stand out. First, pre-2010, the aggregate market elasticity is fairly stable, ranging from 0.21 (a multiplier of 4.76) to 0.47 (a multiplier of 2.13). Price elasticities during this period are far more stable than earnings elasticities. Second, during 2014, there is a large spike in elasticities for both prices and earnings that reverts soon after. The rise in earnings elasticity during 2014 is a consequence of large increase in *price* elasticities measured during this year, which itself is driven by a spike in the household sector’s – and only the household sector’s – price elasticity in 2014. Because the household sector is constructed as a residual to 13-F reporting institutions, and because it relies on an investment universe approach in the cross-section of holdings, this jump more likely represents measurement error rather than a true change in price elasticities. Alternative methods to estimate price elasticities may generate smoother measures of price elasticities and, by extension, earnings elasticities, during this period.


1.4.2 Estimates by Sector

The manager-level estimates from the demand system allow for alternative forms of aggregation that can provide some insight on how different investing sectors react to earnings news. I asset-weight quarterly manager level sensitivities within 6 13-F sectors: Banks, Pension Funds, Investment Advisors, Mutual Funds, Insurance Companies, and the residual “household” sector, and plot the earnings elasticities:

Consistent with the implications of the CARA model, all sectors have nonnegative earnings elasticity over the full sample. There is a common increase across sectors beginning in the 1990s, which begins to decline around 2000. However, investors differ in their trends. for example, investment advisors have small responses until the mid-to-late 1990s, whereas the sensitivity of mutual funds to earnings news starts several years prior. The household sector sees very little time-series variation in its earnings sensitivity.

There are also important differences in the levels of sensitivities to prices and earnings.

Figure 1.3: Earnings Elasticities by Sector



This figure shows the aggregate earnings elasticities by sector, 1985Q1-2017Q4. Earnings elasticities are estimated investor-by-quarter using equation (1.17) over rolling four-quarter periods, and are aggregated using AUM weights. Standard errors are bootstrapped using the method detailed in Appendix section 1.8.2.

In Figure (1.4), I cross-sectionally standardize each sector's price and earnings elasticity, and take the time-series mean of these standardized earnings elasticities $\overline{\beta_{Sjt}}$ and price elasticities $\overline{\zeta_{Sjt}}$ within sector. Among institutional investors, there is a positive correlation between price elasticities and earnings elasticities. Of these 13-F investors, investment advisers and insurance companies are the more earnings sensitive, with an average elasticity half a standard deviation above average.⁷ Pension funds are the least earnings elastic and price inelastic, which would be expected of less attentive or active traders. Investment advisers, meanwhile, are earnings and price elastic, consistent with the "rational" responses implied by the CARA model.

Compared to most institutional investors, the residual household sector is price elastic

7. Because estimates of earnings elasticity requires estimates of price elasticities, they are can be sensitive to estimation choices for ζ_{it} (and Δq_{it}). Sector-level estimates in which ζ_{it} is estimated without a requirement that demand is downward sloping yield the largest earnings and price elasticities for investment advisers among institutional investors.


and earnings elastic. In line with previous literature (Kojen and Yogo (2019)), the household sector is among the most price elastic among sector types, with a price elasticity nearly 0.5 standard deviations above average.⁸ But unlike for institutional sectors, this price elasticity does not coincide with higher earnings elasticity: households have the among the lowest sensitivities to earnings news, with magnitudes only slightly above zero. A takeaway from Figure (1.4) is that price and earnings elasticities are not common across investors, as would be expected from a standard homogeneous investor agent model, but rather heterogeneous. Broadly speaking, investor types fall into one of four quadrants: price and earnings elastic (investment advisers, including hedge funds); price elastic and earnings inelastic (household sector); earnings elastic and price inelastic (banks, mutual funds, and insurance companies); and earnings and price inelastic (pension funds). To the extent stocks also differ in the relative proportion of their shares outstanding held by these different investor types, this investor heterogeneity will induce a stock-level heterogeneity in average elasticities, a property that will be explored in the subsequent section.

How important are price elasticity adjustments in estimating the demand for earnings? Recall that if price elasticities are zero, then there is no identification problem, in the sense that the reduced-form estimates perfectly recover the structural estimates. Intuitively, for an investor that has no response to price ($\zeta_{it} = 0$), the entire response to the earnings news can be attributed to the news itself rather than the contemporaneous price impact. For sectors with elasticities close to 0 – such as the insurance sector – this is indeed the case. Fig. (1.5) panel (a) compares the time series estimates of β_{it} to the reduced-form estimates $\tilde{\beta}_{it}$ one would obtain if they ran the following regression of trades on earnings surprises:

$$\Delta q_{int} = \alpha_{it} + \tilde{\beta}_{it} SUE_{nt} + \Psi'_{it} \Delta \mathbf{X}_{nt} + u_{int} \quad (1.18)$$

8. The price elasticity estimates for the household sector are still well below the price elasticities implied by standard asset pricing models.

Figure 1.4: Price vs. Earnings Elasticities by Sector



This figure shows relative price and earnings elasticities by 13-F sector type. Each quarter, I compute the asset-weighted price and earnings elasticity within sector, and then standardize these measures by quarter (i.e. across the investor types). The displayed results are the time-series means of these standardized measures, i.e. the axes represent the average number of standard deviations of a sector's elasticity relative to average. I omit 2014Q1-Q4 due to the large spike in household elasticity associated with the “investment universe” approach to measuring elasticities.


While $\tilde{\beta}_{it}$ is always below β_{it} (since price elasticities are positive), the two estimates are close in levels and trends.

However, for investors with price elasticities sufficiently far from zero, such as the household sector, the wedge between reduced-form and structural elasticities can be much larger. Fig. (1.5) panel (b) compares the reduced-form coefficient to the structural coefficient for the household sector, where the reduced-form coefficient $\tilde{\beta}_{it}$ is estimated using (1.18) on the household sector's trades. Not only is the wedge between the two coefficients much larger, but the sign differs as well, with $\tilde{\beta}_{it} < 0$ and $\beta_{it} > 0$ in most periods. That the reduced-form coefficient is negative would otherwise suggest that this sector has *negative* sensitivity to earnings news. Rather, the household has positive demand sensitivity to earnings news, but is sufficiently price elastic that it sells shares in response to the price increase induced by the rest of the market's comparatively larger earnings elasticities. As noted, this problem is not addressed by controlling for returns: the coefficient from a regression of the household sector's demand on earnings surprises remains negative even after controlling for contemporaneous log stock returns over the quarter. While markets appear to be much more inelastic in the data than in standard asset pricing theory, the importance of adjusting estimates is still large: on average, the asset-weighted reduced-form elasticity $\tilde{\beta}_{St}$ is 22% the magnitude of the asset-weighted structural elasticity β_{St} , and ranges from -25% of the magnitude (when reduced-form estimates are negative) to 59% the magnitude (in 2001). Equivalently, the difference between $\tilde{\beta}_{St}$ and β_{St} is on average 77% of the magnitude of β_{St} . A full time-series of the discrepancy can be found in appendix figure (1.13).

1.4.3 *Decomposition of the Earnings Passthrough*

An advantage of analyzing the pricing of earnings news using a demand system is that it provides a formula for the aggregate earnings passthrough in terms of investor-level price and earnings elasticities. Empirically, the passthrough of earnings surprises to stock returns

Figure 1.5: Reduced-form vs. Structural Earnings Elasticities



This figure compares reduced-form earnings elasticities, $\tilde{\beta}_{it}$, to structural elasticities, β_{it} . $\tilde{\beta}_{it}$ is estimated as in equation (1.18), and the adjustment to recover β_{it} uses bounded elasticities estimated from (1.13). In panel (a), manager level estimates within the Insurance sector are AUM-weighted to obtain a sector level measure. Panel (b) shows the analogous exercise for the household sector. Coefficients represent estimates over four-quarter rolling periods, with bootstrapped standard errors.

exhibits substantial time series variation; figure (1.6) plots the rolling coefficient from cross-sectional pricing regressions of quarterly stock returns on standardized unexpected earnings.

Figure 1.6: Time Series of Raw Earnings Passthrough



This figure shows the time series of coefficients (β) from a 3-quarter rolling regression $Ret_{nt} = \alpha_t + \beta_t SU E_{nt} + \epsilon_{nt}$, where Ret_{nt} is the total return over the quarter. Standard errors are clustered by quarter; dates correspond to the midpoint of the 3-quarter window.

According to the model, this time series variation could be driven by any combination of three factors. First, it could be due to changing *fundamental* elasticities, as investors have become more sensitive to earnings news. This could occur, for example, if earnings surprises became more or less informative about future growth prospects over time, or if consensus expectations by sell-side analysts better reflected the expectations of asset managers. A second account could be due to changing *price* elasticities, that is, if investors became more or less responsive to price changes. An example of institutional changes that could drive changing price elasticities is the growing share of passive investors, who tend to have price elasticities close to zero, or momentum investors, who have negative price elasticities. If more momentum investors enter the market, even with no change in earnings sensitivity, the passthrough from a given earnings surprise will be larger because these momentum investors

will augment the initial price impact of the earnings news. Finally, the change could be due to a *residual* term, capturing everything from estimation error to error in the underlying assumptions used to derive the model.⁹

Equation (1.6) implies that in the limit and under the model assumptions, the earnings passthrough, γ_{nt} is:

$$\gamma_{nt} = \frac{\beta_{Snt}}{\zeta_{Snt}} \quad (1.19)$$

where β_{Snt} and ζ_{Snt} are the within-stock share-weighted earnings and price elasticities respectively. Taking cross-sectional expectations of (1.19) and then applying log transformations, we can express the log average treatment effect as:

$$\begin{aligned} \log E_t [\gamma_{nt}] &= \log E_t [\beta_{Snt}/\zeta_{Snt}] \\ \log \hat{\gamma}_t &= \log \hat{\beta}_{St} - \log \hat{\zeta}_{St} + \hat{u}_t \end{aligned} \quad (1.20)$$

where u_t denote a residual arising from a Jensen term, model approximation error, and finite sample estimation.

Equation (1.20) not only provides an intuitive decomposition of the (log) earnings passthrough as a difference of (log) asset-weighted price and fundamental elasticities, but also allows for an additional test of the estimated asset demand system. Since the earnings passthrough (estimated from pricing regressions) should equal the ratio of share-weighted elasticities (estimated from holdings regressions), comparing the implied passthrough to the actual passthrough is an untested moment. I construct the demand-implied analog of the earnings passthrough by share-weighting investor-level β_{it} and ζ_{it} within-stock each quarter and

9. These assumptions include linear demand curve with investor-by-quarter specific coefficients; zero mean expectation of latent demand conditional on the earnings surprise and characteristics; similarity between the price elasticities implied by the investment-universe instrument and those that occur on earnings-news days.

taking the ratio of these objects.

Figure (1.7) shows the results from estimates. The solid line indicates the actual log earnings passthrough, estimated from cross-sectional pricing regressions of log returns on earnings surprises; the dotted line, by contrast, shows the demand-implied analog from averaging the ratio of share-weighted elasticities across stocks, and taking a log. The demand-implied passthrough approximates the actual passthrough in level in changes. In particular, it captures the low frequency rise in the 1990s, the high frequency variation around the turn of the century, and the drop and subsequent recovery around the recession. However, it overshoots the magnitude of the earnings passthrough throughout the 2000's.


In addition to showing the demand-implied passthrough, Figure (1.7)(b) also shows the relative contribution of the two weighted share-weighted elasticities. Log aggregate earnings elasticities, shown in red bar, sum with (negative) log aggregate price elasticities, shown in blue bars, and the residual (bronze bars) to yield the demand-implied coefficient. While the aggregate price elasticities are fairly constant over the time series, the aggregate earnings response moves in tandem with the pricing passthrough, suggesting that earnings elasticities play a comparatively larger role in explaining the time-series variation of γ_t . To quantify these relative contributions, I estimate a time series variance decomposition of the log earnings passthrough, by taking covariances of each component of equation (1.20) with $\log \gamma_t$:

$$1 = \underbrace{\frac{Cov(\log \hat{\beta}_{St}, \log \hat{\gamma}_t)}{Var(\log \hat{\gamma}_t)}}_{0.96} + \underbrace{\frac{Cov(-\log \hat{\zeta}_{St}, \log \hat{\gamma}_t)}{Var(\log \hat{\gamma}_t)}}_{-0.13} + \underbrace{\frac{Cov(\hat{u}_t, \log \hat{\gamma}_t)}{Var(\log \hat{\gamma}_t)}}_{0.17}$$


Close to 96% of variation in the earnings passthrough is driven by changing earnings elasticities; 17% is driven by the residual term; and a smaller -13% is driven by changing price elasticities. While the specific magnitudes are somewhat dependent on the estimation approach (e.g., the size of the rolling window, construction of Δq_{int}), the relative magnitudes

Figure 1.7: Earnings Passthrough Decomposition

(a) Actual vs. Demand-Implied Log Mean Earnings Passthrough



(b) Decomposition of Log Mean Earnings Passthrough



These figures compare the earnings passthrough estimated from pricing regressions to the earnings passthrough implied by taking the ratio of share-weighted elasticities estimated from the demand system. In panel (a), the solid black line plots the time series of $\log \gamma_t$ from quarterly regressions of $\Delta p_{nt} = \alpha_t + \gamma_t SUE_{nt} + \mathbf{\Gamma}'_t \Delta \mathbf{X}_{nt} \epsilon_{nt}$, where Δp_{nt} is the quarterly log return, SUE_{nt} is the standardized earnings surprise, and $\Delta \mathbf{X}_{nt}$ is the change in control characteristics. The dotted line plots $\log E_t [\beta_{Snt} / \zeta_{Snt}]$, where β_{Snt} and ζ_{Snt} are calculated by share-weighting, within stock, estimates of β_{it} and ζ_{it} from equations (1.11) and (1.13) respectively. Parameters are estimated over rolling four quarter windows. Panel (b) presents an analogous plot where $\log \gamma_t$ (the solid black line) is decomposed into the sum of $\log E_t [\beta_{Snt}]$, $-\log E_t [\zeta_{Snt}]$, and a residual u_t . The blue bars plot the earnings response $\log E_t [\beta_{Snt}]$, the red bars plot the price response $-\log E_t [\zeta_{Snt}]$, and the bronze bars plot the residual u_t .

consistently point to the primacy of changing earnings elasticities. These estimates are not surprising given the previously discussed stability of price elasticities over time, but they suggest that the vast majority of the variation earnings passthrough is well captured by the aggregate price and fundamental elasticities implied by the demand system.¹⁰

1.5 Asset Pricing Implications

So far, the analysis in this paper has considered heterogeneity in elasticities across investors and over time. But the implications for this variation in elasticities for asset pricing are not immediately clear: heterogeneity across investors may be unimportant to the extent that these differences “wash out” in the cross-section, and the documented time-series variation in elasticities does not reveal the economic reason for this evolution. In this section, I analyze the role that heterogeneous elasticities play in cross-sectional asset pricing. I first connect investor-level elasticity estimates to stock-level estimates, by exploiting variation in the ownership structure of stocks. The logic for this approach is that stocks inherit the sensitivities of their investors. If owners of an asset play a more salient role in pricing an asset at a point in time than investors who do not own the asset, then variation in the sensitivities of a stock’s owners will determine its response to earnings news. I first confirm that these stock-level measures of elasticity account for some variation in heterogeneous passthroughs. In section 1.5.1, I show that stocks with earnings-sensitive investors see a larger response to earnings news than stocks with less earnings-sensitive investors, while price-sensitive stocks see smaller passthroughs than price insensitive stocks. These effects are not spanned by many stock characteristics that explain expected returns and heterogeneous passthroughs.

In section 1.5.2, I then examine whether differences in earnings responses persist in the

10. The small contribution of price elasticities to time series *variation* in the log earnings passthrough does not imply that price elasticities are unimportant for matching the *levels* of the earnings passthrough. Because price elasticities feature in the denominator of the earnings passthrough, small changes in price elasticity are still able to generate large changes in the passthrough.

longer run. If ownership-induced differences in earnings passthroughs reflect frictions in external investors’ ability to respond to news at the same frequency as the owners of a stock, then passthroughs may be corrected as this friction is eased over time. However, ownership-differences could alternatively reflect unobserved characteristics that are informative about future earnings growth, in which case they can reflect rational pricing of the same news. Frictions may also be sufficiently strong such that differences in pricing are not corrected over a horizon observable to the econometrician. To test these channels, I form a set of two portfolios that each bet on long-run correction to ownership-induced responses to news. First, an “earnings sensitivity” combines a long-short portfolio that bets on reversal in earnings sensitive stocks (those expected to overreact to earnings news), with a long-short portfolio that bets on momentum in earnings insensitive stocks (those expected to underreact to earnings news). Second, a “price sensitivity” portfolio combines a long-short portfolio that bets on reversal in price inelastic stocks (those expected to overreact) with a long-short portfolio that bets on momentum in price elastic stocks (those expected to underreact). Each portfolio generates positive returns unconditionally and relative to standard factor models, consistent with the misreaction channel of ownership-related pricing of earnings news.

1.5.1 *Heterogeneous Earnings Passthrough*

A long literature in accounting endeavors to understand “earning response coefficients” (ERCs), i.e. why the market responds more strongly to the same earnings beat for some firms compared to others. In this literature, heterogeneous responses are often attributed to differences in characteristics of the underlying *firm*; for example, differences in market beta, capital structure, earnings quality, growth opportunities, price informativeness, and analyst dispersion (Scott and Scott (2015)). I consider an alternative account of heterogeneous earnings responses based on the sensitivities of a stock’s owners. Specifically, I compare the passthrough of earnings surprises for stocks with more earnings sensitive investors to those

with less earnings sensitive investors.

For each stock-quarter, I construct a measure of the weighted earnings and price elasticities of its owners. That is, I construct measures $\check{\beta}_{Snt}$ and $\check{\zeta}_{Snt}$ that give the share-weighted elasticities of the owners of stock n . First, at the investor level, I construct the average sensitivity of investor i over the preceding year $\bar{\beta}_{it} := \frac{1}{4} \sum_{k=0}^3 \beta_{it-k}$, $\bar{\zeta}_{it} := \frac{1}{4} \sum_{k=0}^3 \zeta_{it-k}$. To aggregate to the stock level, I take weighted averages of investor-level earnings sensitivities at the stock level:

$$\check{\beta}_{Snt} := \sum_i w_{int} \bar{\beta}_{it}, \quad w_{int} := \frac{Q_{int}}{\sum_i Q_{int}} \quad (1.21)$$

I construct an analogous measure of price elasticities,

$$\check{\zeta}_{Snt} := \sum_i w_{int} \bar{\zeta}_{it}, \quad w_{int} := \frac{Q_{int}}{\sum_i Q_{int}} \quad (1.22)$$

An important feature of these stock-level measures is that they use very little information on either a stock's past responses to earnings, or on the investor's response to that stock's history of beating earnings. Because investors hold many stocks, the measures of $\check{\beta}_{Snt}$ and $\check{\zeta}_{Snt}$ will be primarily driven by how investors respond to *other* stocks in their portfolio (in other words, omitting stock n from the computation of the investor-level sensitivities used to generate $\check{\beta}_{Snt}$ and $\check{\zeta}_{Snt}$ will have little effect on the investor-level or stock-level sensitivities).

I then estimate earnings response panel regressions of the form:

$$\Delta p_{nt} = \alpha_t + \alpha_n + \delta_1 SUE_{nt} + \delta_2 \check{\beta}_{Snt-1} + \delta_3 SUE_{nt} \times \check{\beta}_{Snt-1} + \mathbf{\Gamma}'(\mathbf{X}_{nt} + \mathbf{X}_{nt} \times SUE_{nt}) + \epsilon_{nt} \quad (1.23)$$

where Δp_{nt} is quarterly log returns. The coefficient of interest is δ_3 : if $\delta_3 > 0$, stocks with more earnings-sensitive investors see a higher passthrough from a given earnings surprise.

In all specifications, I include stock and date fixed effects and control for the same set of characteristics used in estimation: firm size (log book equity), profitability, asset growth, market beta, and dividends to book equity. I use the lagged measure of share-weighted earnings elasticities, $\check{\beta}_{Snt-1}$, such that the sorting variable is known ex ante and does not use any information from the contemporaneous earnings surprise. Finally, I cross-sectionally standardize $\check{\beta}_{Snt-1}$ so that δ_3 does not capture time series differences in the aggregate earnings passthrough.

Table (1.1) displays the results from estimating versions of (1.23). In column (1), log returns are regressed on earnings surprises without any interactions; as expected, stocks with higher earnings surprises have higher log returns (i.e., there exists a positive earnings passthrough). But crucially, this passthrough depends on the earnings sensitivity of its owners. Column (2) indicates that stocks with higher cross-sectional earnings elasticity see a higher earnings passthrough as well. Column (3) controls for the share weighted price elasticity, $\check{\zeta}_{Snt-1}$, as well as its interaction with the earnings surprise, $SUE_{nt} \times \check{\zeta}_{Snt-1}$. Consistent with the implications of the static model, in which higher price elasticity translates into a smaller earnings passthrough (since investors are more sensitive to price increases and dampen any price response from the earnings news), the coefficient on the interaction of earnings surprise with price elasticity is negative. In other words, the earnings passthrough is smaller for stocks that with more price elastic investors.

In column (4), I show that these effects are are not spanned by other characteristics that commonly explain heterogeneous earnings response coefficients. In addition to the controls used in estimation with the demand system (log book equity, profitability, asset growth, market beta, and dividends to book equity), I include measures of value (log book-to-market) and leverage (assets-to-book equity). I include these control variables as standalone characteristics and as interactions of each with the earnings surprise, for a total of 14 control

Table 1.1: Earnings Passthrough and Stock Level Sensitivities

	(1)	(2)	(3)	(4)	(5)	(6)
SUE_{nt}	3.814*** (5.774)	4.038*** (5.535)	4.156*** (5.540)	6.611*** (8.887)	4.180*** (5.238)	5.867*** (6.494)
$SUE_{nt} \times \check{\beta}_{Snt-1}$		0.493** (2.834)	0.472* (2.400)	0.326* (2.593)		
$SUE_{nt} \times \check{\zeta}_{Snt-1}$			-0.312*** (-3.785)	-0.420** (-3.162)		
$SUE_{nt} \times \check{\gamma}_{Snt-1}$					0.657* (2.611)	0.646** (3.052)
Num.Obs.	233 762	233 762	233 762	233 762	233 762	233 762
Controls	✓	✓	✓	✓	✓	✓
Date FE	✓	✓	✓	✓	✓	✓
Stock FE	✓	✓	✓	✓	✓	✓
Controls $\times SUE_{nt}$				✓		✓

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

This table shows estimates from regressions of the form $\Delta p_{nt} = \alpha_t + \alpha_n + \delta_1 SUE_{nt} + \delta_2 \check{\beta}_{Snt-1} + \delta_3 SUE_{nt} \times \check{\beta}_{Snt-1} + \mathbf{\Gamma}'(\mathbf{X}_{nt} + \mathbf{X}_{nt} \times SUE_{nt}) + \epsilon_{nt}$. All columns use date and stock fixed effects, include the earnings surprise SUE_{nt} , and control for log book equity, log book-to-market, asset growth, profitability, dividends to book equity, CAPM beta, and assets-to-book equity. Column (2) adds $\check{\beta}_{Snt-1}$ and the $\check{\beta}_{Snt-1} \times SUE_{nt}$, and column (3) additionally adds $\check{\zeta}_{Snt-1}$ and $\check{\zeta}_{Snt-1} \times SUE_{nt}$, where $\check{\beta}_{Snt-1}$ and $\check{\zeta}_{Snt-1}$ are constructed as in (1.21)-(1.22) and are then cross-sectionally standardized within quarter. Column (4) includes interactions of the seven control variables with SUE_{nt} . Column (5) includes $\check{\gamma}_{Snt-1}$ and $\check{\gamma}_{Snt-1} \times SUE_{nt}$, where $\check{\gamma}_{Snt-1}$ is the cross-sectionally standardized $\check{\beta}_{Snt-1}/\check{\zeta}_{Snt-1}$. t -statistics associated with standard errors clustered two-way at the date and stock level are in parentheses.


variables beyond the date and stock fixed effects. The coefficient on the interaction of the earnings elasticity with SUE_{nt} remains positive and similar in magnitude, while the interaction with price elasticity remains negative and significant. Finally, in columns (5) and (6) I include the “total” heterogeneous passthrough implied by the demand system, $\check{\gamma}_{Snt-1} := \check{\beta}_{Snt-1}/\check{\zeta}_{Snt-1}$. Consistent with the model, passthroughs are larger when the ratio of share-weighted earnings elasticities to share-weighted price elasticities is larger.

That interactions of elasticity measures with earnings news retain their sign and significance in the presence of alternative characteristics suggests that the ownership measures are not spanned by existing characteristics. However, they do not speak to the relative importance of these measures to explaining heterogeneous passthroughs. In figure (1.8), I compare the in-sample, within-adjusted R^2 of the earnings passthrough implied by different models. The left-most bars indicate the R^2 from a base model that only uses characteristic controls and fixed effects. Adding interactions of the seven stock-level characteristics with the earnings surprise slightly improves the adjusted R^2 , consistent with the notion that these characteristics account for some variation in the earnings response coefficient. Finally, I also add the interaction of the elasticity measures ($\check{\beta}_{Snt}$ and $\check{\zeta}_{Snt}$) with the earnings surprise, and chart the R^2 in the red bars. Here the improvement is large relative to only adding the stock characteristic variables. These results point to an additional channel driving heterogeneous earnings passthroughs. In addition to characteristics associated with the underlying firm and investor expectations, heterogeneous ownership structures drive variation in responses to earnings news.

1.5.2 Earnings Surprises and Misreaction

Cross-sectional variation in investor sensitivity to earnings news could arise from investors sorting into stocks that differ in the price-informativeness of their earnings news. However, it could also reflect a role for over or underreaction in the dynamic response of returns to earning

Figure 1.8: Improvements to R^2 from Elasticity Measures



This figure shows the within-adjusted R^2 from different models associated with equation (1.23). The “base” model shows the R^2 from estimating $\Delta p_{nt} = \alpha + SUE_{nt} + \mathbf{\Gamma}' \mathbf{X}_{nt} + \epsilon_{it}$ where α is α_t in the model with date fixed effects and is $\alpha_t + \alpha_n$ in the date and stock fixed effects model. $\mathbf{X}_{nt}$ includes log book equity, log book-to-market, asset growth, profitability, dividends to book equity, CAPM beta, and assets-to-book equity. The “+ Chars” model shows the R^2 from estimating $\Delta p_{nt} = \alpha + SUE_{nt} + \mathbf{\Gamma}' (\mathbf{X}_{nt} + \mathbf{X}_{nt} \times SUE_{nt}) + \epsilon_{nt}$. Finally, the “+ Chars & Elast” shows the R^2 from estimating $\Delta p_{nt} = \alpha + SUE_{nt} + \mathbf{\Gamma}' (\mathbf{X}_{nt} + \mathbf{X}_{nt} \times SUE_{nt}) + \mathbf{\lambda}' (\mathcal{E} + \mathcal{E} \times SUE_{nt}) + \epsilon_{nt}$, where $\mathcal{E} := (\check{\beta}_{Snt-1}, \check{\zeta}_{Snt-1})$.

news, particularly if sensitivities are orthogonal to future fundamentals. In this section, I address whether the heterogeneous passthroughs documented in the previous section are consistent with the market “correctly” pricing earnings news.

The central finding of this exercise is that variation in earnings sensitivity does not reflect rational investor sorting into stocks, but misreaction to earnings news. Consider a stock with earnings-insensitive investors, that realizes a large positive earnings surprise in quarter t . Because the investors of this stock do not respond much to earnings news, the return response is likely to be small. However, as time passes, sensitive external investors may internalize the news and bid up the price of the asset. Likewise, a stock with an extremely earnings sensitive investor base may overreact to news, generating a stock response that is later corrected. These effects imply a pattern of earnings momentum for stocks held by earnings-insensitive investors, and a pattern of earnings reversal for stocks held by earnings-sensitive investors.

To test this channel, I construct 25 cross-sectional portfolios as the intersection of 5 cross-sectional *SUE* quintiles and 5 cross-sectional earnings-sensitivity quintiles on the set of I/B/E/S covered stocks. For stocks in the highest quintile of earnings surprises, the candidate strategy goes long stocks in the lowest quintile of earnings sensitivity and short stocks in the highest quintile of earnings sensitivity. For stocks in the lowest quintile of earnings news, the strategy is reversed: long stocks in the highest quintile of earnings sensitivity, short stocks in the lowest quintile of earnings sensitivity. These portfolios are formed at the end of the quarter, and are held through the subsequent quarter. Formally, let $R_{nt+1}^{j,k}$ denote the return

of a stock in SUE quintile j and $\check{\beta}$ quintile k ,

$$\begin{aligned}
R_{t+1}^{GoodNews} &:= \overbrace{\sum_n w_{nt}^{5,1} R_{nt+1}^{5,1}}^{\text{Good News + Earnings Insensitive}} - \overbrace{\sum_n w_{nt}^{5,5} R_{nt+1}^{5,5}}^{\text{Goods News + Earnings Sensitive}} \\
R_{t+1}^{BadNews} &:= \overbrace{\sum_n w_{nt}^{1,5} R_{nt+1}^{1,5}}^{\text{Bad News + Earnings Sensitive}} - \overbrace{\sum_n w_{nt}^{1,1} R_{nt+1}^{1,1}}^{\text{Bad News + Earnings Insensitive}}
\end{aligned}$$

For the equal-weighted portfolio, the stocks are weighted equally within each leg of the portfolio, i.e. $w_{nt}^{j,k} = 1/\sum_n \mathbf{1}\{n \in j, k\}$. For the value-weighted portfolio, the stocks are weighted by lagged market equity in each leg of the portfolio, i.e. $w_n^{j,k} = me_{nt}^{j,k} / \sum_n me_{nt}^{j,k}$. The total earnings sensitivity portfolio return is then formed as the simple average of these two portfolios:

$$R_{t+1}^\beta = \frac{1}{2} \left(R_{t+1}^{GoodNews} + R_{t+1}^{BadNews} \right) \quad (1.24)$$

As detailed above, one way of thinking about this portfolio is a pair of “convergence” trades on two stocks that receive the same earnings news, but differ in the sensitivity of their owners. $R_{t+1}^{GoodNews}$ is a convergence trade conditional on good news, while $R_{t+1}^{BadNews}$ is a convergence trade conditional on bad news. An equivalent way of thinking about the returns on the total portfolio is as a pair of momentum and reversal strategies. That is, we can equivalently write the returns on the total portfolio as:

$$\begin{aligned}
R_{t+1}^\beta &= \frac{1}{2} \left(R_{t+1}^{Insensitive} + R_{t+1}^{Sensitive} \right), \\
R_{t+1}^{Insensitive} &= \sum_n w_{nt}^{5,1} R_{nt+1}^{5,1} - \sum_n w_{nt}^{1,1} R_{nt+1}^{1,1} \\
R_{t+1}^{Sensitive} &= \sum_n w_{nt}^{1,5} R_{nt+1}^{1,5} - \sum_n w_{nt}^{5,5} R_{nt+1}^{5,5}
\end{aligned}$$

The earnings insensitive portfolio, $R_{t+1}^{Insensitive}$, conditions on earnings insensitive investors and goes long stocks that have received good news and short stocks that have received bad news: this is similar to a post-earnings announcement drift portfolio, in that it bets on momentum in these stocks, but at a lower frequency. The earnings sensitive portfolio, $R_{t+1}^{Sensitive}$, conditions on an earnings sensitive investors, and bets long stocks that have received bad news and short stocks that have received good news, i.e. it bets on reversal in these stocks. Expressing the total portfolio this way highlights the fact that there should be no aggregate loading of this portfolio on the well-documented tendency of stocks to exhibit earnings momentum (Novy-Marx (2015), Chan, Jegadeesh, and Lakonishok (1996)). Because the portfolio takes an equal weight in momentum and reversal it does not simply load on this momentum factor.


Figure (1.9) illustrates the timing of the portfolio. The stock-level elasticities are estimated using ownership data prior to the release of news using data on $t - 4$ from $t - 1$. This not only ensures the strategy is feasible – since 13-F reporting requirements allow a delay between the end of the quarter and the period of disclosure – but also avoids using the end-of-period ownership shares, which are endogenous to responses to the earnings announcement itself. Between periods $t - 1$ and t the firm realizes its earnings news, and is sorted into SUE quintiles based on its relative surprise during this period. At the end of period t , the portfolios are formed using quintiles from the earnings news as of t and the elasticities as of $t - 1$. The portfolio is then held for a quarter from t to $t + 1$.

It is important to note that theories of over- and under-reaction do not directly speak to the timing over which correction must occur. When to form the portfolio and how long to hold it are ultimately empirical questions that depend on the degree of overreaction and correction.¹¹ In this exercise, the portfolio is formed at the *end* of the quarter in which the earnings announcement occurs, even though the earnings announcement can happen at

11. The momentum vs. reversal portfolio, and the good earnings news vs. bad earnings news portfolios may also differ from each other in their dynamics.

any point during the quarter. The portfolio is then held for a single quarter, even though misreaction can be corrected over a shorter horizon (including before the end of the quarter) or longer horizons (e.g. a year). While coarse, this approach has several advantages. First, since portfolios are formed once per quarter and are unique to the entire quarter, they reduce turnover relative to a portfolio that is resorted at the frequency at which earnings announcements occur (e.g. monthly). Second, since the elasticities are constructed using quarterly data on trades and holdings, they measure sensitivities at this frequency and suggest an analogous period for misreaction. Third, positive returns formed even on coarse measures – quarterly formation, and quarterly holding period – likely reflect conservative measures of the true returns achievable. More sophisticated strategies could analyze the timing over which over/underreaction peaks, the timing with which it corrects, and the differences in these properties over the four legs of the strategy.

Figure 1.9: Portfolio Timing



Before tracking the returns of the aggregate portfolio, I illustrate graphical evidence of misreaction patterns by tracking the *future* returns of stocks in different quantiles of earnings sensitivity after very positive (and negative) earnings news. In Fig. (1.10) I condition on stocks in the highest (panel (a)) and lowest (panel (b)) quintiles of earnings news respectively, and I track the cumulative returns in the *subsequent* quarter. For stocks in the highest quintile of earnings news, displayed in panel (a), the future returns are inversely related to earnings sensitivity: among stocks with good earnings news, those with the most earnings sensitive investors have the lowest future returns, consistent with contemporaneous overreaction that is later corrected. By contrast, high earnings news stocks with the

least earnings-sensitive investors have the highest returns. This pattern is consistent with contemporaneous *under*reaction by these investors, that is later corrected.


As noted, an alternative explanation for this pattern is that stocks held by earnings sensitive investors have properties associated with higher average returns. According to this narrative, stocks with the most sensitive investors might be expected to underperform stocks with high sensitive investors in any quarter. To test this, Fig. 1.10b considers the same strategy, but for stocks in the lowest quintile of earnings news; that is, those stocks that severely underperform analyst earning expectations. For these stocks, the pattern is reversed: low earnings news stocks with the most earnings-sensitive investors have high future returns, compared to those with the least earnings sensitive investors. This pattern is again consistent with stories of under and overreaction. High-earnings sensitive investors that see poor earnings news overreact to the news, causing a contemporaneous drop in stock price that is too large. The next quarter, this overreaction is corrected by other investors, leading to large future returns. For low-earnings sensitive investors, the logic is similar: these investors underreact to the low earnings news, causing a drop in the stock price that is too low. The subsequent quarter, the remaining investors correct this underreaction to poor news, and drive down the stock price further.

Figure (1.11) plots the cumulative raw returns of the aggregate earnings sensitivity portfolio for the 1990-2017 sample. To ensure that differences are not driven by correlations with stock-level price elasticities, I first orthogonalize $\check{\beta}_{Snt-1}$ on $\check{\zeta}_{Snt-1}$ in each quarter and form the earnings portfolio on the intersection of the earnings surprise and the orthogonalized earnings sensitivity. This earnings-sensitivity portfolio generates large positive returns, with a quarterly return of 1.57 ($t = 4.80$) for the equal weighted portfolio, and 1.28% ($t = 3.27$) over the 27 year period.


Since this portfolio bets equally on reversal on momentum within each SUE quintile – and since the total portfolio return is an average of the two SUE quintile portfolio returns

Figure 1.10: Future Returns by SUE and Beta Quintile

(a) Future Returns of **Highest** SUE Quintile Stocks




(b) Future Returns of **Lowest** SUE Quintile Stocks



This figure shows future returns of stocks sorted on earnings news and investor earnings sensitivity. At the end of each quarter t , stocks are sorted into the intersection of 5 quintiles of earnings news that quarter (on SUE_{nt}) and 5 quintiles of earnings sensitivity measured at the beginning of the quarter ($\check{\beta}_{Snt-1}$), where $\check{\beta}_{Snt-1}$ is first cross-sectionally orthogonalized on $\check{\zeta}_{Snt-1}$. Panel (a) conditions on stocks in the top quintile of earnings news (SUE quintile 5) and looks at the (equal-weighted) average returns the subsequent quarter ($t + 1$) for stocks in each quintile of $\check{\beta}_{Snt-1}$. The lines chart the cumulative returns from this strategy. Panel (b) conditions on stocks in the lowest quintile of earnings news (SUE quintile 1) and similarly plots the cumulative average return in $t + 1$ of stocks associated with each $\check{\beta}_{Snt-1}$ quintile. Consistent with overreaction, for good earnings news, future returns are inversely related with earnings sensitivity, while for bad earnings news, future returns are positively related to earnings sensitivity.

Figure 1.11: Earnings-Sensitivity Portfolio Returns



This figure shows the cumulative returns to the earnings sensitivity portfolio described in equation (1.24). The red line plots the cumulative quarterly equal-weighted returns, while the blue line plots the cumulative quarterly value-weighted returns. Since the elasticity measures use data from 4 quarters prior to the portfolio formation quarter, the sample period begins in 1990Q4 and extends to the 2017Q4.

Table 1.2: Earnings Sensitivity Portfolio Returns

	Equal-Weighted				Value-Weighted			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
α	0.016*** (4.792)	0.014*** (4.274)	0.015*** (4.464)	0.017*** (4.704)	0.013*** (3.272)	0.011*** (2.756)	0.011*** (2.748)	0.010** (2.390)
MKT_t		0.054 (1.305)	0.051 (1.149)	0.033 (0.706)		0.075 (1.522)	0.059 (1.096)	0.071 (1.268)
SMB_t			-0.040 (-0.530)	-0.057 (-0.751)			0.062 (0.695)	0.074 (0.808)
HML_t			-0.074 (-1.398)	-0.093* (-1.707)			-0.023 (-0.363)	-0.011 (-0.164)
UMD_t				-0.060 (-1.419)				0.040 (0.773)
Num.Obs.	108	108	108	108	108	108	108	108

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

This table shows the quarterly alpha associated with the earnings sensitivity portfolio described in equation (1.24). Columns (1)-(4) show the alpha of the equal-weighted portfolio while columns (5)-(8) show the alpha of the value-weighted portfolio. Robust t-statistics are shown in parentheses.

– the positive returns observed are not simply a case of earnings momentum. Nor are the returns explained by differential loadings on common factors. In table (1.2), I evaluate the alpha of this strategy against the market, Fama-French 3, and FF3/Carhart momentum model. The alpha remains positive and significant across specifications (with small loadings on existing factors). The smaller returns on the value-weighted portfolio (compared to the equal-weighted portfolio) are consistent with the notion that the market is more accurate at correctly pricing the earnings news of large stocks.

The model of investor demand does not only relate earnings passthroughs to earnings elasticities, but also relates them to price elasticities. In particular, when price elasticities are larger, earnings passthroughs are smaller. The implication for a trading strategy based on misreaction is that very price sensitive stocks see a muted contemporaneous reaction to earnings news that may be expected to grow in the following period. Price insensitive

stocks do not have sufficient active capital to trade against earnings-sensitive investors, which would lead to a too-large response that reverts in the following period. As before, I form four strategies that together combine to make a “price sensitivity” portfolio as:

$$\begin{aligned}
R_{t+1}^{\zeta} &= \frac{1}{2} \left(R_{t+1}^{GoodNews} + R_{t+1}^{BadNews} \right), \\
R_{t+1}^{GoodNews} &= \overbrace{\sum_n w_{nt}^{5,5} R_{nt+1}^{5,5}}^{\text{Good News + Price Sensitive}} - \overbrace{\sum_n w_{nt}^{5,1} R_{nt+1}^{5,1}}^{\text{Goods News + Price Insensitive}} \\
R_{t+1}^{BadNews} &= \overbrace{\sum_n w_{nt}^{1,1} R_{nt+1}^{1,1}}^{\text{Bad News + Price Insensitive}} - \overbrace{\sum_n w_{nt}^{1,5} R_{nt+1}^{1,5}}^{\text{Bad News + Price Sensitive}}
\end{aligned} \tag{1.25}$$

where $R_{nt+1}^{j,k}$ denotes the return of a stock in SUE quintile j and price elasticity quintile k . As before, a challenge with this approach is that earnings elasticities are correlated with price elasticities in the cross-section. Because the household sector has large price elasticities and small earnings elasticities, high household sector ownership of a stock will induce variation both in price elasticities and in earning elasticities that will work in opposite directions. To ensure that the price sensitivity portfolio is not indirectly picking up variation in earnings sensitivity, I first cross-sectionally residualize the stock-level measure of $\check{\zeta}_{Snt}$ on $\check{\beta}_{Snt}$, and use the residuals as the sorting variable for the price-elasticity quintile. Because these residuals are independent from $\check{\beta}_{Snt}$ (by construction) and from SUE_{nt} (by random assignment of earnings news) the sorts of this portfolio are independent from the sorts of earnings news portfolio in the limit.¹² Moreover, since the sorts for this portfolio only depend on relative price elasticities, not their magnitude, I use measures of price elasticity that are not bounded below by zero in estimation in order to induce maximum variation in price sensitivities.

12. An alternative approach is to conduct a triple sort by taking the intersection of portfolios sorted on SUE , β , and ζ . Unfortunately, there is not a sufficient number of stocks in the cross-section to deal with this additional dimensionality.

Table 1.3: Price Sensitivity Portfolio Returns

	Equal-Weighted				Value-Weighted			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
α	0.014*** (4.776)	0.012*** (4.094)	0.012*** (3.930)	0.012*** (3.979)	0.012** (2.617)	0.008* (1.817)	0.009* (1.871)	0.010* (1.982)
MKT_t		0.076** (2.130)	0.070* (1.804)	0.061 (1.506)		0.159*** (2.906)	0.165*** (2.765)	0.153** (2.455)
SMB_t			0.040 (0.617)	0.032 (0.484)			-0.053 (-0.529)	-0.064 (-0.628)
HML_t			0.016 (0.352)	0.007 (0.151)			-0.035 (-0.495)	-0.047 (-0.642)
UMD_t				-0.029 (-0.784)				-0.039 (-0.675)
Num.Obs.	108	108	108	108	108	108	108	108

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

This table shows the quarterly alpha associated with the price sensitivity portfolio described in equation (1.25). Columns (1)-(4) show the alpha of the equal-weighted portfolio while columns (5)-(8) show the alpha the value-weighted portfolio. Robust t -statistics are shown in parentheses.

Table (1.3) displays the cumulative returns from the “price sensitivity” portfolio 1990-2017. As with the earnings sensitivity portfolio, both the equal- and value-weighted portfolios display positive and significant alpha over the sample period. The four-factor equal-weighted alpha is 1.36% equal-weighted ($t = 4.78$) and 1.17% value-weighted ($t = 2.62$). The returns of both portfolios suggest that extremes of earnings and price sensitivities are associated with market mispricing in the cross-section that is resolved in the time-series.

1.6 Extending the Approach to Other Variables

As figure (1.1) makes clear, earnings news is among the most salient characteristics for explaining asset prices. However, there are many variables that affect investor demand (and hence prices), and which are likely to be heterogeneous in their effects on investor demand. In this section, I sketch out an approach for applying my methodology to a broader set of characteristics and price-relevant news. Because the demand system delivers heterogeneous

treatment effects through the ownership channel, the goal of this approach is not only to understand how investors differ in the types of news they attend to, but also to improve methods of return predictability by allowing for heterogeneous passthroughs at the stock level.

Consider a candidate set of characteristics $\mathbf{X}_{nt}$ that are relevant to investor demand. As a first stage, one could estimate these by regressing returns on changes in characteristics. Suppose characteristic X_{nt} is found to affect to average returns, i.e.

$$\Delta p_{nt} = \alpha_t + \beta \Delta X_{nt} + \epsilon_{nt}, \quad (1.26)$$

While β measures the average treatment of effect of characteristic ΔX_{nt} – i.e. it provides evidence that ΔX_{nt} matters for investor demand on average – it is plausible that X_{nt} is heterogeneous in its effect on investor demand. To estimate these heterogeneous sensitivities, one could replicate the approach in equations (1.16)-(1.17): first, estimate price elasticities ζ_{it} using a chosen price instrument (e.g. the investment universe approach). Second, construct the compensated trades Δq_{it}^* , by adjusting investor trades for their price component, $\Delta q_{int}^* = \Delta q_{int} + \zeta_{it} \Delta p_{nt}$. Third estimate the investor-level sensitivities, δ_{it} , to innovations in X_{nt} by running a regression of compensated trades on ΔX_{nt} :

$$\Delta q_{int}^* = \alpha_{it} + \delta_{it} \Delta X_{nt} + u_{int} \quad (1.27)$$

Fourth, asset-weight investor elasticities – both price and earnings – across investors to obtain market-level elasticities, δ_{Snt} and ζ_{Snt} . Finally, construct the demand-implied heterogeneous version of β in (1.26) as:

$$\beta_{nt} := \frac{\delta_{Snt}}{\zeta_{Snt}} \quad (1.28)$$

As noted, this approach delivers two quantities of interest. The first is the investor($\times$ quarter)

measure of elasticities, which speak to heterogeneity in investors' demand curves. For example, we may be interested in which investors attend to news about ESG, share repurchases, etc, and how this varies over time and by investor type (e.g. active vs passive). The δ_{it} identify this component of demand not driven by the contemporaneous price impact of the rest of the market. Second, this approach delivers heterogeneous coefficients in the return forecasting regression, and thus improve our understanding of return predictability. Just as allowing for heterogeneous passthroughs improved the R^2 in the earnings response coefficient regressions, one could estimate regressions of the form:

$$\Delta p_{nt} = \alpha_t + \theta_1 \Delta X_{nt} + \theta_2 \beta_{nt} + \theta_3 \beta_{nt} \times \Delta X_{nt} + \epsilon_{nt} \quad (1.29)$$

where β_{nt} is as defined in (1.28) and θ_3 captures the heterogeneous treatment effect of ΔX_{nt} implied by the demand system. Moreover, if extremes of sensitivity represent over- and under-reaction, they may be predictive of future returns as well, as the market corrects contemporaneous mispricing.

It is important to note that this approach is not solely limited to accounting characteristics or to the cross-section of equities. Investors attend to many news not found in traditional (e.g. Compustat) datasets – monetary policy shocks and other macroeconomic indicators, retail sentiment and news mentions, etc. – and the effect of these on the cross-section of stocks likely depends on the ownership. In addition, this approach can be extended to other asset classes, such as fixed income, where characteristics can including changes in ratings, volatility and inflation expectations, and tax policy.

1.7 Conclusion

Understanding how different investors respond to fundamentals is central to understanding the pricing of earnings news. This paper documents that the observed general equilibrium

responses of investors to earnings news reflect not only their sensitivity to the earnings news itself, but also how they react to the price impact induced by the market's response. When price elasticities also vary across investors, two investors may share similar general equilibrium responses but differ in their sensitivity to earnings.

In the first half of this paper, I develop a method for recovering earnings elasticities by adjusting each investor's total trading response for their response for prices, and then measuring the sensitivity of this compensated measure to earnings news. Intuitively, this method uses countefactual price elasticities to estimate how an investor would respond to the difference in pricing between the investor's perceived assessment of the earnings news and the market's. Using quarterly 13-F holdings data, I apply this method to estimate investor-by-quarter price and earnings elasticities over a 30-year period. These estimates reveal significant heterogeneity across investor types on both measures of elasticities, with the largest margin of variation between most institutional sectors (price inelastic, earnings elastic); investment advisers (price and earnings elastic); and the residual household sector (price elastic, earnings inelastic).

Because stocks vary in their ownership structure, *investor*-level heterogeneity in sensitivity to prices and fundamentals translates to *stock*-level variation in price and earnings sensitivity. In the second half of this paper, I show that these stock-level measures explain many properties associated with the pricing of earning news. More earnings sensitive and less price sensitive stocks see larger contemporaneous passthroughs from a given earnings surprise. This differential sensitivity appears to capture mispricing rather than unobserved characteristics that would justify a differential response for rational reasons. Ultimately, this cross-sectional mispricing is resolved in the time series: extremes of earnings sensitivity tend to revert the following quarter, such that strategies that bet on momentum in earnings insensitive (price sensitive) stocks and reversal in earnings sensitive (price insensitive) stocks generate profitable returns. Together, these results suggest that the ownership structure

of stocks, through heterogeneous sensitivities to prices and fundamentals, determines the pricing of earnings news in the market.

The importance of adjusting holdings responses for price elasticities is particularly important for price-relevant news (such as earnings news), but it applies more generally to any news that affects investor demand. The methodology in this paper could be extended to understanding investor demand for other sorts of news, such as news about macroeconomics, sentiment and uncertainty, ESG, regulatory and legal landscapes, and corporate actions. As with the analysis in this paper, this approach can be used not only to understand investor heterogeneity in demand for each type of news, but also to decompose the (average) passthrough of this news to aggregate prices and track the cross-sectional and time-series implications of stock-level heterogeneous responses.

Finally, while this paper focuses on the estimation and asset pricing implications of heterogeneous elasticities, it does not speak to the sources of this variation. Why earnings elasticities have changed over time (and what has been the impact of e.g., the rise in passive investing); how the estimated elasticities relate to the perceived persistence of future earnings growth and other behavioral models; and what information – beyond standardized earnings surprises – is contained in the earnings announcement is a ripe area for future research.

1.8 Appendix

1.8.1 Specialized Model with Homogeneous Earnings Passthrough

This section discusses estimation of a model with homogeneous earnings passthroughs. In effect, if elasticities are independent from relative ownership shares, then variation in ownership will not induce variation in the sensitivities of stocks. In this case, one can recover the structural earnings elasticity by adjusting the reduced-form coefficient by the average earnings passthrough identified from pricing regressions. Estimation of this model requires

the following proposition:

Proposition: Let s_{it} denote any investor-by-quarter component, and let η_{int} be the residual in $S_{int} = s_{it} + \eta_{int}$, where $S_{int} = Q_{int} / \sum_i Q_{int}$. If $\forall n \in N, \forall t \in T, \eta_{int} \perp (\beta_{it}, \zeta_{it})$, then $\hat{\gamma}_t$ in the cross-sectional regression $\Delta p_{nt} = \delta_t + \gamma_t SUE_{nt} + \epsilon_{nt}$ is consistent estimator of $\beta_{Snt} / \zeta_{Snt}$.

Proof. Imposing market clearing on (1.1),

$$\begin{aligned}
0 &= \sum_i (s_{it} + \varepsilon_{int}) (\alpha_{it} + \beta_{it} SUE_{nt} - \zeta_{it} \Delta p_{nt} + u_{itn}) \\
&= \alpha_{St} + \beta_{St} SUE_{nt} - \zeta_{St} \Delta p_{nt} + u_{Stn} \\
&\quad + \left(\sum_i \varepsilon_{int} \alpha_{it} + \sum_i \varepsilon_{int} \beta_{it} SUE_{nt} - \sum_i \varepsilon_{int} \zeta_{it} \Delta p_{nt} + \sum_i \varepsilon_{int} u_{itn} \right) \\
\Delta p_{nt} &= \frac{1}{\zeta_{St} + \sum_i \zeta_{it} \varepsilon_{int}} \times \\
&\quad \left[\alpha_{St} + \beta_{St} SUE_{nt} + u_{Stn} + \left(\sum_i \varepsilon_{int} \alpha_{it} + \sum_i \varepsilon_{int} \beta_{it} SUE_{nt} + \sum_i \varepsilon_{int} u_{itn} \right) \right]
\end{aligned}$$

Suppose $\forall n \in \mathcal{N}, t \in \mathcal{T}, \epsilon_{int} \perp (\zeta_{it}, \beta_{it})$. In words, this says that knowing an investor's elasticity is uninformative for whether they hold more or less of a given stock (in advance of earnings surprises). This condition implies that $\sum_i \zeta_{it} \varepsilon_{int} = 0$ and $\sum_i \beta_{it} \varepsilon_{int} = 0$. Then we have:

$$\begin{aligned}
\Delta p_{nt} &= \frac{\alpha_{St}}{\zeta_{St}} + \frac{\beta_{St}}{\zeta_{St}} SUE_{nt} + \frac{u_{Stn}}{\zeta_{St}} + \frac{1}{\zeta_{St}} \left(\sum_i \varepsilon_{int} \alpha_{it} + \sum_i \varepsilon_{int} \beta_{it} SUE_{nt} + \sum_i \varepsilon_{int} u_{itn} \right) \\
&= \frac{\alpha_{St}}{\zeta_{St}} + \frac{\beta_{St}}{\zeta_{St}} SUE_{nt} + \frac{u_{Stn}}{\zeta_{St}} + \frac{SUE_{nt}}{\zeta_{St}} \underbrace{\sum_i \varepsilon_{int} \beta_{it}}_{=0} + \underbrace{\left(\frac{1}{\zeta_{St}} \sum_i \varepsilon_{int} u_{itn} + \frac{1}{\zeta_{St}} \sum_i \varepsilon_{int} \alpha_{it} \right)}_{\check{u}_{nt}} \\
&= \frac{\alpha_{St}}{\zeta_{St}} + \frac{\beta_{St}}{\zeta_{St}} SUE_{nt} + \frac{u_{Stn}}{\zeta_{St}}
\end{aligned}$$

□

This proposition decomposes ownership shares into an investor-by-quarter component, s_{it} and a residual component, η_{int} . A natural candidate for the investor-by-quarter component s_{it} is the fractional ownership, or AUM-weights in the market: i.e. $s_{it} = A_{it} / \sum_i A_{it}$, where $A_{it} = \sum_n P_{nt} Q_{int}$ is the AUM held in 13-F securities. The proposition then asserts that if the deviations of ownership shares from these AUM-weights are uncorrelated with investor-level elasticities, the earnings passthrough coefficient is a consistent estimator of the ratio of asset-weighted elasticities.

The assumption allows estimation of the following specialized model. First, one can estimate a “quantity” regression using investor holdings data to obtain investor-by-quarter reduced-form elasticities, $\tilde{\beta}_{it}$:

$$\Delta q_{int} = \alpha_{it} + \tilde{\beta}_{it} SUE_{nt} + \boldsymbol{\Psi}'_{it} \Delta \mathbf{X}_{nt} + \epsilon_{int} \quad (1.30)$$

Second, one can estimate price elasticities ζ_{it} using (1.13). Finally, one can directly use the pricing passthrough coefficient γ_t , by running a series of cross-sectional pricing regressions of log returns on earnings surprises:

$$\Delta p_{nt} = \alpha_t + \gamma_t SUE_{nt} + \boldsymbol{\Upsilon}'_t \Delta \mathbf{X}_{nt} + \varepsilon_{nt} \quad (1.31)$$

Under these assumptions, one can then recover the structural earnings elasticity as:

$$\hat{\beta}_{it} = \hat{\tilde{\beta}}_{it} + \hat{\zeta}_{it} \hat{\gamma}_t$$

1.8.2 Bootstrap Details

I bootstrap standard error using the following procedure over 4-quarter rolling windows:

1. In iteration k , sample with replacement from the set of stock-quarters over $(t - 4) : t$. Denote this set $\Omega_t^{(k)}$.
2. Estimate the GMM regression, i.e. equation (1.13), on the full sample of manager i 's holdings in quarter t .
3. Construct Δq_{int}^* on the set of stock quarters in $\Omega_t^{(k)}$ as in (1.16).
4. For manager i , construct the k th draw of β_{it} as in (1.17):

$$\Delta q_{int}^{*(k)} = \alpha_{it}^{(k)} + \beta_{it}^{(k)} SUE_{nt} + \Psi_{it}'^{(k)} \Delta \mathbf{X}_{nt} + \varepsilon_{int}^{(k)}, \quad \forall n, t \in \Omega_t^{(k)}$$

5. Aggregate to 13-F sector level (j) as:

$$\beta_{Sjt}^{(k)} = \sum_{i \in \mathcal{J}(j)} \frac{AUM_{it}}{\sum_{i \in \mathcal{J}(j)} AUM_{it}} \beta_{it}^{(k)}$$

6. Repeat steps (1)-(5) K times to obtain the empirical distribution of β_{Sjt} .

1.8.3 Dynamic Model and Earnings Momentum

Consider a dynamic demand curve in which investor demand depends not only on current earnings news and prices ($SUE_{nt}, \Delta p_{nt}$) but also those that occurred at a lag: SUE_{nt-1} , Δp_{nt-1} . We can write investor demand as:

$$\Delta q_{int} = \alpha_{it} + \beta_{it}^{st} SUE_{nt} + \beta_{it}^{lt} SUE_{nt-1} - \zeta_{it}^{st} \Delta p_{nt} - \zeta_{it}^{lt} \Delta p_{nt-1} + u_{int} \quad (1.32)$$

where “ st ” and “ lt ” superscript denote “short-term” and “long-term” respectively.

Imposing market clearing, we solve for prices:

$$\Delta p_{nt} = \frac{1}{\zeta_{Snt}^{st}} \left(\alpha_{Snt} + \beta_{Snt}^{st} SUE_{nt} + \beta_{Snt}^{lt} SUE_{nt-1} - \zeta_{Snt}^{lt} \Delta p_{nt-1} + u_{Snt} \right) \quad (1.33)$$

Substituting back into investor demand,

$$\begin{aligned} \Delta q_{int} &= \underbrace{\left(\alpha_{it} - \frac{\zeta_{it}^{st}}{\zeta_{Snt}^{st}} \alpha_{Snt} \right)}_{\tilde{\alpha}_{int}} + \underbrace{\left(\beta_{it}^{st} - \zeta_{it}^{st} \frac{\beta_{Snt}^{st}}{\zeta_{Snt}^{st}} \right)}_{\tilde{\beta}_{int}^{st}} SUE_{nt} + \underbrace{\left(\beta_{it}^{lt} - \zeta_{it}^{st} \frac{\beta_{Snt}^{lt}}{\zeta_{Snt}^{st}} \right)}_{\tilde{\beta}_{int}^{lt}} SUE_{nt-1} \\ &\quad - \underbrace{\left(\zeta_{it}^{lt} + \frac{\zeta_{Snt}^{lt}}{\zeta_{Snt}^{st}} \right)}_{\tilde{\zeta}_{int}^{lt}} \Delta p_{nt-1} + \underbrace{\left(u_{int} - \zeta_{it}^{st} \frac{u_{Snt}}{\zeta_{Snt}^{st}} \right)}_{\tilde{u}_{int}} \\ &= \tilde{\alpha}_{it} + \tilde{\beta}_{int}^{st} SUE_{nt} + \tilde{\beta}_{int}^{lt} SUE_{nt-1} - \tilde{\zeta}_{int}^{lt} \Delta p_{nt-1} + \tilde{u}_{int} \end{aligned} \quad (1.34)$$

As before, reduced-form coefficients are structural coefficients adjusted by a ratio of elasticities:

$$\tilde{\beta}_{int}^{st} = \beta_{it}^{st} - \zeta_{it}^{st} \frac{\beta_{Snt}^{st}}{\zeta_{Snt}^{st}} \quad (1.35)$$

$$\tilde{\beta}_{int}^{lt} = \beta_{it}^{lt} - \zeta_{it}^{st} \frac{\beta_{Snt}^{lt}}{\zeta_{Snt}^{st}} \quad (1.36)$$

$$\tilde{\zeta}_{int}^{lt} = \zeta_{it}^{lt} + \frac{\zeta_{Snt}^{lt}}{\zeta_{Snt}^{st}} \quad (1.37)$$

Inspecting (1.32), the ratios of share-weighted elasticities, in turn, correspond to stock-specific passthroughs. $\beta_{Snt}^{st}/\zeta_{Snt}^{st}$ is the contemporaneous earnings passthrough; $\beta_{Snt}^{lt}/\zeta_{Snt}^{st}$ is the passthrough of lagged earnings (controlling for lagged returns) to current prices, i.e. “fundamental momentum”; and $-\zeta_{Snt}^{lt}/\zeta_{Snt}^{st}$ is the passthrough of lagged return to current

returns, i.e. “price momentum”. Fundamental momentum – the momentum associated with higher earnings in the past, controlling for past prices – is higher when investors are more sensitive to earnings last quarter or less sensitive to prices today. Price momentum depends on the term structure of price elasticities; if investors are less *long-term* price elastic – e.g., if they extrapolate past returns in which case $\zeta_{Snt}^{lt} < 0$, then the passthrough is larger. If, at the same time the market is more short-term price elastic, it can better absorb demand shocks contemporaneously, dampening any momentum effects.

Unlike in the static model, the reduced-form regression in (1.34) is not solely a function of earnings surprises. Nor can elasticities be recovered simply by estimating (1.34) with earnings surprises and *realized* past returns. Instead, it requires instruments for price changes Δp_{nt-1} as well. This is because market clearing was imposed only at time t , not $t - 1$, so $Cov(\Delta p_{nt-1}, \tilde{u}_{int}) \neq 0$ (e.g. if there are correlated demand shocks over time). One way to solve this is using instruments for longer term price changes, as in Huebner (2023), who uses flow-induced trading. Under this approach, one could measure long and short term price elasticities using chosen price instruments; adjust the trades this quarter for the term structure of elasticities, and then regress these compensated trades on current and past earnings surprises.

Alternatively one can impose market clearing at $t - 1$ as well and solve for prices Δp_{nt-1} , i.e. solve for the endogenous variable. Since Δp_{nt-1} is a function of Δp_{nt-2} , this introduces the need for an instrument for $t - 2$ prices as well:

$$\Delta p_{nt-1} = \frac{1}{\zeta_{Snt-1}^{st}} \left(\alpha_{Snt-1} + \beta_{Snt-1}^{st} SUE_{nt-1} + \beta_{Snt-1}^{lt} SUE_{nt-2} - \zeta_{Snt-1}^{lt} \Delta p_{nt-2} + u_{Snt-1} \right)$$

$$\Delta q_{int} = \tilde{\alpha}_{it}^{(1)} + \tilde{\beta}_{int}^{st(0)} SUE_{nt} + \tilde{\beta}_{int}^{lt(1)} SUE_{nt-1} + \tilde{\beta}_{int}^{lt(2)} SUE_{nt-2} - \tilde{\zeta}_{it}^{lt(1)} \Delta p_{nt-2} + \tilde{u}_{int}^{(1)}$$

One can iterate recursively to solve for current prices as:

$$\Delta p_t = \tilde{\alpha} + \frac{\beta_S^{st}}{\zeta_S^{st}} SUE_{nt} + \sum_{k=1}^{\infty} \tilde{\beta}^{(k)} SUE_{nt-k} + \tilde{u}_{nt}$$

where

$$\begin{aligned}\tilde{\beta}^{(k)} &= \frac{1}{\zeta_S^{st}} (-1)^k \left[\left(\frac{\zeta_S^{lt}}{\zeta_S^{st}} \right)^k \beta_S^{st} - \left(\frac{\zeta_S^{lt}}{\zeta_S^{st}} \right)^{k-1} \beta_S^{lt} \right] \\ \tilde{\alpha} &= \frac{1}{\zeta_S^{st}} \sum_{k=0}^{\infty} (-1)^k \left(\frac{\zeta_S^{lt}}{\zeta_S^{st}} \right)^k \alpha_{St-k} \\ \tilde{u}_{nt} &= \frac{1}{\zeta_S^{st}} \sum_{k=0}^{\infty} (-1)^k \left(\frac{\zeta_S^{lt}}{\zeta_S^{st}} \right)^k u_{St-k}\end{aligned}$$

In the limit, no price instruments are needed; however, estimation may suffer from low power.

1.8.4 Additional Figures

Figure 1.12: Bias and Precision of Standardized Unexpected Earnings





Figure 1.13: Discrepancy Between $\tilde{\beta}_{St}$ and β_{St}



This figure shows the time series of the magnitude of the difference between the asset-weighted reduced-form earnings sensitivity, $\tilde{\beta}_{St}$, and the asset-weighted structural elasticity, β_{St} . $\tilde{\beta}_{St}$ is computed by estimating (1.18) for each investor $\times$ quarter, and aggregating using AUM weights. β_{St} is computed by estimating β_{it} according to (1.17) and aggregating using AUM weights. The y -axis plots the difference between these quantities as a percentage of the magnitude of the structural earnings elasticity, i.e. it plots $(\tilde{\beta}_{St} - \beta_{St}) / \beta_{St}$. On average across the time series, the difference is 77% of β_{St} .

CHAPTER 2

INFLATION EXPECTATIONS AND STOCK RETURNS

2.1 Introduction

Do stock prices rise or fall when inflation expectations increase and if so, why?¹ A novice investor may notice that inflation expectations enter the nominal discount rate and conclude that stock prices must fall when expectations rise. A sophisticated investor may push back that stocks are ultimately a claim on income generated by real assets. As a result, dividends should rise with inflation, canceling out the effect of the discount rate. Finally, an even more sophisticated investor may argue that it is not so straightforward, as risk-premia and real interest rates could also change when inflation expectations move. For example, stock prices may fall if the central bank tightens monetary policy in response to rising expectations. Understanding the response of stock prices to changes in inflation expectations calls for careful empirical study.

Consistent with previous work (Campbell, Shiller, and Viceira (2009), Viceira (2013), Gourio and Ngo (2020), and Duffee (2023)) documenting a positive correlation between aggregate stock returns and inflation expectations, we first show that stock prices unambiguously rise when inflation expectations increase, and that this relationship is a highly stable and invariant characteristic of the asset class. We regress daily stock returns on changes in the 10-year breakeven inflation rate measured using Treasury Inflation-Protected Securities (TIPS), and find that a 1 basis point increase in the 10-year breakeven inflation rate is associated with an 11 basis point increase in the overall stock market value. Given that the standard deviation of daily changes in the 10-year breakeven inflation rate is around 3.6 basis points, the magnitude is non-trivial. This positive relationship is stable across time,

1. This chapter reflects joint work with Manav Chaudhary.

sub-samples, and the cross-section of stocks, and the results are virtually unchanged when we use inflation swaps, rather than TIPS, to measure expectations. The positive relationship remains intact even in the post-pandemic high inflation episode, when stock returns and realized inflation appeared to move in opposite directions. While the positive comovement establishes that stocks are a hedge against changes in inflation expectations, it does not explain why stocks provide this hedge. The remainder of the paper focuses on this question.

We zoom in on CPI release dates to better understand the causal effect of investors updating their inflation expectations on stock prices. CPI release dates help us identify the causal effect for two reasons. First, we can attribute changes in stock returns and expectations on CPI days to the inflation information released in the print (high-frequency identification assumption). Second, since the CPI print contains only direct information about inflation, it seems reasonable that any effect of the information shock on beliefs about other variables works through the agents' revision of inflation expectations (exclusion restriction). In other words, we are assuming that agents use the information contained in the inflation print to revise their beliefs about future inflation, which they may then use to update their beliefs about other relevant economic variables such as economic growth and the path of monetary policy. Assuming the high-frequency identification assumption and the exclusion restriction are satisfied, regressing returns on the unanticipated changes in expectations reveals the causal effect. The CPI-day regression estimates mirror closely the results from the unconditional daily regression. We find that a 1 basis point increase in the 10-year breakeven inflation rate causes a 12 basis point increase in the overall stock market value. This relationship is also positive and stable across the cross-section of stocks, over sub-samples, and is essentially unchanged for alternative expectation measures. While these estimates are silent about the precise conclusions investors may draw about the macroeconomy after updating their future inflation beliefs, the sharp parallel between the unconditional relationship

and the causal estimates suggests that the hedging property is a structural feature of equities.

Next, we decompose the channels driving the positive causal effect of the high inflation information shock. Asset pricing identities imply that positive returns must be driven by some combination of higher real cashflows, lower risk-free rates, or lower expected excess returns. We employ four different approaches that together point to the primacy of the expected excess return channel. We first sort stocks into portfolios by their market beta, and find that stocks with larger market beta have monotonically higher loadings on inflation expectations. Under a factor model where the market return is the only factor, we derive a simple estimator for how portfolios should respond to changes in inflation expectations. We find that the two-stage estimates from this approach mirror the direct cross-sectional estimates; in other words, our results are consistent with a model where returns are solely driven by innovations to the market factor implied SDF. The next two approaches attempt to directly quantify the contribution of cashflow, risk-free rate and excess return expectation channels. Following Bernanke and Kuttner (2005), in the first approach we estimate a six-variable, six-lag VAR to decompose the return response into the three channels. Nearly the entire effect of inflation expectations on stock returns is attributed to expected excess returns, with the contribution from future real rate and future dividends statistically indistinguishable from zero. Our third approach takes advantage of the high-frequency setting to directly control the mediating channels. Controlling for contemporaneous changes in treasury yields and market-based cash-flow expectation measures (from Gormsen and Koijen (2020)) does little to change the response to inflation expectations. Of the 12 basis point rise in aggregate market value in response to a 1 basis point rise in expectations, virtually all of it is from changes in expected excess returns with the contributions from changes in cash flow expectations and from changes in future risk-free rate expectations statistically indistinguishable from zero. Our final approach looks at the cross-section of returns. We

estimate the cross-sectional responses over a battery of accounting variables that we would expect to be associated with the real value of cashflows. We find virtually no patterns based on a firm’s leverage, profitability, or relative share of real and nominal assets. Moreover, the positive relationship is significant and consistent across all industries, indicating that it is likely driven by broader factors like risk premia, rather than specific firm cash flows. While all four methods have shortcomings, they jointly provide heuristic evidence that the effect works primarily through changes in expected excess returns.

Reduced-form results—even causal ones—do not speak to underlying economic mechanisms. For example, our results do not explain *why* rising inflation expectations cause expected excess returns to decline, and if this decline in expected returns reflects changes to risk premia or simply non-fundamental price impacts. We argue that the positive relationship reflects fundamental changes in risk, as information about higher future inflation signals higher future output. In other words, when investors develop higher inflation expectations, they assign lower probabilities to bad states of the world.² This improved outlook would reduce risk premia and thus raise stock prices David and Veronesi (2013).³

To assess the plausibility of the economic mechanism, we test if positive information about future inflation does in fact predict higher future output. The empirical challenge in uncovering this relationship is that the dynamics of aggregate output take several months (or years) to realize after the information shock, and their behavior depends on the evolution of other macroeconomic and financial variables. To overcome this challenge, we estimate a monthly frequency VAR containing output and other relevant endogenous variables such as

2. The effect on equities is largest for negative inflation expectations shocks in low-inflation economies, suggesting that expected disinflation may raise the probability of bad states of the world more than expected inflation raises the probability of good states.

3. This interpretation is consistent with higher inflation being associated with higher output since the 2000s Campbell, Pflueger, and Viceira (2020).

the price level, inflation expectations, and the central bank’s policy rate. We then include our high frequency shock measure as an exogenous variable (justified by the same identification assumptions we used in the CPI-day stock return regressions). We find that the behavior of the variables is broadly consistent with the information shock revealing higher future inflation: (i) the realized price level rises in response and stays elevated, and (ii) the central bank raises rates in the short-run and brings them back down in the long-run. Most importantly, consistent with our economic mechanism, positive information about future inflation predicts higher output and lower risk. An unexpected 10 basis point rise in the 10-year breakeven expected inflation predicts a 0.7% increase in output over the following year. The dynamics of VIX, and future realized variance of stock returns, show that higher expected inflation is associated with reduced volatility as well. Overall, this finding is consistent with investors using information about higher future inflation to revise their probabilities of future good states of the world.

Our paper contributes most directly to the literature studying the relationship between stock returns and inflation expectations Fama and Schwert (1977), Fama (1981), and Chen, Roll, and Ross (1986). This older literature found a negative correlation between stock returns and expected inflation, i.e., stocks did not provide a hedge against changes in inflation expectations. We build on these studies in several ways. First, since there have been many economic changes over the past few decades, most notably in the joint distribution of inflation and output (see e.g., Campbell, Pflueger, and Viceira (2020)), the sample periods of the past studies are no longer representative of the current economic regime. Our paper is among the first to document the relationship between inflation expectations and stock returns for the post-2000s period. Second, past studies typically measured expectations using either econometric models or survey-based measures. Model-based approaches make strong assumptions about the expectations process of investors, while survey-based measures are

often infrequent and may not represent the expectations of the aggregate investor class. We overcome these limitations by using market-based inflation expectations measures, which provide a better quality and higher frequency proxy for changes in investor beliefs. As a result, we achieve significantly higher statistical power than past methods. Finally, the high-frequency (daily) nature of these contracts reduces the risk of confounders leading to misspecification by allowing us to focus on CPI release days.

Several asset pricing models disagree about the sign and magnitude of the relationship between inflation expectations and equity prices David and Veronesi (2013), Campbell, Sunderam, and Viceira (2017), Campbell, Pflueger, and Viceira (2020), and Bhamra et al. (2021). Our results provide a yardstick for assessing these predictions. Furthermore, while our decomposition result—that the causal effect predominantly functions through the expected excess return channel—does not reveal the underlying economic mechanism, it does provide a guiding principle for the channels through which this underlying mechanism must function. For example, given that the effect is positive and the risk-free rate channel plays a limited role, it is unlikely that an explanation based on monetary policy reaction can reconcile the relationship between equity returns and changes in inflation expectations.

Our work is also linked to the broader empirical literature studying the relationship between stock returns and realized inflation (Lintner (1975), Bodie (1976), Jaffe and Mandelker (1976), Nelson (1976), Fama and Schwert (1977), Fama (1981), Chen, Roll, and Ross (1986), Bekaert, Wang, and Tille (2010), Ang, Briere, and Signori (2012), Boons et al. (2020), and Roussanov, Liu, and Fang (2021)). Theoretically, the predicted effect of realized inflation and stock returns, while not clear cut, is less ambiguous than the effect of expected inflation. Unlike inflation expectations, realized inflation does not appear in the nominal discount rate. Hence, *prima facie*, we would expect realized inflation to affect (nominal) returns positively

to the extent that dividends are claims on real goods, and so rise with inflation. Contrary to this prediction, the literature has overwhelmingly documented a negative exposure of stock returns to realized inflation.⁴ That said, the negative relationship between stock returns and realized inflation has turned positive since the 2000s—David and Veronesi (2013), Campbell, Pflueger, and Viceira (2020), and Boons et al. (2020) models suggest that the sign flip is consistent with higher inflation now proxying for higher growth. Roussanov, Liu, and Fang (2021) shed further light on the positive sign flip, showing that it is predominantly due to the correlation of returns and energy inflation changing sign, and not the correlation of returns and core inflation, which continues to be negative.⁵ A final advantage of using inflation expectations is that, by virtue of being expectations, they strip out much of idiosyncratic noise that appears ex-post in realized inflation series, allowing greater statistical power.

Our paper also relates to the recent literature attempting to determine if investors are aware of and are trying to insure against specific risk factors. Our revealed preferences approach complements the Chinco, Hartzmark, and Sussman (2021) survey-based method. By looking at price changes in response to high-frequency shocks to specific risk factors, we can show whether investors are trying to insure against those risks. Our paper also contributes to a growing literature emphasizing the importance of using subjective expectations, rather than objective realizations, in models of asset pricing (De La O and Myers (2021), Nagel and Xu (2022)).

Finally, our paper provides a new method for decomposing the response of prices to

4. Fama (1981) suggested the negative relationship could be due to high inflation proxying for low growth, whereas Modigliani and Cohn (1979) suggested that the negative correlation could be a result of investors incorrectly using nominal discount rates for valuing real cash flows (money illusion).

5. Related to this literature that directly looks at the relationship between returns and inflation, Weber (2015) shows that inflation risk affects equity returns via a sticky-price channel, namely, firms that adjust their product prices infrequently carry a return premium.

high-frequency shocks into contributions from changes in expected excess returns, changes in cash flow expectations, and changes in the risk-free rate using market-based expectation measures. Our method is complementary to the existing VAR methodology proposed by Bernanke and Kuttner (2005). Our method directly controls for mediating factors to determine their contribution; this does not suffer from the bad control problem due to the stronger identification assumptions the high-frequency setting allows. Similar methods have been used in epidemiology and psychology but not in financial economics as far as we know MacKinnon et al. (2002).

The remainder of the paper is structured as follows. Section 2.2 discusses the data used in our analysis. Section 2.3 provides the top-line result of the relationship between stock returns and changes in inflation expectations. Section 2.4 zooms in on scheduled CPI releases to estimate the causal effect of changes in inflation expectations on stock returns. Section 2.5 assesses the contributions of the various channels driving the stock price response. Section 2.6 provides a economic interpretation of our results, relating inflation expectations to output growth. Section 2.7 concludes the paper.

2.2 Data

In this section, we describe data sources, the construction of our measures, and present summary statistics.

Inflation expectations: Our primary measure of inflation expectations uses yields on treasuries and treasury-implied protected securities (TIPS) to back out the breakeven inflation rate. These securities adjust principal and coupon payments on government-issued treasuries by the realized change in the consumer price index over the corresponding period. Investors holding TIPS thus earn the real (inflation-adjusted) yield on treasuries, and the difference between these yields and the yields on nominal US treasuries of the same maturity

pins down risk-neutral expectations of inflation over the time period. We obtain daily data on TIPS-implied breakeven inflation expectation from the Federal Reserve Board’s TIPS file (for more detail, see Gürkaynak, Sack, and Wright (2010)), and compute daily changes in the daily inflation expectation as our primary variable of interest. Our series comprises 5,276 daily observations on each annual maturity of 2 to 20 years from January 5, 1999 to January 1, 2022.

Since the market for TIPS sometimes suffers from limited liquidity, we also use inflation expectations derived from swaps as a robustness check for all our results. We obtain inflation swaps at 1, 5, and 10 year maturities from Bloomberg. Collectively these have a shorter history than TIPS, and span 4,481 daily observations from July 21, 2004 to November 21, 2021. Our two inflation expectations measures are highly correlated in both levels ($corr(\mathbb{E}_t[\pi_t^{TIPS,10Y}], \mathbb{E}_t[\pi_t^{Swap,10Y}]) = 0.96$) and changes ($corr(\Delta\mathbb{E}_t[\pi_t^{TIPS,10Y}], \Delta\mathbb{E}_t[\pi_t^{Swap,10Y}]) = 0.74$).

Returns: Daily stock returns on the value-weighted US stock index are collected from the Center for Research in Security and Prices (CRSP). We augment this dataset with several other sources of equity returns. For the November 2021-2022 period analyzed at the end of Section 2.3, we use daily returns on the SP500 Total Return Index, accessed from Yahoo Finance. For portfolio constructions, we use firm characteristics from Compustat, the set of anomaly portfolios from the Open Source Asset Pricing Database Chen and Zimmermann (2022), and Fama-French 30 industry portfolios and monthly risk-free rates from Ken French’s website.

Dates: We separately analyze the inflation expectation - stock return relationship on sample dates corresponding to inflation news releases. We collect dates of Consumer Price

Index (CPI) releases from the Bureau of Labor Statistics.

Other Measures: We obtain real yields from GSW TIPS file (Gürkaynak, Sack, and Wright (2010)), and cash flow growth expectations from changes in dividend strips (Gormsen and Koijen (2020)).

Table 2.1 presents summary statistics on the main inflation expectations measures, the value-weighted stock return, and the real 10-year yield.

Table 2.1: Summary statistics of main measures

	Mean (basis points)	Standard deviation (basis points)	Observations
$\Delta \mathbb{E}_t [\pi_t^{TIPS,10Y}]$	-0.0	3.6	5028
ret_t	3.6	123.5	5536
$\Delta \mathbb{E}_t [\pi_t^{Swap,10Y}]$	-0.0	3.8	4295
$r_t^{TIPS,10Y}$	-0.1	4.7	5028

This table shows the summary statistics for measure of daily changes in inflation expectations and returns. Changes in both TIPS- and Swap-implied measured of inflation expectations are mean 0 over the sample. The sample period is 2000-2021 for TIPS measures, and 2004-2021 for Swap-implied measures.

2.3 Stocks prices rise hand-in-hand with inflation expectations

In this section we verify and show the robustness of the finding that stock prices rise when inflation expectations rise. To capture the overall effect of changes in inflation expectation on returns, we regress daily CRSP value-weighted returns on daily changes in the TIPS implied

10-year breakeven inflation rate:

$$ret_t = \alpha + \beta_\pi \Delta \mathbb{E}_t \left[\pi_t^{10Y} \right] + \varepsilon_t \quad (2.1)$$

We choose TIPS over swaps for our baseline measure due to its longer history. Furthermore, we use the 10-year tenor because it aggregates inflation expectations across the term structure. The longer horizon also suggests a cleaner measure of inflation expectations: as documented by Bahaj et al. (2023), so-called fundamental expectations – i.e., as opposed to liquidity shocks – account for the majority of price variation of inflation-linked products at longer horizons. As we later show, the overall conclusion is unchanged if we use alternative expectations measures or tenors.

Column 1 of table 2.2 presents the estimates for our baseline regression using the full post-2000 sample. For a 1 basis point increase in inflation expectations, stocks on average provide an 11.2 basis point return. The estimate is highly significant. Since the standard deviation of the change in 10-year TIPS expectations is 3.6 basis points, these results imply that a one standard-deviation increase in daily inflation expectations is associated with a nearly one third standard deviation increase in daily stock returns. In Column 3, we re-estimate equation 2.1 using inflation swaps to measure inflation expectations and obtain very similar estimates.

Our baseline estimate in column 1 uses the 10-year measures of inflation expectation. In column 2, we explore the relevant horizon at which expectations operate by examining innovations to short-dated and long-dated inflation-linked contracts separately. Specifically, we decompose the 10-year breakeven rate into the 5-year breakeven rate, $\mathbb{E}_t \left[\pi_t^{0-5Y} \right]$, and the component from the forward 5-to-10 year breakeven rate $\Delta \mathbb{E}_t \left[\pi_t^{5-10Y} \right]$. Column 2 shows that higher expected inflation is associated with higher returns at both the short and long

Table 2.2: The exposure of stocks to changes in inflation expectations

	TIPS		Inflation Swaps	
	(1)	(2)	(3)	(4)
	$ret_{market,t}$	$ret_{market,t}$	$ret_{market,t}$	$ret_{market,t}$
$\Delta \mathbb{E}_t[\pi_{t,TIPS}^{10year}]$	11.16*** (13.03)			
$\Delta \mathbb{E}_t[\pi_{t,TIPS}^{0-5years}]$		11.40*** (13.45)		
$\Delta \mathbb{E}_t[\pi_{t,TIPS}^{5-10years}]$		7.44*** (5.01)		
$\Delta \mathbb{E}_t[\pi_{t,Swap}^{10year}]$			10.25*** (8.16)	
$\Delta \mathbb{E}_t[\pi_{t,Swap}^{0-5years}]$				10.93*** (9.08)
$\Delta \mathbb{E}_t[\pi_{t,Swap}^{5-10years}]$				7.36*** (4.25)
N	5,021	5,021	4,155	4,086
R^2	0.11	0.11	0.11	0.12

t statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$


This table displays the results from a regression of the daily value-weighted stock return on change in daily 10 year breakeven inflation expectations. Columns (1) and (2) measure breakeven inflation expectations using TIPS, while columns (3) and (4) use inflation swaps. Units of both the dependent and independent variable are in basis points, i.e. a coefficient of 10 indicates that a 1 basis point increase in inflation expectations is associated with a 10 basis point increase in daily stock returns. t -statistics are based on HAC standard errors with 4-lags.

ends of the term structure. There is some evidence that the effect is stronger in the short end ($\hat{\beta} = 11.40$) than the long end ($\hat{\beta} = 7.44$). Nevertheless, more relevant to our paper, the positive exposure of stock to changes in inflation expectations is not limited to a single expectations horizon but is robust across the term structure. Column 4 repeats the exercise using the swap-based inflation expectation measures and confirms the findings of column 2.

Next, we explore whether the stable positive relationship observed for stock market returns is also present in the cross-section of stock returns. To test this, we re-estimate equation 2.1 on the value-weighted return of each of 30 Fama-French industries. Figure 2.1 plots the industry-portfolio specific coefficients β_π . While there is heterogeneity among industries, all 30 industries have positive and significant exposure to inflation expectations that are similar in magnitude to the market portfolio. Hence, the positive exposure of market returns to inflation expectations does not seem to be driven by a few stocks but rather is a trait across stocks. The observed heterogeneity is consistent with Roussanov, Liu, and Fang (2021): we find that coal, chemicals, and oil—all energy-adjacent industries—are among the most sensitive to changes in expected inflation.

Given the dynamic macroeconomic environment post-2000—including changing nominal-real covariance, a recession, and a zero-lower bound period—a natural concern in studying the conditional effect of inflation is whether it is driven by a specific sub-period. We therefore check whether the positive results are robust across the time series. We estimate a 750 calendar day (approximately two years) rolling window of equation 2.1 and plot the coefficients in figure 2.2. Despite the relatively short rolling window used for the analysis, the stability and significance of the estimates are striking. Apart from a brief period during the financial crisis, the relationship has been positive and robust since 2000. Additionally, the rolling regression using swap-implied inflation expectations yields an essentially identical


Figure 2.1: Inflation Betas Across 30 Fama-French Industries



This figure illustrates the response of industry portfolios to changes in expected inflation. Each point represents the coefficient from the bivariate regression of the value-weighted stock return in the corresponding industry on the daily change in 10-year TIPS-implied expected inflation. Point estimates and 95% robust confidence intervals are plotted for each of the Fama-French 30 industries.

conclusion to the TIPS-based results.

Figure 2.2: Rolling Inflation Expectations Beta



This figure shows the coefficients from daily rolling regressions with 750-calendar day windows of value weighted stock returns on changes in daily inflation expectations. The blue line uses TIPS-based measures, while the red line uses swap-based measures. The shaded areas represent 95% robust confidence intervals.

2.4 Causal Effect of Changes in Inflation Expectations

We now estimate the causal effect of inflation expectations on stock returns. The estimates in the previous section summarized the joint behavior of stock returns and inflation expectations we observe in equilibrium. While valuable for assessing if stocks hedge changes in inflation expectations, the results from the unconditional analysis do not reveal the underlying structural relationship.

To properly identify the causal effect of inflation expectations, an ideal experiment would shock investor beliefs about inflation and then trace out the subsequent effect on returns. The obvious empirical challenge is that we cannot selectively shock investors' expectations. Zooming in on CPI dates offers a quasi-experimental setting that approximates randomly shocking investors' beliefs about future inflation.

2.4.1 Identification

Our main approach to causal analysis is to re-run our main specification (see equation 2.1) on CPI days. This approach will reveal the causal relationship if the following two identification assumptions are satisfied,

1. **High-frequency identification assumption:** the only shock that drives changes in inflation expectations and stock returns in the estimation window is the CPI information shock.
2. **Exclusion restriction:** any effect of the information shock on other variables works through inflation expectations.


The high-frequency identification assumption is needed to ensure that our observed effect is due only to the CPI information shock and not other confounders. We use daily, rather than intraday, windows to ensure sufficient time for stock prices to fully adjust to the inflation information shock.⁶ The exclusion restriction is needed to ensure that the observed return response is attributable to the change in inflation expectations. The exclusion restriction assumes that agents use the CPI print only to infer information about inflation, and so any subsequent outcomes are a result of changes in inflation expectations. Crucially, the assumption does *not* imply that agents do not update their beliefs about other macroeconomic

6. We do not believe our results will be significantly impacted by using intraday data. Since the estimates are highly significant and robust under various sub-samples, any intra-day confounders would need to be present for almost every sub-sample to explain away the day-level result.

variables, only that these updates are due to revisions in inflation expectations. For example, after a CPI print agents update their inflation expectations, as a result of which, they may also update other beliefs (e.g., monetary policy, growth, etc.). This exclusion restriction seems plausible given the limited nature of the information contained in CPI reports—it contains only descriptive information about aggregated prices with essentially no analysis.

The causal structure implied by these two identifying assumptions is captured by the DAG in figure 2.3. First, the inflation news shock is the sole driver of the causal effect—reflecting the high-frequency identification assumption. Second, changes in inflation expectations mediate the effect of the shock on all other variables—reflecting the exclusion restriction. Assuming the assumptions are satisfied, the β_π obtained from estimating our main specification on CPI days will give the total causal effect of the change in inflation expectations. The total effect aggregates the direct and indirect effects from the mediating channels of changing (i) cash flow expectations, (ii) interest rate expectations, and (iii) excess return expectations.

Figure 2.3: CPI day regression DAG



2.4.2 CPI release day summary statistics

CPI announcement days are scheduled monthly and released by the Bureau of Labor Statistics. The dates of these announcements are known in advance, and the announcement occurs at 8:30am before market open. In our sample from January 2000 to December 2021, we have

274 CPI releases.

Table 2.3 compares the changes in inflation expectation on CPI days to non-CPI days. As a first sense check of the identification strategy, we find that the average change in inflation expectations is approximately zero, and it is no different on CPI vs. non-CPI days. The zero estimate is consistent with the CPI release leading to unanticipated changes in inflation expectations. As a second sense check, we find that the standard deviation of changes in inflation expectations is significantly higher on CPI release days. If CPI releases contain information of particular relevance for inflation expectations, then we would expect to see higher variance in inflation expectations on CPI release days.⁷

Table 2.3: Inflation expectation measures CPI vs non-CPI days

	All	Non-CPI	CPI	Diff	<i>p</i> -value (Diff)
$\Delta \mathbb{E}_t \left[\pi_t^{TIPS, 10Y} \right]$					
Mean	-0.025	-0.022	-0.087	-0.066	0.8
St. Deviation	3.6	3.6	4.3	0.66	5.1e-05***
Observations	5021	4747	274		
$\Delta \mathbb{E}_t \left[\pi_t^{Swap, 10Y} \right]$					
Mean	7.9e-05	8e-03	-0.14	-0.15	0.63
St. Deviation	3.9	3.9	4.5	0.62	0.0014**
Observations	4155	3937	218		

The units are in basis points. The final two columns report the difference in mean and standard deviations for the given measure on CPI vs. non-CPI days. For means, the *p*-value is associated with a standard t-test test statistic. For standard deviations, the *p*-value is associated with an *F* test for equal variances.

2.4.3 Causal estimates

In table 2.4, we present the results from the CPI day regressions. Column 1 presents our baseline specification. We find a positive and statistically significant coefficient of 11.88

7. Barring periods of significant market stress in 2008 Q3/4 and 2020 Q1, of the five largest single day changes in 10-year breakeven inflation expectations, three occurred on CPI release days.

($t = 4.43$). In other words, a 1 basis point increase in inflation expectations causes a nearly 12 basis point rise in stock prices. Since the standard deviation of the change in 10-year TIPS expectations on CPI days is 4.3 basis points, these magnitudes are large even on news days. As before, column 3 confirms that the estimates are broadly similar when we use swaps to measure inflation expectations.

Table 2.4: CPI day effect of changes in inflation expectations on returns

	TIPS		Inflation Swaps		Sample Split	
	(1)	(2)	(3)	(4)	(5)	(6)
	$ret_{market,t}$	$ret_{market,t}$	$ret_{market,t}$	$ret_{market,t}$	$ret_{market,t}$	$ret_{market,t}$
$\Delta \mathbb{E}_t[\pi_{t,TIPS}^{10year}]$	11.88*** (4.43)				11.13** (2.92)	12.96*** (3.57)
$\Delta \mathbb{E}_t[\pi_{t,TIPS}^{0-5years}]$		12.07*** (4.32)				
$\Delta \mathbb{E}_t[\pi_{t,TIPS}^{5-10years}]$		9.80 (1.97)				
$\Delta \mathbb{E}_t[\pi_{t,Swap}^{10year}]$			10.53** (3.22)			
$\Delta \mathbb{E}_t[\pi_{t,Swap}^{0-5years}]$				10.51** (3.27)		
$\Delta \mathbb{E}_t[\pi_{t,Swap}^{5-10years}]$				10.95* (2.23)		
N	274	274	218	216	137	137
R^2	0.17	0.17	0.17	0.17	0.12	0.28

t statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

This table displays the results from a regression of the CPI day value-weighted stock return on change in daily 10 year breakeven inflation expectations. Columns (1) and (2) measure breakeven inflation expectations using TIPS, while columns (3) and (4) use inflation swaps. Columns (5) and (6) run the TIPS baseline specification but split the sample in half, pre and post Nov 2010 respectively. Units of both the dependent and independent variable are in basis points, i.e., a coefficient of 10 indicates that a 1 basis point increase in inflation expectations is associated with a 10 basis point increase in daily stock returns. t -statistics are based on HAC standard errors with 4-lags.


Columns 2 and 3 decompose the effect across the term structure. We lose statistical power; however, the estimates are still positive and have a similar magnitude both on the short-end and long-end of the expectations term structure. Overall, the positive exposure of stock to changes in inflation expectations is not limited to operating at a single horizon.

Columns 5 and 6 present the sensitivity of the estimates to using different sub-samples. We split the observations in half around the midpoint of November 2010. Hence, each of the two sub-samples has 137 observations. Column 5 presents the estimates for pre-November 2010, and column 6 presents the estimates for post-November 2010 (including November). The estimates in the two sub-samples are remarkably similar and significantly positive. Hence, a few outliers do not seem to be driving our estimates.

Like the unconditional analysis, we re-estimate our baseline specification for each of the 30 Fama-French industry portfolios. Figure 2.4 plots the industry-specific β_π . The individual coefficients are in the neighborhood of the estimate for overall market returns. While there is some heterogeneity, all 30 industries have positive and significant exposure to changes in inflation expectations. The positive effect of higher expected inflation on stock prices does not appear to depend on the industry or sector analyzed.

We next check for evidence of a reversal in the stock returns following the initial response to changes in the inflation expectations. If the price increase reflects fundamental information, prices should stay elevated after changes in inflation expectations. However, if the price change is for non-fundamental reasons — such as temporary mispricings — we would see the price eventually reverse. We check for reversal by regressing cumulative returns h days

Figure 2.4: Inflation Betas Across 30 Fama-French Industries




This figure illustrates the response of industry portfolios to changes in expected inflation. Each point represents the coefficient from the bivariate regression of the value-weighted stock return in the corresponding industry on the daily change in 10-year TIPS-implied expected inflation on CPI release dates. Point estimates and 95% robust confidence intervals are plotted for each of the Fama-French 30 industries.

after the CPI release on the change in inflation expectations on CPI dates,

$$ret_{t+h} = \alpha + \beta_{\pi}^h \Delta \mathbb{E}_t \left[\pi_t^{10Y} \right] + \varepsilon_t, \quad h \in \{0, 1, 2, 3, \dots\} \quad (2.2)$$

Figure 2.5 plots the cumulative forward return. If anything, the cumulative return in the day following the CPI release date is even higher, not lower. At a horizon of 10 days, the point estimates are in line with the same-day effect and show no evidence of reversal. This persistence suggests that CPI releases reveal fundamental information.

Figure 2.5: Reversal



This figure illustrates the persistent effects of the inflation insurance premium at horizons up to 10-days. The black line represents cumulative forward returns n days after a change in inflation expectations (i.e., it connects the point-estimates of β_{π}^h from equation 2.2), while the shaded area shows the 95% robust confidence interval. The sample comprises CPI release dates.

Our baseline results indicate that higher inflation expectations cause equity prices to rise.

Furthermore, this effect is highly robust—the effect is significantly positive regardless of the measure of inflation expectations, the date sample analyzed, the horizon of expectations, and the sector of focus.

2.4.4 *Asymmetric Responses: Disinflation Risk*

A natural question to ask is whether the effect of inflation expectations on stock returns is asymmetric. Several papers have identified that household inflation expectations are more responsive to inflation than to disinflation (Nagel and Yan (2022), Baqaee (2020)), and a longer literature has identified the asymmetric role of inflation on economic fundamentals (for example, due to downward nominal wage rigidity). It is plausible that the effect of inflation on stock returns could vary heterogeneously – both by the current expected inflation environment, and by the sign of the innovation in inflation expectations. To motivate this analysis, we begin by splitting the sample into one of four categories based on the level of expected inflation and the sign of the change in daily breakevens. We classify each CPI day as above or below median in 10 year inflation expectations, and as a positive or negative change in expected inflation on the CPI day. The intersection of these categories provides four subgroups over which to separately estimate the effects.

Table 2.5 shows the results from this analysis. While all coefficients load positively, the effect is strongest for reductions in inflation expectations in environments where the level of expected inflation over the next 10 years is already low; the effect is more than twice as large as the unconditional average. For both high and low expected inflation environments, negative innovations to expected inflation have a larger (i.e. more negative) effect on stock prices than positive innovations to expectations. And for both positive and negative changes in expected inflation, the effect is stronger in low expected inflation environments than in high ones. These patterns are robust to using inflation swaps in place of TIPS, omitting the financial crisis period (2008-2009), and considering all days rather than restricting to CPI

days alone.

Table 2.5: Inflation Beta by 2x2 Sort

	(1)	(2)	(3)	(4)	(5)
		High, +	High, -	Low, +	Low, -
$\Delta \mathbb{E}_t[\pi_{t,TIPS}^{10year}]$	11.73*** (4.40)	7.19 (1.81)	11.53* (2.27)	9.10 (1.93)	24.27** (3.03)
N	277	71	67	65	74
R^2	0.16	0.04	0.12	0.06	0.27

t statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

This table shows the coefficients from regressions of value-weighted returns on changes in daily breakeven over four subsamples of CPI release days. “High” (“Low”) days are ones in which the 10-year expected inflation is greater (less) than its median. “+” (“-”) are CPI days in which the change inflation expectations was positive (negative). Robust t -statistics are in parentheses.

These results may also shed light on the specific pricing risks that investors incorporate when adjusting inflation expectations. While the analysis has framed positive inflation as good news about economic growth, Table 2.5 suggests that the primary effect is through avoiding a *deflationary* environment. Environments in which the economy moves closer to deflation – when expected inflation is already low, and moves lower – is when the stock market prices the largest effect. There are a number of reasons why deflation could be a pertinent risk, including recession risk and zero-lower bound risk. For example, Bok, Mertens, and Williams (2022) develop a model in which the prospect of a binding zero-lower bound makes a low interest rate (low inflation) more costly, as the central bank is unable to lower rates further. This mechanism could give rise to asymmetry featured in table 2.5 if the ZLB constrains the central bank from stimulating the economy. However, in appendix section 2.8.3, we control directly for changes in the option-implied (i.e., risk-neutral) probability of hitting the zero-lower bound, and verify that including this control does little to change the expected inflation beta. We revisit the implications of inflation for economic growth and

interest rates in section 2.6.

2.4.5 *Inflation Surprises vs. Inflation Expectation Surprises*

Why is it that stocks appear to fall in response to high inflation news, but respond positively to positive innovations in inflation expectations? Indeed previous literature has documented a negative relationship between shocks to inflation and stock returns. For example, Gil de Rubio Cruz et al. (2022) document a robust negative relationship between inflation surprises – calculated as the difference between realized CPI and survey-based expectations – and the cross-section of equity returns. In this subsection, we reconcile the opposite signs on inflation surprises and inflation *expectations* surprises on stock returns. We argue that these opposite signs are consistent, as the positive impact of breakeven changes on stock returns is driven by the part of the change that is orthogonal to the inflation surprise.

In Table 2.7, we measure the effect of CPI surprises on aggregate stock returns. We calculate CPI surprises using the actual reported CPI less the survey-based expectation of the CPI announcement. We explore five measures of CPI – total CPI (both year-over-year and month-over-month); CPI ex-food and energy (also YoY and MoM); and seasonally-adjusted CPI. Consistent with prior literature, all inflation surprises are associated with a negative impact on stock returns, though the statistical significance is marginal.

That stock returns covary positively with changes to inflation expectations and negatively with inflation surprises could suggest that inflation surprises and changes in breakeven are negatively correlated. In Panel B of table 2.7, we show that this is not the case. Positive CPI surprises are associated with statistically significant increases in expected inflation for virtually all CPI surprise measures. In other words, when the CPI reading exceeds the survey-based expectation, investors tend to revise upwards their expectations of future infla-

tion. An implication of these opposite signed results is that the positive impact of inflation expectations on stock returns is not primarily driven by the impact of inflation surprises on stock returns, but rather contains distinct informational content relevant to equities. To see this, observe that we can express changes in breakeven as linear in the CPI surprise and a component orthogonal to CPI surprise:

$$\Delta\mathbb{E}[\pi_t] = \beta \text{Surprise}_t + \eta_t$$

Taking covariances with returns, we have:

$$\underbrace{\text{Cov}(\Delta\mathbb{E}[\pi_t], \text{ret}_t)}_{(+)} = \underbrace{\beta}_{(+)} \underbrace{\text{Cov}(\text{Surprise}_t, \text{ret}_t)}_{(-)} + \underbrace{\text{Cov}(\eta_t, \text{ret}_t)}_{(+)}$$

where the signs of the first three terms are from tables 2.4 and 2.7, and the positive sign on $\text{Cov}(\eta_t, \text{ret}_t)$ follows by requirement. To sustain the positive covariance between changes to expected inflation and stock returns, stocks must respond very positively to the component of inflation expectations uncorrelated with surprises. Our paper is therefore consistent with the literature documenting a negative inflation-stock return gap: changes to inflation expectations contain information distinct from inflation surprises that more than compensate for the negative impact of these surprises. These results also explain the challenge in finding large and statistically robust impacts of inflation shocks on stock returns. The direct channel of inflation shocks on stock returns is negative, but is mitigated by the indirect channel of inflation shocks on expectations.

The distinctive effects of inflation and inflation expectations can also explain the response of equities to inflation news in the post-Covid period. During 2021-2023, realized inflation in the US alone has risen above 6%, with some months witnessing year-over-year changes as large as 9%. More pertinently, during this period, when realized inflation has surprised to

the upside, stock prices have tended to decline sharply. A natural question is whether the relationship between stock returns and expected inflation has now turned negative.

While equities have responded negatively to realized inflation, we show that this effect is not driven by long-term inflation expectations. In fact, the 2 to 10 year forward inflation rate is a positive and significant predictor of stock returns. To show this, we repeat the analysis from the previous section, restricting our sample to the last two years (November 2021 through October 2023), when inflation was high.⁸

Table 2.6 presents the results from this exercise. Columns (1) and (2) show the coefficients from a regression of stock returns on the unconditional daily changes in inflation expectations at 10-year and 2-year tenors. While the coefficients are smaller and insignificant, the point estimates are positive. Columns (3) and (4) provide evidence that the lower point estimates are not due to a changing relationship between long-term inflation expectations and stock returns. In column (3), we show that changes in the 2-10 year forward rate are significant and positive predictors of stock returns. And restricting to the 25 CPI days in the post-2021 sample, the point estimate rises to a (marginally insignificant) 14.69.

Why have the coefficients remained positive in the post-pandemic high inflation episode? And why are the responses stronger with forward rates than with spot rates? We argue that both of these results are consistent with the negative relationship between inflation and equity returns and the positive relationship between inflation expectations and stock returns. The short-end of the inflation expectations term structure is highly correlated with realized inflation. If realized inflation, in turn, has a negative impact on equities, then changes in

8. To focus on the most recent stock return data possible, we use daily returns of the S&P500 Total Return index from Yahoo Finance. We continue to measure inflation expectations using treasury-implied breakeven rates.

inflation expectations associated with changes in inflation may offset, leading to coefficients close to zero. Forward inflation expectations, by contrast, net out the near term changes in realized inflation, and instead focus on the longer-term inflation expectations. These forward expectations may better capture the economic regime. To illustrate this, Figure 2.6 plots the year-on-year change in headline CPI against the 2-year, 5-year and 10-year instantaneous inflation forwards constructed using TIPS and treasuries. The figure shows that, despite realized inflation rising to 9%, inflation expectations have remained anchored at around 2%. The forward rates for the three tenors are virtually indistinguishable, suggesting that the recent increase in inflation is perceived by the market as a purely temporary phenomenon operating in the same inflation regime. Changes in inflation expectations isolated from spot realized inflation – the forward expected inflation – continue to have a significant and positive effect on stock returns.

Table 2.6: Baseline Regression for November 2021 to October 2023 Sample


	(1)	(2)	(3)	(4)
$\Delta \mathbb{E}_t[\pi_{t,TIPS}^{10year}]$	2.06 (1.60)			
$\Delta \mathbb{E}_t[\pi_{t,TIPS}^{2year}]$		0.23 (0.23)		
$\Delta \mathbb{E}_t[\pi_{t,TIPS}^{2-10year}]$			2.36* (1.98)	14.69 (1.44)
CPI days	No	No	No	Yes
N	475	475	475	25
R^2	0.01	0.00	0.01	0.09

t statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

This table shows the regression of the S&P500 returns on breakeven inflation expectations of different tenors. t-statistics are based on HAC standard errors with 4-lags.

Figure 2.6: Realized Inflation vs Inflation Expectations



This figure shows the time series of forward rates on breakeven inflation across different tenors. 2-year, 5-year, and 10-year forward rates are calculated using the differences in implied breakevens for corresponding TIPS contracts. For comparison, we also plot the year-over-year realized inflation series in dark blue.

Table 2.7: Inflation Surprises vs. Inflation Expectation Surprises

	MoM	YoY	Core (MoM)	Core (YoY)	Core (SA)
	(1)	(2)	(3)	(4)	(5)

Panel A: CPI Surprises and Stock Returns

	<i>Dep. Variable: $ret_{market,t}$</i>				
CPI Surprise	-0.950 (0.905)	-0.126 (0.848)	-1.902* (0.750)	-0.520 (0.848)	-0.396 (0.369)
Observations	302	220	300	219	108

Panel B: CPI Surprises and Changes in Exp. Inflation

	<i>Dep. Variable: $\Delta \mathbb{E}_t[\pi_{t,TIPS}^{10Y}]$</i>				
CPI Surprise	3.169 (2.783)	4.871* (2.227)	10.31*** (2.663)	9.659*** (2.485)	3.909*** (1.121)
Observations	279	222	278	221	110

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

This table contrasts the relationship between inflation surprises and stock returns with the relationship between changes in inflation expectations and stock returns. Panel A shows the results of a regression of value-weighted stock returns on 5 different measures of CPI surprises: month-over-month CPI (MoM), year-over-year CPI (YoY), CPI ex-Food and Energy month-over-month (Core (MoM)), CPI ex-Food and Energy year-over-year (Core (YoY)), and Core CPI Seasonally Adjusted (Core (SA)). Surprises are calculated using survey data from Bloomberg. Panel B shows the results from a regression of changes in TIPS-implied breakeven on each of these same 5 CPI surprise measures. Robust standard errors are in parentheses.

2.4.6 Changes in beliefs about future expectations or changes in term premia?

A natural question when using market-based rather than survey-based inflation expectations is to what extent the measure is contaminated by changes in term premia. For example, when new information about inflation is published on CPI days, do changes in market-based inflation expectations predominantly reflect changes in investor beliefs about future inflation, or changes in the term premium?

In the appendix 2.8.2 we conduct three exercises that suggest that changes in term premia are not driving our results. First, using Fama and Bliss (1987) regressions, we show that unconditionally changes in term premia do not play a substantial role in driving changes breakeven-inflation rates. Second, we directly control for changes in term premia, estimated using Kim and Wright (2005) three-factor affine term structure model, and find that they do not significantly alter our estimates. In terms of the direct effect of term premia changes on stock prices, while it is significant for the full sample, it is insignificant when we condition on CPI days. Finally, consistent with investor beliefs changing, we show that changes in market-based inflation expectations are a significant predictor of inflation forecast revisions by analysts.

2.5 Disentangling the Channels

The estimates in the previous section establish that changes in inflation expectations cause stock prices to rise. However, the estimates do not reveal the channel by which the causal relationship works. According to the present discount identity (Campbell and Shiller (1988)), the stock price rise must be due to changes in expectations about (i) dividends, (ii) real rates, or (iii) excess returns. In this section, we use four different approaches to assess the

prominence of these three channels. While each of the methods has its shortcomings, all suggest that the rise in stock prices following higher inflation expectations is mainly due to lower future expected excess returns.

We start by providing heuristic evidence showing that a using the market's β_{mkt}^π and a stocks i 's β_i^{mkt} is a good predictor of the stocks β_i^π . Our next two methods take advantage of the exogeneity of our inflation expectations shock to decompose the response of market returns. We first apply the VAR method proposed by Bernanke and Kuttner (2005) to decompose the response of stock returns into the three channels. We find the effect is almost entirely driven by changes in future excess returns. Next, we use market based measures of dividend and real rate expectations to decompose the response of β_π into its respective channel. This exercise also suggests the primacy of changes in expected excess returns driving the price response. The final method assess if variables typically associated with firm-payout policy can explain the cross-sectional variation in β_π . We find little-to-no relationship, suggesting changes in dividend expectations do not seem to be an important channel.

2.5.1 *Market factor implied β_π*

Our first approach provides heuristic evidence for the role of changes in expected excess returns in driving the observed cross-sectional variation in β_π . We study the estimates that should be observed if returns are driven solely by innovations to the market factor; this is an exercise similar in spirit to the one done by Bernanke and Kuttner (2005) for monetary policy shocks. In and of itself, our empirical exercise is atheoretic, as it does not take a stance of whether or not the market factor reflects risk compensation. It only asks if changes in the market factor can explain the observed heterogeneity in β_π . However, if CAPM were the *true* risk model, our exercise can be interpreted as assessing if changes in risk premia

explain the response of stock prices to changes in inflation expectation.

To that end, suppose that returns are described by the one-factor model,

$$r_{i,t} = \alpha + \beta_i^{mkt}(r_{mkt,t} - r_t^f) + \varepsilon_{i,t}$$

If changes in inflation expectations only affect prices through market returns, then stock i 's response to inflation expectation is given by,

$$\hat{\beta}_i^\pi = \hat{\beta}_i^{mkt} \cdot \hat{\beta}_{mkt}^\pi$$


where $\hat{\beta}_{mkt}^\pi$ is the response of market returns to changes in inflation expectations. In other words, the response of stock i to changes in inflation expectations should be directly proportional to the aggregate market response to those changes. Due to estimation and approximation noise, this equation will not hold exactly; however, if the magnitudes of $\hat{\beta}_i^\pi$ and $\hat{\beta}_i^{mkt}\hat{\beta}_{mkt}^\pi$ are close, it provides indicative evidence for the underlying assumptions.

Figure 2.7 compares the two estimates across the deciles of inflation beta. Consistent with the role of expected excess returns, the market beta implied estimates are close to those obtained through the direct regression. (In all cases, the market beta implied estimates are well within the 95% robust confidence intervals of the direct estimates).

2.5.2 VAR decomposition of β_π

We employ a VAR to decompose the market return to inflation expectation shocks into contributions from changes in expectations of (i) dividends, (ii) real rates, and (iii) excess returns. This exercise is akin to those first employed by Campbell and Ammer (1993) to

Figure 2.7: Market beta implied estimates vs. direct estimates of β^π



This figure shows the comparison of directly-estimated expected inflation betas with market beta implied expected inflation betas. The black dots (direct estimates, $\hat{\beta}_i^\pi$) show the coefficients from regressing the daily stock returns of the given portfolio on changes in 10-year breakeven on CPI days. Each point represents a different decile of stocks by market beta, and the associated error bars show 95% robust confidence intervals. The red dots (market beta implied estimates, $\hat{\beta}_i^{mkt} \cdot \hat{\beta}_{mkt}^\pi$) show the product of each portfolio's market beta with the aggregate response of the value-weighted market portfolio to changes in expected inflation (i.e., 11.88). Sorted market portfolios are obtained from Open Asset Pricing Data Chen and Zimmermann (2022), and the market Beta for each portfolio is estimated using monthly data post-2000. Inflation expectations are measured using TIPS.

decompose stock returns and then modified by Bernanke and Kuttner (2005) to decompose the response of stock prices to monetary policy shocks.

The exercise relies on the Campbell and Shiller (1988) insight that we can express changes in excess equity returns as the sum of the expectations of discounted future dividends, the real interest rate, and future excess returns:

$$e_{t+1}^y = \tilde{e}_{t+1}^d - \tilde{e}_{t+1}^r - \tilde{e}_{t+1}^y$$

where $\tilde{e}_{t+1}^x := \mathbb{E}_{t+1} \left[\sum_{j=0}^{\infty} \rho^j \Delta x_{t+1+j} \right] - \mathbb{E}_t \left[\sum_{j=0}^{\infty} \rho^j \Delta x_{t+1+j} \right]$ for variable x , and ρ is the intertemporal discount factor. Under a given value for the intertemporal discount factor, we can decompose the response of stock prices into their three channels by estimating the response of various expectations to changes in inflation expectations.

For our empirical exercise, we follow the standard implementation. We set $\rho = 0.9962$, and use a monthly frequency six-variable and six-lag VAR. The VAR includes CRSP value-weighted excess returns and 1-month real rates (measured as the 1-month T-bill return deflated using CPI), and treats dividends as a residual. In addition to excess returns and real rates, we include four predictor variables: 1-year treasury yield, log CPI, log IP, and smoothed price-dividend ratio for the CRSP value-weighted index.⁹

As in Bernanke and Kuttner (2005), we first estimate the unconditional VAR. We begin estimation in 1990 to take account of the recent structural break in the correlation between bond and equity returns Campbell, Pflueger, and Viceira (2020). We then regress the one-step ahead forecast errors from the VAR on the CPI day changes in inflation expectations.

9. Dividends are smoothed using a three month backward looking moving average. The price dividend ratio is also linearly de-trended to factor in changes in payout policy overtime.

We adhere to the Bernanke and Kuttner (2005) two-step estimation procedure as it allows us to achieve more stable estimates for the unconditional VAR in spite of a shorter time series for our shock measure. We then use the parameter estimates to decompose the response of stock prices to our inflation expectation shock into contributions from changes in (i) expected excess returns, (ii) expected real rates, and (iii) expected dividends. Finally, we evaluate standard errors using block bootstrap.

Table 2.8 summarizes the results from the VAR decomposition exercise. Consistent with our univariate regressions, current excess returns rise by 14bp in response to a 1pp CPI-day change in inflation expectations. Almost all of this rise can be attributed to a fall in future expected returns. In fact the estimated contribution from changes in expected real rates and expected dividends cannot be distinguished from zero. This finding is particularly stark given that we treated dividends as a residual in the VAR.

Table 2.8: VAR decomposition of excess return response

	Total effect	Contribution		
	Current excess return	Future excess returns	Future real rates	Future dividends
Estimates	14.11*** (13.91)	14.27*** (2.80)	0.16 (0.07)	-0.32 (-0.15)
<i>N</i> first step	383	383	383	383
<i>N</i> second step	264	264	264	264

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

This table shows the VAR decomposition of the excess return response to changes in inflation expectations. Following the method of Bernanke and Kuttner (2005), excess returns are predicted using monthly data on a six-lag, six-variable VAR. The variables are excess stock returns, real interest rates, 1-year treasury yields, log CPI, log IP, and a smoothed price-dividend ratio for the CRSP value-weighted index. Estimation begins in 1990 and runs through 2022. Block-bootstrap t-statistics in parentheses.

Overall, the VAR decomposition suggests that the main driver of the stock price response to changes in inflation expectations is due to changes in expected future excess returns.

2.5.3 Mediation Regression Decomposition of β_π

Next, we use market-based measures for the cash flow growth expectations and real rates to control for the mediating channels directly. We run the following three specifications,

$$ret_t = \alpha_0 + \beta_\pi \Delta \mathbb{E}_t[\pi_t^{10year}] + \varepsilon_{0,t} \quad (2.3)$$

$$ret_t = \alpha_1 + \beta_\pi^{cf+er} \Delta \mathbb{E}_t[\pi_t^{10year}] + \gamma \Delta r_t^{nominal} + \varepsilon_{1,t} \quad (2.4)$$

$$ret_t = \alpha_1 + \beta_\pi^{er} \Delta \mathbb{E}_t[\pi_t^{10year}] + \gamma \Delta r_t^{nominal} + \lambda \Delta \mathbb{E}_t[cf_t] + \varepsilon_{2,t} \quad (2.5)$$

Equation 2.3 is identical to our main specification. In this, β_π captures the total effect from channels (as shown in the DAG in figure 2.3 and estimated empirically in the previous section). Equation 2.4 controls for the change in the nominal risk-free rate, which we measure using the 10-year treasury yield.¹⁰ By controlling for the nominal yield, the coefficient on inflation in this regression captures the effect of inflation not associated with changes in the risk-free component of the nominal discount rate. Thus this coefficient, β_π^{cf+er} , reflects the combined effects of changes in expected dividends and expected excess returns. Finally, equation 2.5 additionally controls for changes in dividend growth expectations, which we measure using 24-month dividend strips for the S&P500 Gormsen and Kojen (2020).¹¹ Overall, controlling for changes in real rates and changes in dividend growth expectations restricts β_π^{er} to operate only through the remaining channel, the expected excess return channel.

Comparing coefficients across the three regressions allows us to recover the effect from each channel. The cash flow effect is identified as the difference between the combined

10. We use nominal yields since the cash flow growth measure is in nominal terms. Hence, the nominal yield captures the component of the (nominal) discount rate affected by the (nominal) risk-free rate.

11. Ideally, we would like to use longer horizon growth expectations, but these are unavailable. That said, longer-term growth rates are notoriously difficult to forecast, so investors likely assume they are equal to their unconditional mean. We would not expect long-run dividend growth expectations to move around significantly.

cash flow and excess return effect and the individual excess return effect: $(\beta_{\pi}^{cf+er} - \beta_{\pi}^{er})$. Likewise, the interest rate effect is identified as the difference between the cumulative effect of inflation and the combined cash flow and excess return channel: $(\beta_{\pi} - \beta_{\pi}^{cf+er})$. Together these provide estimates of the three channels:

$$\beta_{\pi}^r = \beta_{\pi} - \beta_{\pi}^{cf+er} \quad (2.6)$$

$$\beta_{\pi}^{cf} = \beta_{\pi}^{cf+er} - \beta_{\pi}^{er} \quad (2.7)$$

$$\beta_{\pi}^{er} = \beta_{\pi}^{er} \quad (2.8)$$

We provide a proof for the consistency of these estimators in appendix 2.8.1. Note that the constituent effects sum to the total effect

$$\begin{aligned} \beta_{\pi}^r + \beta_{\pi}^{cf} + \beta_{\pi}^{er} &= (\beta_{\pi} - \beta_{\pi}^{cf+er}) + (\beta_{\pi}^{cf+er} - \beta_{\pi}^{er}) + (\beta_{\pi}^{er}) \\ &= \beta_{\pi}, \end{aligned}$$

so that we divide each coefficient by the total effect to compute the share of the cumulative inflation premium attributed to each channel.

Before turning to the results, several clarifying points merit discussion. As a caveat on the interpretation, since our dividend growth measure is taken under the risk-neutral measure, and since the risk-free rate measure includes term premia, including them as controls at least partially controls for changes in risk premia also. We may therefore introduce some attenuation bias in β_{π}^{er} . However, changes in term premium do not seem to be associated with stock returns on CPI release days (see appendix section 2.8.2). Second, to successfully control the mediating channels, we require the high-frequency identification assumption and the exclusion restriction to be satisfied—which is why we focus on CPI days for this formal decomposition exercise. In fact, the high-frequency setting is critically what allows us to

avoid the bad control problem (see appendix section 2.8.1).

In table 2.9, we show the results from this decomposition for both TIPS-based and swap-based inflation expectations measures on CPI announcement dates. Reading across columns (1)-(3), adding in controls for interest rates and changes in expected cash flows increases the cumulative inflation coefficient from 11.88 to 12.06. Using our decomposition, the interest rate channel accounts for -0.878 of the 11.88 coefficient, the cash flow channel accounts for 0.703 of the 11.88, and the remainder we attribute to the expected excess return channel.

As a robustness check, in columns (4)-(6), we perform the same decomposition with the swap-based inflation expectations measure. Although the point estimates are different, for both measures the coefficient on expected excess returns—i.e., the component of returns not due to changes in expected cash flows or interest rates—is significantly different from 0. At the same time, we cannot reject the zero effect null hypothesis for the cash flow and interest rate channels. This decomposition provides strong evidence that high inflation expectations increase returns by reducing future expected excess returns and complement the results in the previous subsections.

Since industries differ in their duration (i.e., their response to changes in interest rates), it is helpful to consider whether the decomposition explored above applies cross-sectionally. In figure 2.8, we re-estimate the inflation betas on the 30 Fama-French industries with and without controlling for changes in 10-year real interest rates on these CPI days. The point estimates and ordering of industries is almost entirely unchanged, and all 30 industries retain a significant inflation premium—due to data limitations, we do not have stock-specific dividend growth expectation measures. For completeness, we show an analogous approach showing the unconditional *interest* rate betas (i.e., the coefficients from a bivariate regression of industry returns on changes in 10-year interest rates). These two exercises demonstrate

Table 2.9: Decomposition of Inflation Premium by Channel

	TIPS			Swap		
	(1)	(2)	(3)	(4)	(5)	(6)
$\Delta \mathbb{E}_t[\pi_{t,TIPS}^{10year}]$	11.88*** (4.43)	12.76*** (4.06)	12.06*** (3.63)			
$\Delta \mathbb{E}_t[\pi_{t,Swap}^{10year}]$				10.53** (3.22)	10.07** (2.87)	10.58** (3.02)
$\Delta r_t^{10year(Nominal)}$		-1.20 (-0.72)	-2.36 (-1.35)		0.85 (0.46)	-0.15 (-0.10)
$\Delta \mathbb{E}_t[cf_t]$			0.14** (2.64)			0.11 (0.77)
β_i			-0.878			0.451
β_{cf}			0.703			-0.501
β_{rp}			12.059***			10.575**
N	274	274	217	218	214	163
R^2	0.17	0.17	0.24	0.17	0.17	0.20


t statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

This table shows the decomposition of the cumulative inflation expectation premium, β_π , into the interest rate, cash flow, and risk premium channels— β^r , β^{cf} , and β^{er} respectively. Columns (1)-(3) estimate equations (2.3)-(2.3), with the β 's in the table notes formed according to equations (2.6)-(2.8). For the derived coefficients, the test statistics are derived using the Seemingly Unrelated Regressions (SUR) model.

that even in the cross-section, the path of future interest rates is not responsible for the positive causal effect.

Figure 2.8: Inflation Betas By Fama-French Industries with and without Interest Rate Controls



This figure illustrates the response of industry portfolios to changes in expected inflation, with and without controlling for changes in interest rates. The top left replicates the results from figure 2.4, showing the inflation betas from a bivariate regression of industry-level value-weighted stock returns on changes in 10-year TIPS breakevens. In the figure below (bottom left), we plot those same inflation betas after controlling for interest rates. The right two figures show the betas on interest rates in the bivariate regression of value weighted industry returns on changes in 10-year rates (top right) and the coefficients on the interest rate changes after controlling for inflation. In all figures, the sample comprises CPI release dates. Point estimates and 95% robust confidence intervals are plotted for each of the Fama-French 30 industries.

2.5.4 *Expected dividend channel and cross-sectional β_π*

Our final approach assesses the importance of the expected dividend channel by seeing if variables typically associated with firm-payout policy can explain the cross-sectional variation in β_π . We examine returns on portfolios sorted on a large number of firm characteristics that capture leverage, coverage, profitability, and other characteristics.¹² For each characteristic, we sort firms into deciles as of each CPI date, construct the value-weighted daily returns within each decile, and regress those returns on changes in daily inflation expectations across the CPI sample.

Figure 2.9 shows the coefficients and point estimates across six specific characteristics: asset tangibility, total payout ratio, book leverage, market leverage, cash ratio, and profit margin. We selected these characteristics because they relate to common conceptions about which types of companies are likely to be “winners” from higher inflation. *Ceteris paribus*, equity holders of firms with high leverage should disproportionately benefit because higher inflation erodes the real value of nominal debt. However, we do not find this in the data—there is little evidence that the estimated β_π is higher for firms with higher leverage regardless of if we measure it as book or market leverage. Similarly, asset tangibility and cash ratio—which measure the firm’s relative quantity of real and nominal assets—have no discernible pattern. We also look at heterogeneity by firm profitability, where we measure profitability as net income over sales. For several reasons, profitability may be relevant to inflation’s effect on equity value. Firms with higher profit margins can capture a larger share of sales; since sales reflect revenue from goods and services sold at nominal prices, high profitability firms may serve as a better hedge. Additionally, profitability may serve as a rough proxy for

12. We look at asset tangibility, book leverage (debt/assets), book leverage (debt/(debt + shareholders equity), common equity, current ratio, debt/ebitda, dividend payout ratio, ebit coverage, ebitda coverage, ebitda ratio, interest ratio, market leverage (long term, short term, and total), profit margin, return on equity, and total payout ratio.

market power. According to popular wisdom, firms with high market power may be better able to pass higher costs to consumers and so can better withstand inflationary shocks (although there is little academic evidence for this channel). Despite these arguments, we find no clear pattern in stock responses across deciles of profitability.

The only pattern that shows some degree of monotonicity is the total payout ratio. Firms with higher payout ratios appear to have smaller increases in returns in response to an increase in inflation expectations. This evidence is consistent with the total payout ratio proxying a firm's stock's "cash likeness". Higher inflation deteriorates the real value of nominally denominated payouts; hence, *ceteris paribus*, the price of these higher payout ratio firms should decline relatively more. Nonetheless, the large standard errors do no longer allow us to reject the null that the highest and lowest decile have different returns.


In short, while every decile portfolio has a positive inflation beta significantly different from zero, it is difficult to conclude that differences in firm financial characteristics account for differences in return response. The chosen firm characteristics reflect variables related to cash flow, and in turn their future dividends. The absence of significant results suggests that changes in expected dividends are not the main driver of the positive effect of inflation expectations on stock returns.

2.6 Towards an Economic Interpretation


The preceding results have established that increases in expected inflation lead to an increase in stock prices, and that this increase is driven by a decline in expected future excess returns. However, the analysis thus far has been model agnostic. That is, while we establish the relationships in reduced-form, we have carefully avoided attaching any structural interpretation to the results.

Figure 2.9: Inflation Betas by Total Payout Ratio and Leverage


(a) Book Leverage




(b) Market Leverage




(c) Asset Tangibility




(d) Cash Ratio



(e) Profit Margin



(f) Total Payout Ratio



These figures show inflation betas on decile sorts of various firm-level financial characteristics. Portfolios for each characteristic are formed as of each CPI day and split into deciles. The points in the figure represent the estimated coefficient and 95% robust confidence intervals from a regression of value-weighted stock returns within each decile on changes in daily inflation expectations. The regression sample comprises CPI release dates, with TIPS-implied inflation measures. Book leverage is defined as debt/assets; market leverage is debt/(debt+share repurchases + common equity); Asset tangibility is defined as PPE/assets; cash ratio is cash and short term investments over assets; profit margin is net income/sales; and total payout ratio is (dividends + share repurchases)/net income.

In this section, we attempt to go a step further and provide an economic interpretation of our results. We argue that our results are consistent with information about higher future inflation, signalling higher future output. In other words, higher inflation expectations result in investors lowering the probabilities they assign to bad states of the world, which in turn lowers risk premia and raises stock prices David and Veronesi (2013). Furthermore, we show that the central mechanism underpinning this economic mechanism does in fact obtain in the data, namely, information about higher future inflation strongly predicts higher future output.

2.6.1 Inflation Expectations and Future Output

The economic interpretation described above provides a clear and testable empirical prediction: if information shocks lead investors to revise their inflation expectations upwards, we should expect higher output in the future as well. The challenge in testing this prediction is that (i) the time series evolution of output is determined through the joint dynamic behavior of several macroeconomic variables, and (ii) an initial exogenous information shock about future inflation is required to trace out the causal response. We solve the former by using a VAR to model the endogenous dynamics of macroeconomic variables, and we solve the latter by leveraging our high-frequency inflation information shock.

Specifically, we estimate a monthly frequency VAR with five endogenous variables and twelve lags, and we estimate it for the sample period of January 2000 through December 2021. For endogenous variables, we include: (i) output, measured using log industrial production (ii) price level, measured using log CPI (iii) inflation expectations, measured using the 10-year breakeven rate, and (iv) policy rate, measured using the 1-year ZCB treasury rate, and (v)

and implied volatility, measured using the VIX.¹³ Finally, we include unanticipated changes in the 10-year breakeven inflation rate on CPI days (cumulated to monthly frequency) as an exogenous variable in the VAR. We represent the VAR as:

$$Y_t = \phi z_t + \sum_{j=1}^{12} A_j Y_{t-j} + u_t \quad (2.9)$$

where Y_t is a vector containing our five endogenous variables and z_t is the unanticipated changes in the 10-year breakeven inflation rate on CPI days, cumulated up to monthly frequency.

Figure 2.10 displays the response of endogenous variables for a 10 basis point, information-induced shock to inflation expectations—assuming our shock satisfies the high frequency identification and exclusion assumptions, these IRFs can be interpreted as the causal response to the information shock. Mechanically, expected inflation rises instantly by approximately 10 basis points. Consistent with rational expectations, higher (realized) inflation appears over the following months, with the price level rising approximately 0.25 percentage point and remaining at that level for several months. We also observe the reaction of monetary policy to elevated inflation; the short-term interest rate rises by roughly 12 basis over the subsequent year, after which it declines back to steady state.

The most relevant figure for the proposed economic mechanism is captured in the response of output and volatility to the rise in inflation expectations. Panel (d) of figure 2.10 traces

13. We use the 1-year treasury rather than the Fed Funds rate because for a large portion of the sample, the Fed Funds was at the zero lower bound and hence had little to no variation. We use industrial production to measure output, rather than GDP, because it is available at the monthly level (GDP is only available quarterly). We use VIX to measure volatility since it best represents investors' expectations of future volatility. However, since VIX is a risk-neutral measure, in Appendix section 2.8.4 we re-estimate the IRF using realized volatility to show that the change in VIX is not driven by CPI-day changes in the variance risk premium.

the response of industrial production to a 10 basis point increase in inflation expectations. Output rises by 0.7% over the next year, and stays above steady state for at least three years.

At the same time, implied volatility, measured with VIX, drops more than 2 percentage points immediately, and remains depressed for nearly 18 months. In other words, higher inflation expectations positively predict higher growth and negatively predicts future volatility.


Overall, our finding is consistent with investors using information about higher future inflation, to revise up their probabilities of future good states of the world, thus reducing the compensation they demand for holding risk. This explanation is in line with David and Veronesi (2013), but also Fama (1981) “proxy” explanation for why higher realized inflation generated negative stock returns during the stagflation era—higher inflation at the time predicted lower future output. It also sheds light on the persistent positive correlation even in the post-pandemic high inflation episodes. While the anchored long-term expectations illustrated in figure 2.6 do not directly speak to the inflation-growth covariance, they suggest that investors expect a similar inflation regime in the future.

2.7 Conclusion

This paper documents that since the 2000s, stocks have provided a hedge against changes in inflation expectations. A 1 basis point positive change in inflation expectations is associated with an 11 basis point higher return. This relationship is highly significant and robust to the choice of the inflation expectations measure and estimation period, and the effect is present across the cross-section of stocks. The positive relationship remains intact in the post-pandemic high inflation episode.

Zooming in on CPI releases dates, we find that the positive association is causal. Information about higher future inflation causes investors to raise their inflation expectations;

Figure 2.10: Impulse Responses to an Increase in Expected Inflation



This figure shows the IRF for a response to a 10 basis point exogenous increase in CPI day inflation expectations. The top row shows the response of the level of 10-year breakevens and one-year treasury yields; the second row shows the response of log CPI and log industrial production, and panel (e) shows the response of VIX. The results are based on a 12-lag monthly VAR spanning January 2000 to December 2021.

stock prices, in turn, rise as well. Furthermore, using several complementary approaches, we find the rise in prices is driven by investors lowering their expectations about future excess returns. First, Bernanke and Kuttner (2005) style VAR decomposition and also high-frequency mediation regressions attribute the rise in stock market returns almost entirely to changes in expected excess returns. Secondly, variables traditionally associated with firm cashflows are not systematically associated with stocks' cross-sectional response to changes in inflation expectations.

Overall, our findings are consistent with investors using information about higher future inflation to revise up their probabilities of future good states of the world, thus reducing the compensation they demand for holding risk David and Veronesi (2013). This mechanism reconciles (i) why there is positive association between changes in inflation expectations and stock returns, (ii) why information about higher future inflation causes stock prices to rise, and (iii) why price response is explained by lowering future excess returns (which from the mechanism's lens reflects a decline in risk premium). Finally, consistent with the economic mechanism, we show that information about future inflation predicts higher future output.

2.8 Appendix

2.8.1 Decomposition Result

Suppose x effects y through mediating variables $z_1, \dots, z_n$ such that,

$$z_1 = \alpha_1 + \beta_1 x + \varepsilon_1$$

$$z_2 = \alpha_2 + \beta_2 x + \varepsilon_2$$

$$\vdots$$

$$z_n = \alpha_n + \beta_n x + \varepsilon_n$$

$$y = \gamma_0 + \sum_{i=1}^n \gamma_i z_i + \eta$$

We can substitute all the mediating effects into the final equation to see the total effect (reduced form effect) of y on x ,

$$y = a + bx + \epsilon$$

where $a = \gamma_0 + \sum_{i=1}^n \alpha_i \gamma_i$, $b = \sum_{i=1}^n \beta_i \gamma_i$, and $\epsilon = \eta + \sum_{i=1}^n \varepsilon_i \gamma_i$.

Proposition: If $\mathbb{E}[x\epsilon] = 0$ then OLS consistently estimates the total effect b of x on y .

Corollary: If x is a shock that is orthogonal to y then OLS consistently estimates the total effect b of x on y .

Proposition: The mediated effect of x on y through z_1 (i.e. $\beta_1 \gamma_1$) can be consistently estimated by regressing y on x and controlling for $z_2, \dots, z_n$ if,

1. $\mathbb{E}[x(\varepsilon_1, \eta)'] = 0$

2. $\mathbb{E}[z_i(\varepsilon_1, \eta)'] = 0$ for all $i \geq 2$

Proof. Suppose you run the following regression,

$$y = \delta_0 + \delta_1 x + \sum_{i=2}^N \delta_i z_i + u$$

Substituting in the true model we can see that: $\delta_0 = \gamma_0 + \alpha_1 \gamma_1$, $\delta_1 = \beta_1 \gamma_1$, $\delta_i = \gamma_i$ for all $i \geq 2$, and $u = \varepsilon_1 + \eta$. OLS consistently estimates the parameters if $\mathbb{E}[(x, z_2, \dots, z_n)'u] = 0$. This is guaranteed if the moment conditions in the proposition are satisfied. $\square$

Corollary: If x is orthogonal to y and the only shock to occur, then all mediating effects $\beta_i \gamma_i$ can be estimated.

If x is the only shock to occur, then all changes in the mediating variables can be attributed to x i.e. $\varepsilon_i \approx 0$. Hence, orthogonality of errors across different mediating variables is satisfied. Note this assumption is stronger than only assuming orthogonality. If we were only to assume orthogonality then controlling for mediating factors would introduce a bad control bias Angrist and Pischke (2009).

In most quasi-experimental studies, the only shock assumption is too strong. However, for high-frequency identification the assumption is plausible. In fact, this is typically how we think of high-frequency identification. We focus on a narrow window of time, where we believe the only piece of information released.

2.8.2 Robustness to Changes in Term Premium

In this section, we show that the positive response of stock prices to inflation information is (predominantly) driven by changes in inflation expectations and not changes in term premia. Changes in breakeven inflation rates can be driven by changes in inflation expectations or changes in term premia embedded in treasuries and TIPS. We begin by using Fama and Bliss (1987) style regressions to assess the extent to which unconditional changes in breakeven rates reflect changes in expectations vs. term premia. We find that the role inflation expectations is large, and for longer horizons it is difficult to rule out all the variation is driven by changes in inflation expectations. Next, we turn to assess if term premium changes can explain away our baseline result for CPI days. Unfortunately, we cannot use the Fama and Bliss (1987) regressions to assess the CPI-day drivers, instead we include changes in Kim and Wright (2005) term premium measure as a control in our CPI day regressions and find our point estimates are essentially unchanged.

Expectation Hypothesis

The expectations hypothesis (EH) for bonds states that the expected excess returns on long-term bonds over short-term bonds are constant over-time. We can then link the EH for nominal treasuries and TIPS to breakeven inflation using the following definition,¹⁴

$$y_{nt}^{\pi} := y_{nt}^n - y_{nt}^r \quad (\text{breakeven inflation rate})$$

where y_{nt}^k and f_{nt}^k are the time t yields and forward rates for ZCB bonds maturity in n periods, respectively. And $k = n$ refers to nominal bonds, $k = r$ refers to real bonds, and $k = \pi$ refers to break-even inflation.

14. As is standard in the bond literature testing EH, we will work with log returns.

More formally EH can be stated as,

$$\mathbb{E}_t \left[y_{tn}^k - \frac{1}{n} \sum_{i=0}^{n-1} y_{1,t+i}^k \right] = \theta_n^k \quad (2.10)$$

From the perspective of breakeven inflation, this says that the current breakeven inflation rate is equal the expected path of one-period realized inflation plus a time-invariant constant. As should be apparent from the definitions above, a failure of EH amounts to there being an additional time varying error term. For example, EH would fail if in practice we have,

$$\mathbb{E}_t \left[y_{tn}^k - \frac{1}{n} \sum_{i=0}^{n-1} y_{1,t+i}^k \right] = \theta_n^k + \varepsilon_{n,t}^k \quad \text{s.t.} \quad \text{Var}(\varepsilon_{n,t}^k) > 0$$

This error term is often referred to as term premium, and could reflect rational risk premium or non-fundamental drivers such as sentiment. The definitions above also highlight why the failure of EH for nominal treasuries and TIPS do not necessarily imply a failure of EH for breakeven inflation. For example, if $\varepsilon_{n,t}^r = \varepsilon_{n,t}^y$, EH will fail for nominal treasuries and TIPS, but will be satisfied for breakeven inflation. To this point, in our paper we use changes in breakeven inflation expectations to proxy for changes in investor beliefs about future inflation. This would be correct if,

$$\text{Var}(\Delta \varepsilon_t^\pi) = \text{Var}(\Delta \varepsilon_{n,t}^n - \Delta \varepsilon_{n,t}^r) \approx 0$$

Fama-Bliss Regressions

While there are numerous tests for EH, a common approach is to test for dynamic constraints imposed on bond prices by EH—Fama and Bliss (1987) is one such test. To develop this

test, we need to define the one-period ahead breakeven forward,

$$f_{nt}^{\pi} = (n+1)y_{n+1,t}^{\pi} - ny_{nt}^{\pi} \quad (2.11)$$

By combining the forward rate identity with the definition of EH we can get a dynamic constraint between breakeven forwards and the expected future one-period ahead breakeven rate.

Proposition: *Under EH the breakeven inflation forward rate is the expected future one-period ahead breakeven rate in the future (plus a constant),*

$$f_{nt}^{\pi} = \phi_n^{\pi} + \mathbb{E}_t[y_{1,t+n}^{\pi}] \quad (2.12)$$

Proof:

$$f_{nt}^{\pi} = (n+1)y_{n+1,t}^{\pi} - ny_{nt}^{\pi} \quad (\text{equation 2.11})$$

$$f_{nt}^{\pi} = (n+1)\mathbb{E}_t[y_{n+1,t}^{\pi}] - n\mathbb{E}_t[y_{nt}^{\pi}] \quad (\text{take expectations})$$

$$f_{nt}^{\pi} = \underbrace{(n+1)\theta_{n+1}^{\pi} - n\theta_n^{\pi}}_{:=\phi_n^{\pi}} + \mathbb{E}_t \left[\sum_{i=0}^n y_{1,t+i}^{\pi} - \sum_{i=0}^{n-1} y_{1,t+i}^{\pi} \right] \quad (\text{by equation 2.10})$$


After subtracting the current one-period breakeven rate from both sides of 2.12, we can run the following time-series regression to test if breakeven inflation satisfies the expectations hypothesis,

$$y_{1,t+n}^{\pi} - y_{1t}^{\pi} = \alpha_n + \beta_n(f_{nt}^{\pi} - y_{1t}^{\pi}) + u_{t+n}^{\pi} \quad (2.13)$$

If EH is satisfied we will find $\beta_n = 1$.

We run regression 2.13 at monthly frequency for maturities $n = 1, 2, \dots, 10$ years. As our short-rate rate we use the one-year breakeven inflation rate. As figure 2.11 shows, overall

Figure 2.11: Breakeven inflation Fama-Bliss regressions



we find that β is close to 1. While for maturities less than 5-years we can reject the $\beta = 1$, we cannot do so for longer maturities. In spite of that, the coefficients are generally close to one, and much closer to one than found in the literature for nominal treasuries and TIPS. Hence, overall it seems like majority of the changes in the 10-year breakeven inflation rate reflect changes in investor beliefs, rather than changes in premia.

Controlling for term premium

We look at both measures of inflation expectations (TIPS and Swaps) and both sample periods (all days post-2000 and CPI release days alone). In columns (1)-(4), we first look unconditionally across the full post-2000s sample. In the first column, we re-estimate our main regression specification from equation (2.1), and add a control for changes in the Kim and Wright (2005) 10-year term premium in column 2. As expected, changes in term premia are positively associated with changes in stock returns. However changes in term premia do not themselves explain the response of equities to inflation expectations: for both TIPS and

swaps, the inflation betas remain positive and significant, though smaller in magnitude.

In columns (5)-(8) we zoom in on CPI days. Under our identification assumption, on these days all changes in stock returns should operate through inflation expectations and so the contribution from the term premium should be small. Consistent with this assumption we find that controlling for the term premium does virtually nothing to the point estimates of inflation expectations. Moreover, the point estimates on term premia are insignificant and close to 0. These suggest that changes in term premia are not driving the return response.

Table 2.10: Inflation Expectations Premium After Controlling for Term Premium

	All Days				CPI Days			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\Delta \mathbb{E}_t[\pi_{t,TIPS}^{10year}]$	11.16*** (12.82)	7.69*** (7.31)			11.88*** (4.43)	12.60*** (3.90)		
$\Delta \mathbb{E}_t[\pi_{t,Swap}^{10year}]$			10.25*** (9.06)	7.78*** (6.42)			10.53** (3.22)	10.34** (2.98)
$\Delta \text{Term Premium}_t^{10year}$		9.15*** (6.89)		12.62*** (8.11)		-2.19 (-0.55)		0.49 (0.12)
N	5,021	5,020	4,155	3,953	274	274	218	214
R^2	0.11	0.14	0.11	0.17	0.17	0.17	0.17	0.16

t statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

This table shows the effect of changes on inflation expectations of the stock returns after controlling for changes in the term premium. We use the daily change in 10-year term premium on zero coupon bonds, based on the three-factor model in Kim and Wright (2005). The dependent variable is daily value-weighted stock returns. Changes in inflation expectations are measured using both TIPS (columns 1, 2, 5, and 6) or Swaps (columns 3, 4, 7, and 8). Columns (1)-(4) use the full post-2000 daily sample, while columns (5)-(8) restrict to CPI days. The sample period is 2000-2021. Robust t-statistics are in parentheses.

2.8.3 Robustness to Changes in Probability of Hitting Zero Lower Bound

Given the results on deflation in low inflationary periods, we ask if the results could be generated by the probability of hitting the ZLB, as in Bok, Mertens, and Williams (2022). We use the probability of hitting the ZLB, extracted from options prices, as in Bauer,

Mertens, et al. (2019).¹⁵ Matching the data available from Bauer, Mertens, et al. (2019) restricts us to the Jan 2018- September 2019 sample period. Controlling for changes in the risk-neutral probability of hitting the zero-lower bound does little to change the coefficient on changes in expected inflation. The point estimate on the interaction suggests that perhaps the change in probability could augment it, but it is statistically insignificant.

Table 2.11: Expected Inflation Beta by Change in ZLB Proability (2018-2019)

	All Days			CPI Days		
	(1)	(2)	(3)	(4)	(5)	(6)
	$ret_{market,t}$	$ret_{market,t}$	$ret_{market,t}$	$ret_{market,t}$	$ret_{market,t}$	$ret_{market,t}$
$\Delta \mathbb{E}_t[\pi_{t,TIPS}^{10year}]$	18.94*** (7.41)	14.66*** (5.51)	14.57*** (5.41)	20.74** (2.88)	16.80* (2.11)	15.33 (1.83)
$\Delta Pr_t^{3Y}(ZLB)$		-0.20** (-3.03)	-0.20** (-2.93)		-0.19 (-1.47)	-0.46 (-1.85)
$\Delta \mathbb{E}_t[\pi_{t,TIPS}^{10year}] \times \Delta Pr_t^{3Y}(ZLB)$			0.72 (0.40)			15.37 (1.35)
N	408	407	407	22	22	22
R^2	0.15	0.18	0.18	0.36	0.40	0.47

t statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

This table shows the results from a regression of value-weighted stock returns on changes in expected inflation, changes in the probability of hitting the zero lower bound, and their interaction. Construction of the ZLB is described in detail in Bauer, Mertens, et al. (2019). The sample period is January 1, 2018 through September 17, 2019. Robust t-statistics are in parentheses.

15. We are grateful to Michael Bauer for sharing with us his data.

2.8.4 Robustness of Volatility Response to Inflation Expectations

2.8.5 Survey Expectation Changes and Inflation

In this subsection, we show that changes in inflation expectations reflect changes in investor beliefs about future inflation. To do so, we regress inflation forecast revisions on changes in break-even inflation. If changes in breakeven inflation reflect changes in the inflation expectations of investors, then we should see that these changes should predict changes in surveyed inflation expectations.


In terms of constructing forecast revisions, we use the mean CPI forecasts from the Survey of Professional Forecasters (SPF). Denote the forecast made in period t for the average inflation between periods $t + h$ and $t + h + f$ as $F_t^f \pi_{t+h}$, for example, $F_t^4 \pi_{t+h}$ is the forecast made at time t of the average inflation from quarter $t + h$ to quarter $t + h + 4$. Using this notation, we define the forecast revision as,

$$\eta_t^f(\pi_{t+h}) := F_t^f \pi_{t+h} - F_{t-1}^f \pi_{t+h}$$

Since SPF is conducted quarterly, we can exactly calculate forecast revisions for $f = 1$, but not for $f > 1$. Given this data limitation, we can approximate the revisions for $f > 1$ using the approximation $\eta_t^f(y_{t+h}) \approx F_t^f \pi_{t+h} - F_{t-1}^f \pi_{t+h-1}$. For large f , the $F_{t-1}^f \pi_{t+h-1} \approx F_{t-1}^f \pi_{t+h}$, since a single quarter has a small impact on the overall forecast.

To measure new inflation information made available to forecasters in between surveys, we use changes in the 10-year break-even inflation rate. In terms of timing, the SPF is conducted quarterly, and the respondents receive the survey at the end of the first month of a given quarter, e.g., the survey conducted in Q2, is filled out by the respondents in May (at the earliest). Hence, to measure new information in the interim, we use changes in the breakeven rate between the beginning of the second months of each quarter; e.g., for the Q2 survey, we consider the change in the breakeven rate between February 1st and April 30th.

Figure 2.12: Impulse Responses to an Increase in Expected Inflation



This figure shows the IRF for a response to a 10 basis point exogenous increase in CPI day inflation expectations. The top row shows the response of the level of 10-year breakevens and one-year treasury yields; the second row shows the response of log CPI and log industrial production, and panel (e) shows the response of realized volatility, measured as the time-series standard deviation of daily returns in the CRSP value-weighted index. The results are based on a 12-lag monthly VAR spanning January 2000 to December 2021.

We also construct a measure only using changes in the breakeven rate on CPI release days that fall in the interim between survey dates.

Using our measures, we run the regression,

$$\eta_t^f(y_t + h) = \alpha_h^f + \beta_h^f \Delta \mathbb{E}_t[\pi_{t,TIPS}^{10years}] + \varepsilon_{h,t}^f$$

If changes in breakeven expectations contain information on future inflation, we would expect $\beta > 0$.

Table 2.12 presents the results of this expression. First, in spite of the small number of quarterly frequency observations we find significant evidence of realized innovations in breakeven inflation predicted forecast revisions; consistent with breakeven inflation changing in response to changes in investor beliefs about future inflation. Second, for many of the horizons, innovations explain a large proportion of the changes in forecasts (with the exception of three-quarter ahead forecast, they explain around 20% to 30% of the variation). Third, higher f forecasts reflect a larger segment of the inflation belief term structure; hence, all else equal revisions of larger f forecasts should respond more strongly to changes in the 10-year breakeven rate. Consistent with this prediction, changes in the breakeven rate have particularly high predictive power for $\eta_t^f(\pi_{t+4})$. Finally, only using innovations from CPI days, while we lose power the passthroughs are larger, and significant for the one-quarter ahead and one-year ahead revisions.

Table 2.12: Response of SPF Forecast Revisions to Innovations in Inflation Expectations

	All days				CPI days			
	$\eta_t^1(\pi_{t+1})$	$\eta_t^1(\pi_{t+2})$	$\eta_t^1(\pi_{t+3})$	$\eta_t^4(\pi_{t+4})$	$\eta_t^1(\pi_{t+1})$	$\eta_t^1(\pi_{t+2})$	$\eta_t^3(\pi_{t+3})$	$\eta_t^4(\pi_{t+4})$
$\Delta \mathbb{E}_t[\pi_{t,TIPS}^{10year}]$	2.00** (2.95)	0.57* (2.11)	0.21* (2.41)	0.36*** (4.51)	4.99* (2.19)	1.14 (1.63)	0.42 (1.26)	0.82* (2.52)
N	84	84	84	84	84	84	84	84
R^2	0.22	0.22	0.09	0.29	0.10	0.06	0.03	0.11

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

This table shows the results from a regression SPF forecast innovations on changes in 10-year breakeven inflation in the interim between forecasts. Columns (1) to (4) use all changes in breakeven inflation rates, whereas columns (5) to (8) use changes in breakeven inflation only from CPI days. Robust t-statistics in parentheses.

REFERENCES

- Ang, Andrew, Marie Briere, and Ombretta Signori (July 2012). “Inflation and Individual Equities”. In: *Financial Analysts Journal* 68.4, pp. 36–55. ISSN: 0015-198X. DOI: 10.2469/faj.v68.n4.3. URL: <https://doi.org/10.2469/faj.v68.n4.3> (visited on 04/29/2022).
- Angrist, Joshua D and Jörn-Steffen Pischke (2009). *Mostly harmless econometrics: An empiricist’s companion*. Princeton university press.
- Bahaj, Saleem et al. (2023). “The Market for Inflation Risk”. In: *Available at SSRN 4488881*.
- Baqae, David Rezza (2020). “Asymmetric inflation expectations, downward rigidity of wages, and asymmetric business cycles”. In: *Journal of Monetary Economics* 114, pp. 174–193.
- Barber, Brad M et al. (2013). “The earnings announcement premium around the globe”. In: *Journal of Financial Economics* 108.1, pp. 118–138.
- Barberis, Nicholas, Robin Greenwood, et al. (Jan. 2015). “X-CAPM: An extrapolative capital asset pricing model”. en. In: *Journal of Financial Economics* 115.1, pp. 1–24. ISSN: 0304405X. DOI: 10.1016/j.jfineco.2014.08.007. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0304405X14001822> (visited on 11/15/2023).
- Barberis, Nicholas, Andrei Shleifer, and Robert Vishny (1998). “A model of investor sentiment”. In: *Journal of financial economics* 49.3, pp. 307–343.
- Bauer, Michael D, Thomas M Mertens, et al. (2019). “Zero lower bound risk according to option prices”. In: *FRBSF Economic Letter* 2019, p. 24.
- Bekaert, Geert, Xiaozheng Wang, and Cedric Tille (2010). “Inflation risk and the inflation risk premium [with Discussion]”. In: *Economic Policy* 25.64. Publisher: [Center for Economic Studies, Maison des Sciences de l’Homme, Centre for Economic Policy Research, Wiley], pp. 755–806. ISSN: 0266-4658. URL: <https://www.jstor.org/stable/40929921> (visited on 04/30/2022).
- Bernanke, Ben S and Kenneth N Kuttner (2005). “What explains the stock market’s reaction to Federal Reserve policy?” In: *The Journal of finance* 60.3, pp. 1221–1257.
- Bernard, Victor L and Jacob K Thomas (1989). “Post-earnings-announcement drift: delayed price response or risk premium?” In: *Journal of Accounting research* 27, pp. 1–36.
- Bhamra, Harjoat Singh et al. (2021). “High Inflation: Low Default Risk and Low Equity Valuations”. In: *Chicago Booth Research Paper* 18-12.
- Blank, Michael, Spencer Kwon, and Johnny Tang (2023). “Investor composition and overreaction”. In.
- Bodie, Zvi (1976). “Common stocks as a hedge against inflation”. In: *The journal of finance* 31.2, pp. 459–470.

- Bok, Brandyn, Thomas M Mertens, and John C Williams (2022). “Macroeconomic Drivers and the Pricing of Uncertainty, Inflation, and Bonds”. In: *FRB of New York Staff Report* 1011.
- Boons, Martijn et al. (2020). “Time-varying inflation risk and stock returns”. In: *Journal of Financial Economics* 136.2, pp. 444–470.
- Campbell, John Y and John Ammer (1993). “What moves the stock and bond markets? A variance decomposition for long-term asset returns”. In: *The journal of finance* 48.1, pp. 3–37.
- Campbell, John Y, Carolin Pflueger, and Luis M Viceira (2020). “Macroeconomic drivers of bond and equity risks”. In: *Journal of Political Economy* 128.8, pp. 3148–3185.
- Campbell, John Y and Robert J Shiller (1988). “The dividend-price ratio and expectations of future dividends and discount factors”. In: *The Review of Financial Studies* 1.3, pp. 195–228.
- Campbell, John Y, Robert J Shiller, and Luis M Viceira (2009). *Understanding inflation-indexed bond markets*. Tech. rep. National Bureau of Economic Research.
- Campbell, John Y., Adi Sunderam, and Luis M. Viceira (2017). “Inflation Bets or Deflation Hedges? The Changing Risks of Nominal Bonds”. In: *Critical Finance Review* 6, pp. 263–301.
- Chan, Louis K. C., Narasimhan Jegadeesh, and Josef Lakonishok (Dec. 1996). “Momentum Strategies”. en. In: *The Journal of Finance* 51.5, pp. 1681–1713. ISSN: 0022-1082, 1540-6261. DOI: 10.1111/j.1540-6261.1996.tb05222.x. URL: <https://onlinelibrary.wiley.com/doi/10.1111/j.1540-6261.1996.tb05222.x> (visited on 05/17/2024).
- Chen, Andrew Y. and Tom Zimmermann (2022). “Open Source Cross-Sectional Asset Pricing”. In: *Critical Finance Review* 27.2, pp. 207–264.
- Chen, Nai-Fu, Richard Roll, and Stephen A. Ross (1986). “Economic Forces and the Stock Market”. In: *The Journal of Business* 59.3. Publisher: University of Chicago Press, pp. 383–403. ISSN: 0021-9398. URL: <http://www.jstor.org/stable/2352710> (visited on 04/29/2022).
- Chinco, Alex, Samuel M Hartzmark, and Abigail B Sussman (2021). “A new test of risk factor relevance”. In: *The Journal of Finance*.
- Choi, James J and Thomas M Mertens (2019). “Extrapolative expectations and the equity premium”. In: *Available at SSRN 3462056*.
- Cutler, David M, James M Poterba, and Lawrence H Summers (1990). *Speculative dynamics and the role of feedback traders*.
- Dahlquist, Magnus and Markus Ibert (July 2023). *Equity Return Expectations and Portfolios: Evidence from Large Asset Managers*. en. SSRN Scholarly Paper. Rochester, NY. DOI:

- 10.2139/ssrn.3763796. URL: <https://papers.ssrn.com/abstract=3763796> (visited on 11/15/2023).
- David, Alexander and Pietro Veronesi (2013). “What ties return volatilities to price valuations and fundamentals?” In: *Journal of Political Economy* 121.4, pp. 682–746.
- De La O, Ricardo and Sean Myers (2021). “Subjective cash flow and discount rate expectations”. In: *The Journal of Finance* 76.3, pp. 1339–1387.
- Duffee, Gregory R (2023). “Macroeconomic News and Stock–Bond Comovement”. In: *Review of Finance* 27.5, pp. 1859–1882.
- Fama, Eugene (1981). “Stock Returns, Real Activity, Inflation, and Money”. In: *The American Economic Review* 71.4. Publisher: American Economic Association, pp. 545–565. ISSN: 0002-8282. URL: <http://www.jstor.org/stable/1806180> (visited on 05/08/2022).
- Fama, Eugene and G. William Schwert (Nov. 1977). “Asset returns and inflation”. en. In: *Journal of Financial Economics* 5.2, pp. 115–146. ISSN: 0304-405X. DOI: 10.1016/0304-405X(77)90014-9. URL: <https://www.sciencedirect.com/science/article/pii/0304405X77900149> (visited on 04/30/2022).
- Fama, Eugene F and Robert R Bliss (1987). “The information in long-maturity forward rates”. In: *The American Economic Review*, pp. 680–692.
- Fuster, Andreas, Benjamin Hebert, and David Laibson (2012). “Natural expectations, macroeconomic dynamics, and asset pricing”. In: *NBER Macroeconomics Annual* 26.1, pp. 1–48.
- Gabaix, Xavier and Ralph S. J. Koijen (May 2022). *In Search of the Origins of Financial Fluctuations: The Inelastic Markets Hypothesis*. en. SSRN Scholarly Paper. Rochester, NY. DOI: 10.2139/ssrn.3686935. URL: <https://papers.ssrn.com/abstract=3686935> (visited on 11/15/2023).
- Gil de Rubio Cruz, Antonio et al. (2022). “Inflation surprises in the cross-section of equity returns”. In: *Available at SSRN 4280699*.
- Gormsen, Niels Joachim and Ralph Koijen (2020). “Coronavirus: Impact on stock prices and growth expectations”. In: *The Review of Asset Pricing Studies* 10.4, pp. 574–597.
- Gourio, François and Phuong Ngo (2020). “Risk premia at the ZLB: a macroeconomic interpretation”. In.
- Greenwood, Robin and Andrei Shleifer (Mar. 2014). “Expectations of Returns and Expected Returns”. en. In: *Review of Financial Studies* 27.3, pp. 714–746. ISSN: 0893-9454, 1465-7368. DOI: 10.1093/rfs/hht082. URL: <https://academic.oup.com/rfs/article-lookup/doi/10.1093/rfs/hht082> (visited on 11/05/2023).
- Gürkaynak, Refet S, Brian Sack, and Jonathan H Wright (2010). “The TIPS yield curve and inflation compensation”. In: *American Economic Journal: Macroeconomics* 2.1, pp. 70–92.

- Hirshleifer, David, Jun Li, and Jianfeng Yu (2015). "Asset pricing in production economies with extrapolative expectations". In: *Journal of Monetary Economics* 76, pp. 87–106.
- Hong, Harrison and Jeremy C Stein (1999). "A unified theory of underreaction, momentum trading, and overreaction in asset markets". In: *The Journal of finance* 54.6, pp. 2143–2184.
- Huebner, Paul (2023). "The Making of Momentum: A Demand-System Perspective". en. In: *SSRN Electronic Journal*. ISSN: 1556-5068. DOI: 10.2139/ssrn.4395945. URL: <https://www.ssrn.com/abstract=4395945> (visited on 11/16/2023).
- Jaffe, Jeffrey F and Gershon Mandelker (1976). "The" Fisher effect" for risky assets: An empirical investigation". In: *The Journal of finance* 31.2, pp. 447–458.
- Kim, Don H and Jonathan H Wright (2005). "An arbitrage-free three-factor term structure model and the recent behavior of long-term yields and distant-horizon forward rates". In.
- Koijen, Ralph and Motohiro Yogo (Aug. 2019). "A Demand System Approach to Asset Pricing". en. In: *journal of political economy* 127.4. DOI: <https://doi.org/10.1086/701683>.
- Koijen, Ralph S. J., Robert Richmond, and Motohiro Yogo (Aug. 2023). "Which Investors Matter for Equity Valuations and Expected Returns?" en. In: DOI: 10.2139/ssrn.3378340. URL: <https://papers.ssrn.com/abstract=3378340> (visited on 11/15/2023).
- Laarits, Toomas and Marco Sammon (2022). "The Retail Habitat". en. In: *SSRN Electronic Journal*. ISSN: 1556-5068. DOI: 10.2139/ssrn.4262861. URL: <https://www.ssrn.com/abstract=4262861> (visited on 12/04/2023).
- Lintner, John (1975). "Inflation and security returns". In: *The journal of finance* 30.2, pp. 259–280.
- Livnat, Joshua and Richard R Mendenhall (2006). "Comparing the post-earnings announcement drift for surprises calculated from analyst and time series forecasts". In: *Journal of accounting research* 44.1, pp. 177–205.
- MacKinnon, David P et al. (2002). "A comparison of methods to test mediation and other intervening variable effects." In: *Psychological methods* 7.1, p. 83.
- McClure, Charles and Valeri V Nikolaev (2023). "Demand for stocks and accounting information". In: *Chicago Booth Research Paper* 22-06.
- Modigliani, Franco and Richard A Cohn (1979). "Inflation, rational valuation and the market". In: *Financial Analysts Journal* 35.2, pp. 24–44.
- Nagel, Stefan and Zhengyang Xu (2022). *Dynamics of subjective risk premia*. Tech. rep. National Bureau of Economic Research.
- Nagel, Stefan and Zhen Yan (2022). *Inflation Hedging on Main Street? Evidence from Retail TIPS Fund Flows*. Tech. rep. National Bureau of Economic Research.

- Nelson, Charles R (1976). “Inflation and rates of return on common stocks”. In: *The journal of Finance* 31.2, pp. 471–483.
- Novy-Marx, Robert (Feb. 2015). *Fundamentally, Momentum is Fundamental Momentum*. en. SSRN Scholarly Paper. Rochester, NY. URL: <https://papers.ssrn.com/abstract=2572143> (visited on 04/17/2024).
- Roussanov, Nikolai, Yang Liu, and Xiang Fang (2021). “Getting to the Core: Inflation Risks Within and Across Asset Classes”. In: *Jacobs Levy Equity Management Center for Quantitative Financial Research Paper*.
- Scott, William and William Scott (2015). *Financial accounting theory*. en. Seventh edition. Always learning. Toronto: Pearson. ISBN: 978-0-13-298466-9.
- Skinner, Douglas J and Richard G Sloan (2002). “Earnings surprises, growth expectations, and stock returns or don’t let an earnings torpedo sink your portfolio”. In: *Review of accounting studies* 7.2, pp. 289–312.
- Van Der Beck, Philippe (2022). “Identifying Elasticities in Demand-Based Asset Pricing”. en. In: *SSRN Electronic Journal*. ISSN: 1556-5068. DOI: 10.2139/ssrn.4111329. URL: <https://www.ssrn.com/abstract=4111329> (visited on 11/16/2023).
- Viceira, Luis M (2013). “Inflation-indexed bonds”. In: *NBER Reporter* 3, pp. 16–22.
- Weber, Michael (2015). “Nominal rigidities and asset pricing”. In: *Available at SSRN 2478500*.