Expected EPS × Trailing P/E[*]

Itzhak Ben-David[†] and Alex Chinco[‡]

October 23, 2024
[Click here for the latest version]

Abstract

Sell-side analysts describe how they price their own subjective beliefs in the text of each earnings report. We read a sample of 513 reports and find most analysts do not use a discount rate. They multiply a company's expected EPS (earnings per share) times a trailing P/E (price-to-earnings ratio). Trailing twelve-month P/Es explain 91% of the price-target variation in the broader IBES data. This largely backward-looking approach is problematic for the current research paradigm even if analysts are not the marginal investor. We build a simple model that rationalizes this practice, and we show it predicts market reactions to earnings surprises.

Keywords: Earnings Per Share (EPS), Price-To-Earnings Ratio (P/E), Price Target, Sell-Side Analysts, Present-Value Logic

[*]We would like to thank Xavier Gabaix, Sinan Gokkaya, Valentin Haddad, Sam Hartzmark, Pedro Matos, Jeff Meli, Stefan Nagel, Marco Sammon, Amir Sufi, and Jeff Wurgler for helpful comments. This paper has benefited from feedback we received at the NBER SI Asset-Pricing meeting, Indiana University, and Colorado University.

[†]The Ohio State University and NBER. ben-david.1@osu.edu

[‡]Baruch College, City University of New York. alexchinco@gmail.com

Introduction

At the moment, "asset-pricing theory all stems from one simple concept: price equals expected discounted payoff. The rest is elaboration, special cases, and a closet full of tricks. (Cochrane, 2009, page 1)" Every standard model prices a company's stock using some version of the formula below

$$\text{Price}_t = \frac{\mathbb{E}_t[\text{Dividend}_{t+1}] + \mathbb{E}_t[\text{Price}_{t+1}]}{1+r} \tag{1}$$

$\mathbb{E}_t[\text{Dividend}_{t+1}] + \mathbb{E}_t[\text{Price}_{t+1}]$ denotes investors' expected payoff from owning a share, and r is the rate at which they discount this future cash flow.

Much of what we think we know about discount rates comes from studying the earnings forecasts and price targets found in sell-side research. Most market participants do not publicly announce their subjective payoff expectations. Sell-side analysts do. As a result, analysts' forecasts have had an outsized impact on the asset-pricing literature (Kothari, So, and Verdi, 2016).

Figure 1 shows a December 2019 report about Home Depot written by Chris Horvers, a senior analyst at JP Morgan. Chris Horvers starts his report by telling investors to "Overweight" Home Depot in their portfolios—i.e., buy more shares. Then, he sets a price target of $\mathbb{E}_t[\text{Price}_{t+1}] = \$241/\text{sh}$ for Home Depot in December 2020 (one year into the future) based on his view that the company would have earnings per share (EPS) of $\mathbb{E}_t[\text{EPS}_{t+1}] = \$10.48/\text{sh}$ over the next twelve months (FY2020) and $\mathbb{E}_t[\text{EPS}_{t+2}] = \$11.50/\text{sh}$ the year after (FY2021).

The Institutional Brokers' Estimate System (IBES) has tabulated these numbers into an easy-to-use format, and our profession has spent decades pouring over this data to learn about discount rates. It is common to see papers replace the numerator in Equation (1) with IBES data on $\mathbb{E}_t[\text{EPS}_{t+1}]$ and $\mathbb{E}_t[\text{Price}_{t+1}]$ and solve for the implied r (Gebhardt, Lee, and Swaminathan, 2001). Researchers assess the effects of time-varying discount rates by plugging IBES data into Campbell and Shiller (1988a), a multiperiod approximation to Equation (1).

But there is more to sell-side research than the numbers found in IBES. Analysts explicitly state their pricing rule in the text of each report. We read

J.P.Morgan

North America
12 December 2019

The Home Depot

Analyst Day: Key Messages, Model Reads, Management Follow-Up Takeaways

Overall, we remain impressed with HD's culture, willingness to play long ball, and its plans to imp-rove execution and drive market share in…

Overweight

HD, HD US

Price (11 Dec 19): $212

Price Target (Dec-20): $241

Prior (Dec-20): $252

Retailing/Broadlines

Christopher Horvers, CFA

(a) Top of first page

Key Metrics (FYE Jan)

$ in millions	**FY19E**	**FY20E**	**FY21E**
Financial Estimates			
Revenue	110,118	114,352	118,452
Adj. EBITDA	18,034	18,494	19,266
Adj. EBIT	15,763	16,174	17,191
Adj. net income	11,022	11,261	11,987
Adj. EPS	10.05	10.48	11.50
Valuation			
FCFF yield	5.0%	5.1%	6.0%
Dividend yield	2.6%	2.6%	2.7%
EV/Revenue	2.4	2.3	2.2
EV/EBITDA	14.6	14.3	13.8
Adj. P/E	21.1	20.2	18.4

(b) Key Metrics

Figure 1. *Earning report about Home Depot, which was published on December 12th 2019 by JP Morgan. The lead analyst on this report was Chris Horvers.*

what analysts write and find that most do not use a discount rate. Instead, they typically value large public companies by multiplying their short-term EPS forecast times a trailing P/E (price-to-earnings ratio).

A trailing P/E is not a handy heuristic for doing present-value calculations. It is backward-looking! Instead of valuing a company by discounting its entire expected future earnings stream, analysts usually ask themselves: "How has the market generally priced each dollar of the company's earnings *in the past*?" Sure, both valuation methods may sometimes lead to similar prices in certain situations. But the underlying economics is entirely different.

Our findings have massive implications for asset-pricing research even if sell-side analysts are not the marginal investor. It could be that other market participants may do things differently. But researchers have not spent the past 40+ years analyzing data on those other market participants' subjective beliefs. The numerical values in IBES tell us nothing about discount rates if the people responsible for these numbers are not discounting anything.

More generally, it does not make sense for all of asset-pricing theory to revolve around discount rates if we cannot count on market participants actually using one. After documenting that analysts do not typically apply present-value logic, we spend the remainder of the paper building an alternative path forward. We outline a simple trailing P/E model and show it can explain important real-world patterns, such as how prices respond to earnings surprises.

Investment Thesis, Valuation and Risks

The Home Depot, Inc. *(Overweight; Price Target: $241.00)*

Valuation

Our Dec 2020 price target is $241 (down from $252 prior), which is based on ~21.0x our revised 2021E EPS, in line with its three-year average.

Valuation Matrix

	2018	2019E	2020E	2021E
EPS	$9.89	$10.05	$10.48	$11.50
PE	21.4x	21.1x	20.2x	18.4x
Three Year Avg			21.7x	19.0x
Three Year Peak			24.7x	21.2x
Historic Relative PE			1.2x	1.2x
Relative Five Year PE Peak			1.4x	1.3x
			$241.00	
PE	24.4x	24.0x	23.0x	21.0x
EV/EBITDA	16.8x	16.1x	15.5x	14.6x
Upside/Downside			14%	

Figure 2. *How Chris Horvers described calculating his* $241/sh *price target for Home Depot in his December 2019 earnings report for JP Morgan.*

We begin in Section 1 by analyzing the text of 513 sell-side analyst reports about large publicly traded companies from 2003 through 2022. These reports are more than just dry colorless lists of numbers. Analysts explain how they price their own subjective cash-flow expectations. They do not typically use a discount rate. Instead, analysts usually set price targets, $\text{PriceTarget}_t \overset{\text{def}}{=} \mathbb{E}_t[\text{Price}_{t+1}]$, by multiplying their EPS forecast times a trailing P/E ratio.

FINRA Rule 2241 requires "any recommendation, rating, or price target [to be] accompanied by a clear explanation of any valuation method used." Figure 2 shows how Chris Horvers described his $241/sh price target for Home Depot as "~21.0× our revised 2021E EPS, in line with its three-year average"

$$\underbrace{\mathbb{E}_t[\text{Price}_{t+1}]}_{\$241/\text{sh}} = \underbrace{\mathbb{E}_t[\text{EPS}_{t+2}]}_{\$11.50/\text{sh}} \times \underbrace{\left(\frac{1}{3} \cdot \sum_{\ell=0}^{2} \text{TrailingPE}_{t-\ell} \right)}_{21.0} \tag{2}$$

To predict Home Depot's price in December 2020 (end of period $t+1$), he multiplied his EPS forecast for FY2021 (period $t+2$) times a 3-year trailing average P/E based on FY2017, FY2018, and FY2019 (periods $t-2$, $t-1$, and t).

Chris Horvers could have described carefully modeling the company's cash flows from 2022 onward. He could have gone deep into the weeds, outlining precisely how he discounted these expected payoffs. He was more than capable of doing this sort of analysis. He chose not to.

We appreciate that the classic Gordon model says to price stocks with a forward-looking multiple. This model comes from iterating Equation (1) forward and assuming constant dividend growth

$$\text{Price}_t = \frac{\mathbb{E}_t[\text{Dividend}_{t+1}] + \mathbb{E}_t[\text{Price}_{t+1}]}{1+r} \tag{1}$$

`for` $i = (t+1)$ `to` ∞:
`replace` Price_i
`with` $\frac{\mathbb{E}_i[\text{Dividend}_{i+1}]+\mathbb{E}_i[\text{Price}_{i+1}]}{1+r}$

$$= \sum_{h=1}^{\infty} \frac{\mathbb{E}_t[\text{Dividend}_{t+h}]}{(1+r)^h} \tag{3}$$

Assume: $(1+g)^h = \frac{\mathbb{E}_t[\text{Dividend}_{t+h}]}{\text{Dividend}_t}$

$$= \mathbb{E}_t[\text{Dividend}_{t+1}] \times \left(\frac{1}{r-g}\right) \tag{4}$$

The idea is to use $\left(\frac{1}{r-g}\right) = \left(\sum_{h=1}^{\infty} \frac{\mathbb{E}_t[\text{Dividend}_{t+h}]}{(1+r)^h}\right) / \mathbb{E}_t[\text{Dividend}_{t+1}]$ as a shortcut for calculating the present value of an infinite stream of expected future dividends.

Chris Horvers' 21.0× is not $\left(\frac{1}{r-g}\right)$ in disguise. There is nothing forward-looking about his "Valuation Matrix" in Figure 2. He did not scale up his $\mathbb{E}_t[\text{EPS}_{t+2}] = \$11.50$/sh forecast by a factor of 21.0× to capture the present discounted value of Home Depot's earnings stream from 2022 onward. He chose a value of 21.0× to ensure his price target would be "in line" with how the market had priced each $1 of Home Depot's earnings in the recent past.

21.0× was not Chris Horvers' best guess about Home Depot's future P/E ratio, either. He gave the company an "Overweight" rating in big bold letters at the top of the first page of his report (Figure 1). He felt the market had been undervaluing Home Depot. Chris Horvers said in his report he expected the firm to experience "multiples expansion" going forward.

None of this makes Chris Horvers a bad analyst. Chris Horvers is an excellent analyst. He has been named to *Institutional Investor* magazine's All-America research team multiple times, and his December 2019 report about Home Depot is the kind of report that other analysts are striving to produce. We greatly admire Chris Horvers' work. We think other researchers should read it.

Analysts do not price stocks in the way we, as researchers, have been trained to believe. In a sense, it is surprising that so many papers have been written about biased EPS forecasts. Analysts put a lot of effort into getting those numbers right. Chris Horvers spent pages justifying $\mathbb{E}_t[\text{EPS}_{t+2}] = \$11.50/\text{sh}$. Then, when it came time to capitalize this subjective cash-flow expectation, he did not even attempt to apply the "one simple concept" at the heart of theoretical asset pricing. This seems like the more noteworthy point of departure.

Researchers are free to continue pretending that price always equals expected discounted payoff if they want to. You can still use IBES data to estimate implied costs of capital and perform Campbell and Shiller (1988a) decompositions like nothing has changed. But what would be the point? If the imputed values have nothing to do with how assets are actually being priced, then they will not be relevant in the future or helpful for evaluating policy counterfactuals.

In Section 2, we build a simple model that reflects how analysts say they price stocks. Analysts in our model set one-year-ahead price targets for a single stock by multiplying their short-term EPS forecast times a trailing P/E. Investors adjust their demand based on the relative difference between analysts' price target and the current price. When the price target is 1%pt higher than the current price, investors increase their holdings by μ%pt over the next year. The stock's subsequent price growth also adjusts proportionally, increasing by ν%pt on average each time that investor demand goes up by 1%pt.

We use this model to give conditions under which expected EPS times trailing P/E will be correct on average. The key insight is that the equilibrium price in our model is mostly backward-looking. $\mathbb{E}[\text{EPS}]$ is the only forward-looking input. As a result, it can make sense to use a trailing P/E to set price targets because prices themselves are mostly backward-looking.

Behavioral researchers typically study one-step deviations from a well-known rational benchmark (Rabin, 2013). Changing more than one or two inputs usually makes the resulting analysis intractable, and using a trailing P/E is much more than a step-or-two departure from "price equals expected discounted payoff". It could be that the resulting price dynamics are incomprehensible. Our model shows that this is not the case.

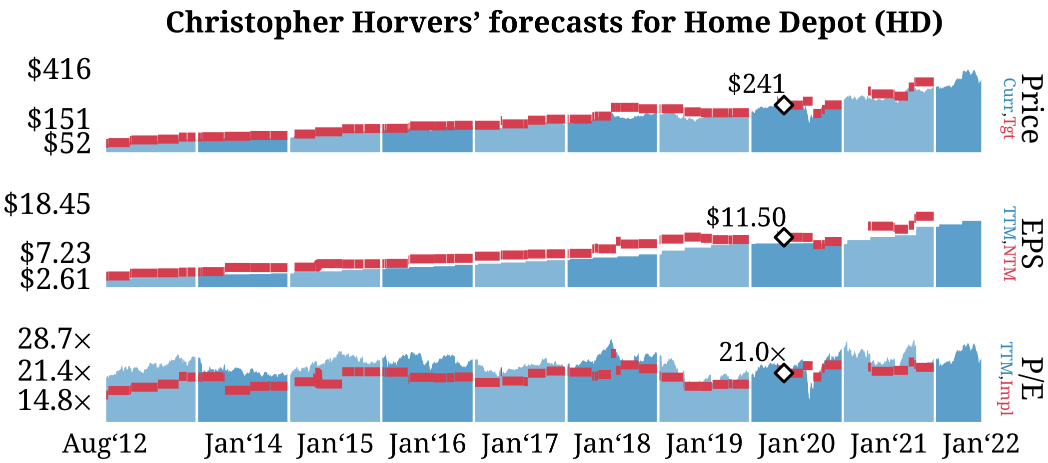

Figure 3. *y-axis labels correspond to the min, median, and max in each panel. (Top Panel) Blue ribbon is Home Depot's closing price on day t in CRSP,* Price_t*. Red line is Chris Horvers' price target* PriceTarget_t *in IBES. (Middle Panel) Blue is the sum of HD's quarterly EPS in IBES over four quarters prior to day t,* EPS_t*. Red is Chris Horvers' EPS forecast for the year following his target date (NTM),* $\mathbb{E}_t[\text{EPS}_{\tau+2}]$*. (Bottom Panel) Blue is HD's TTM P/E ratio,* $\text{TrailingPE}_t = \text{Price}_t \,/\, \text{EPS}_t$*. Red is P/E implied by Chris Horvers' forecasts,* $\text{ImpliedPE}_t = \text{PriceTarget}_t \,/\, \mathbb{E}_t[\text{EPS}_{\tau+2}]$.

The fact that $\mathbb{E}[\text{EPS}]$ is the only forward-looking input to prices also represents a sharp testable restriction. It implies that a piece of news can only affect a firm's return by shifting analysts' price target if it causes them to revise their short-term EPS forecast. Because the firm's trailing P/E is set in stone, the entire effect must operate through this one narrow channel. Whereas, in a Gordon model, news about a firm's long-run EPS growth rate, g, can also predict returns.

Finally, in Section 3, we present two sets of empirical results, each with a different purpose. First, we predict analysts' price targets in IBES to show that our conclusions from reading 513 reports extend to this broader sample. Figure 3 illustrates the general approach. The top panel shows a daily time series of Chris Horvers' price targets as reported in IBES, $\text{PriceTarget}_t \overset{\text{def}}{=} \mathbb{E}_t[\text{Price}_{\tau+1}]$, where $(\tau+1)$ denotes Home Depot's next fiscal year-end date. The bottom panel shows a time series of the P/E ratios implied by Chris Horvers' EPS forecast, $\text{ImpliedPE}_t \overset{\text{def}}{=} \text{PriceTarget}_t \,/\, \mathbb{E}_t[\text{EPS}_{\tau+2}]$. This implied P/E closely tracks Home Depot's trailing twelve-month (TTM) P/E, $\text{TrailingPE}_t \overset{\text{def}}{=} \text{Price}_t \,/\, \text{EPS}_t$.

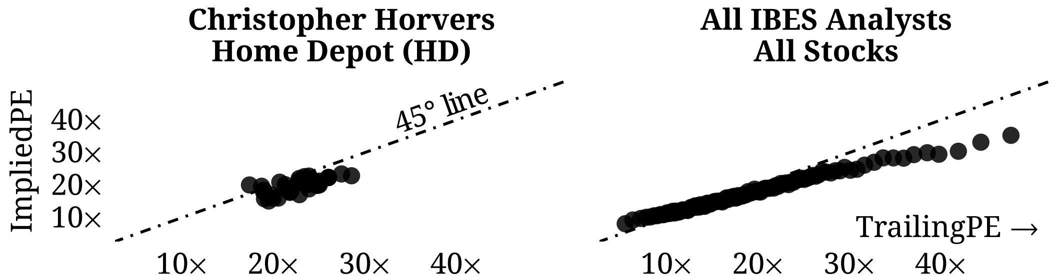

Figure 4. *(Left) Each dot denotes a day on which Chris Horvers updated his price target for Home Depot. x-axis is Home Depot's trailing twelve-month P/E,* $\text{TrailingPE}_t = \text{Price}_t \,/\, \text{EPS}_t$. *y-axis is the P/E ratio implied by Chris Horvers' forecast values,* $\text{ImpliedPE}_t = \text{PriceTarget}_t \,/\, \mathbb{E}_t[\text{EPS}_{\tau+2}]$. *(Right) Binned scatterplot of days on which any IBES analyst updated their price target for any firm.*

Expected EPS times a trailing twelve-month P/E explains $R^2 = 91\%$ of the price-target variation in IBES. However, we know that analysts deploy many variations on this common theme. For instance, Chris Horvers' 21.0× trailing P/E was a trailing three-year average. Every dot in the right panel of Figure 4 does not sit perfectly on the 45° line. But, because our trailing twelve-month calculation captures the essence of what most analysts say they are doing, it predicts their price targets far better than any other model we know of.

We also show that trailing P/Es explain equilibrium outcomes, not just analysts' beliefs. Following an earnings surprise, analysts will revise their EPS forecast for the upcoming year. In our model, they will continue to capitalize this new value of $\mathbb{E}[\text{EPS}]$ using the same trailing P/E. So, among firms with the same size earnings surprise, the subsequent price change should be determined by the firm's trailing P/E at the time of its earnings surprise.

We test this prediction using an approach similar to Fama and MacBeth (1973). First, we group stock-quarter observations into portfolios by the size of their earnings surprise. Then, within each group, we run a separate first-stage regression to estimate the relationship between a stock's trailing P/E and its subsequent price change. Our model says these first-stage estimates should be proportional to the size of the associated earnings surprise. When we evaluate this claim by running a second-stage regression, we find a neat linear fit.

Related Work. This paper is an asset-pricing analog to Ben-David and Chinco (2024). In that paper, we took managers at their word when they said they were EPS maximizers and fleshed out the implications for corporate policies. In this paper, we take sell-side analysts at their word when they say they use trailing P/E ratios and derive the implications for asset prices.

There are numerous papers studying the accuracy of multiples analysis for pricing public equities (Bhojraj and Lee, 2002; Liu, Nissim, and Thomas, 2002; Brav and Lehavy, 2003; Da and Schaumburg, 2011; Bartram and Grinblatt, 2018; Mukhlynina and Nyborg, 2020; Cooper and Lambertides, 2023), IPOs (Kim and Ritter, 1999; Purnanandam and Swaminathan, 2004), and syndicated loan deals (Murfin and Pratt, 2019). We point out that, no matter how accurate they are, a trailing multiple is problematic for the existing research paradigm.

These papers on multiples analysis point to a different connection between asset prices and accounting data. This broad research program has a long history (Basu, 1983; Campbell and Shiller, 1988b; Lamont, 1998; Lewellen, 2004; Kothari, Lewellen, and Warner, 2006; Cready and Gurun, 2010).

Our paper connects to the broader literature on belief formation (Malmendier and Nagel, 2011; Greenwood and Shleifer, 2014; Coibion and Gorodnichenko, 2015; Adam, Marcet, and Beutel, 2017; Bordalo, Gennaioli, Ma, and Shleifer, 2020; Giglio, Maggiori, Stroebel, and Utkus, 2021; Adam, Matveev, and Nagel, 2021; Afrouzi, Kwon, Landier, Ma, and Thesmar, 2023).

There is also a substantial amount of evidence that analysts suffer from predictable biases when making forecasts and capitalizing them into prices (La Porta, 1996; So, 2013; Bouchaud, Krueger, Landier, and Thesmar, 2019; Bordalo, Gennaioli, La Porta, and Shleifer, 2019, 2020, 2024; De la O and Myers, 2021; Charles, Frydman, and Kilic, 2024).

Finally, this paper provides evidence against the discount-rate approach to asset pricing (Cochrane, 2011), which argues that a company's current share price reflects investors' desire to insure themselves against exposure to specific kinds of future aggregate risks. It is hard to find people who think this way in the real world (Chinco, Hartzmark, and Sussman, 2022). We build a simple alternative model that matches what one group of investors actually does.

1 In Their Own Words

The goal of asset-pricing research is to figure out how market participants would price any arbitrary set of risky cash flows. However, in the particular case of sell-side analysts, there is nothing to figure out. For the past two decades, analysts have been legally required to write down their pricing rule in the text of each earnings report. In May 2002, the SEC passed NASD Rule 2711 stating that: "If a research report contains a price target, the [analyst] must disclose in the research report the valuation methods used to determine the price target." In 2015, this rule was superseded by FINRA Rule 2241, which also requires that a "price target [to be] accompanied by a clear explanation." In this section, we examine how analysts explain their own pricing rule and find that most do not use a discount rate. Instead, they rely on trailing P/E ratios.

1.1 Data description

We downloaded 513 earnings reports from Investext in two separate waves. We started with 339 reports written about the 30 largest publicly traded companies at year-end in 2004, 2011, and 2019. This gives us 47 companies in total (see Table 1). For each company in a given year, we include one report written by each brokerage in Table 2.

Based on this first sample, it does not look like many sell-side analysts apply present-value reasoning. However, these are run-of-the-mill reports written by average analysts. Perhaps the best analysts set price equal to expected discounted payoff when writing reports that really matter?

To check whether this is the case, we then downloaded an additional 174 coverage-initiation reports written by 28 sell-side analysts who have been repeatedly named to *Institutional Investor* magazine's All-America team. These analysts are the best of the best (Stickel, 1992), and analysts put a disproportionate amount of effort into coverage-initiation reports (McNichols and O'Brien, 1997), often laying out a general theory for pricing the firm. The average coverage-initiation report in our sample runs 29 pages. 20% are 40+ pages long.

Number of reports about each company (Sample #1)

		2004	2011	2019	Total
1	Abbott Labs	3	4	4	11
2	Adobe			6	6
3	AIG	3			3
4	Altria	3			3
5	Amazon		3	7	10
6	American Express	3			3
7	Amgen	4			4
8	Apple		5	7	12
9	AT&T		3	2	5
10	Bank of America	3		6	9
11	Boeing			5	5
12	Chevron	3	3	7	13
13	Cisco	3	4	6	13
14	Citigroup	2	4	5	11
15	Coca-Cola	3	2	4	9
16	ConocoPhillips		1		1
17	Dell	4			4
18	Disney			3	3
19	eBay	4			4
20	Exxon Mobil	3	2	7	12
21	Facebook			6	6
22	GE	3	3		6
23	Google		4	7	11
24	Home Depot	4		6	10
25	IBM	4	4		8
26	Intel	3	3	5	11
27	Johnson & Johnson	3	3	1	7
28	JP Morgan	2	2	4	8
29	Mastercard			7	7
30	McDonalds		4		4
31	Merck	2	3	3	8
32	Microsoft	4	4	6	14
33	Occidental		3		3
34	Oracle	3	4	6	13
35	Pepsi	3	1	5	9
36	Pfizer	3	4	5	12
37	Philip Morris		2		2
38	Procter & Gamble	3	3		6
39	Qualcomm		4		4
40	Schlumberger		2		2
41	Time Warner	3			3
42	UBS	1			1
43	UnitedHealth			6	6
44	Verizon	3	3	5	11
45	Visa			7	7
46	Walmart	3	3	6	12
47	Wells Fargo	3	3	1	7
	Total	91	93	155	339

Table 1. *Our first sample of documents contains* 339 *sell-side reports written about the largest* 30 *publicly traded companies in 2004, 2011, and 2019.*

Number of reports from each brokerage (Sample #1)

		2004	2011	2019	Total
1	Argus Research	28	30	26	84
2	Cowen and Co	8	14	22	44
3	Credit Suisse	27	25	24	76
4	JP Morgan	28	21	26	75
5	Société Générale		3	8	11
6	Wedbush Securities			10	10
7	Wells Fargo			23	23
8	Wolfe Research			16	16
	Total	91	93	155	339

Table 2. *Our first sample of documents contains* 339 *sell-side reports written by analysts at* 8 *different brokerages.*

Institutional Investor publishes their rankings in October. We read through these issues and recorded which analysts made the All-America team each year. The magazine ranks analysts by GICS sector. For each sector, we identified the 10 analysts with the most years on the All-America team. The 174 documents in our second wave come from All-American analysts on this top-10 list.

Analysts write coverage-initiation reports either when a company is new or when they join a new brokerage. 53 of our 174 coverage-initiation reports (30.5% of the sample) involve companies that went public within the previous three years. If anything, analysts should be more likely to use forward-looking information in this sample because there is less trailing information to go on. Many public firms also initially start out with negative earnings, making it difficult for analysts to apply the formula $\text{PriceTarget} = \mathbb{E}[\text{EPS}] \times \text{TrailingPE}$.

We only include reports written by analysts that can be matched to both IBES and Investext. This is a meaningful restriction. For example, IBES does not include data on Ed Hyman, head of Evercore ISI's research team and the single most capped analyst on *Institutional Investor* magazine's All-America team.

To check the quality of our data, we downloaded all earnings reports in Investext for a subset of analyst-firm pairs. As you can see from Figure 6, the price targets and EPS forecasts in the PDFs perfectly match up with the numbers

Number of reports by each All-American analyst (Sample #2)

		# Reports	Sector
1	Meredith Adler	2	Consumer Discretionary
2	Greg Badishkanian	30	Consumer Discretionary
3	Jamie Baker	8	Industrials
4	Robert Cornell	1	Basic Materials
5	Philip Cusick	2	Media & Entertainment
6	Christopher Danely	3	Technology
7	Robert Drbul	4	Consumer Discretionary
8	John Faucher	3	Consumer Staples
9	Daniel Ford	3	Utilities
10	Michael Gambardella	4	Basic Materials
11	Lisa Gill	1	Health Care
12	John Glass	2	Consumer Discretionary
13	Joseph Greff	7	Consumer Discretionary
14	Tien-tsin Huang	6	Technology
15	Andy Kaplowitz	1	Industrials
16	Andrew Lazar	1	Consumer Staples
17	Greg Melich	3	Consumer Discretionary
18	CJ Muse	6	Technology
19	Joseph Nadol	2	Industrials
20	Himanshu Patel	11	Consumer Discretionary
21	Tycho Peterson	9	Health Care
22	Walter Piecyk	20	Telecommunications
23	Kash Rangan	1	Technology
24	Josh Shanker	2	Financials
25	Andrew Steinerman	4	Financials
26	Brian Tunick	26	Consumer Discretionary
27	Michael Weinstein	6	Health Care
28	Jeffrey Zekauskas	6	Basic Materials
	Total	174	

Table 3. *Our second sample of documents contains* 174 *coverage-initiation reports written by* 28 *different analysts named to Institutional Investor magazine's All-America research team.*

J.P.Morgan

North America Equity Research
18 October 2019

Coca-Cola

Another Strong Quarter; Can KO Keep Up the Pace in 2020?

Neutral

KO, KO US

Price (18 Oct 19): $54.93

Price Target (Dec-20): $59.00

(a) Top of first page

Valuation

Coke is trading at roughly 24.5x our 2020 EPS estimate, a LDD premium compared to our large-cap beverage/HPC multinational average and ~1.5x turn of multiple more expensive than OW-rated PepsiCo. Although the steps the company is taking have improved the pace of growth and profitability, we need comfort that earnings growth beyond 2019 will move closer to the company's LT algorithm of HSD growth or the high end of MSD organic revenue growth is sustainable in the coming years before considering becoming more positive. Our December 2020 price target of $59 is predicated on ~24x our 2021 estimate, broadly in line with the current multiple. While this is above the company's historical valuation, with rates moving lower and organic revenue growth still very strong, we see limited downside to the multiple in the coming months.

(b) Methods section

Figure 5. *Earning report about Coca-Cola, which was published on December 19th 2019 by JP Morgan. The lead analyst on this report was Andrea Teixeira.*

in IBES. Moreover, the P/E ratios implied by these numbers (red lines; bottom panel) line up with the ones found in analysts' reports. We only use these additional reports to ensure the accuracy of our raw numbers.

We take several steps to ensure our 513 observations are representative. In Subsection 3.2, we directly show that PriceTarget = $\mathbb{E}$[EPS] × TrailingPE can account for $R^2 = 91\%$ of the price-target variation in the broader IBES sample. We also note that one cannot fix a selection problem with more data. While Décaire and Graham (2024) and Gormsen and Huber (2024) are able to include tens of thousands of data points, both papers discard observations that do not include a discount rate. Thus, these studies both draw conclusions from a large but highly nonrepresentative sample. In Subsection 1.9, we show our findings line up with the fraction of observations excluded by each paper.

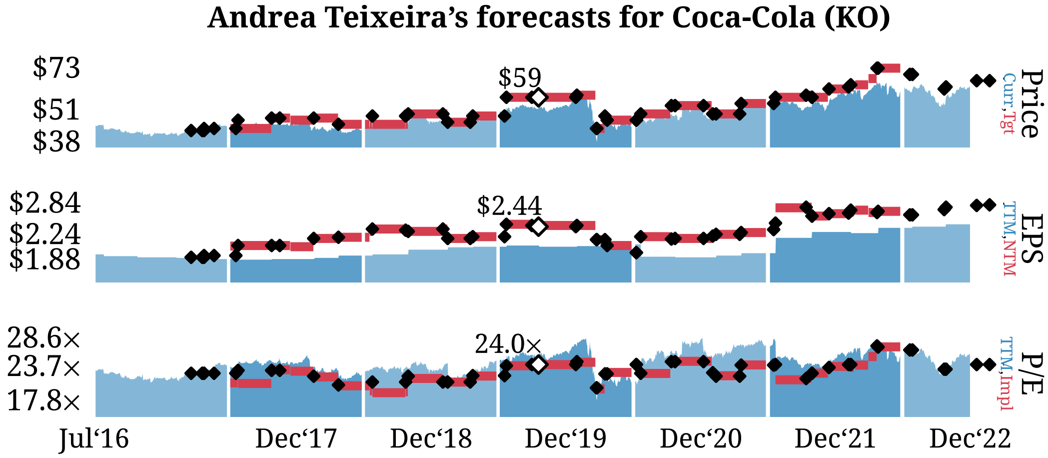

Figure 6. *y-axis shows min, median, and max. (Top) Blue ribbon is Coca-Cola's (KO)'s closing price on day t from CRSP,* Price_t. *Red line is Andrea Teixeira's price target,* $\text{PriceTarget}_t = \mathbb{E}_t[\text{Price}_{\tau+1}]$, *as reported in IBES. (Middle) Blue is KO's trailing twelve-month (TTM) earnings per share (EPS) on day t,* EPS_t, *as reported in IBES. Red is Andrea Teixeira's EPS forecast for the year following her target date,* $\mathbb{E}_t[\text{EPS}_{\tau+2}]$. *(Bottom) Blue is KO's TTM price-to-earnings (P/E) ratio,* $\text{TrailingPE}_t = \text{Price}_t \,/\, \text{EPS}_t$. *Red is the P/E implied by Ms Teixeira's forecasts,* $\text{ImpliedPE}_t = \text{PriceTarget}_t \,/\, \mathbb{E}_t[\text{EPS}_{\tau+2}]$. *White diamonds are values from the October 2019 Andrea Teixeira report about Coca-Cola shown in Figure 5. This report belongs to our* 513 *document sample. Black diamonds are values from other Andrea Teixeira reports about Coca-Cola not in our* 513 *report sample.*

1.2 Most Analysts Rely On Trailing P/E Ratios

Multiples analysis is the standard way to value large public companies. Table 4 shows that analysts used some form of multiples analysis in 94.5% of our sample (485 out of 513 reports). Price-to-earnings (P/E) was the most common multiple and was listed in the methods section 76.8% of the time.

Analysts set a price target based on a multiple of earnings before interest, taxes, depreciation, and amortization (EBITDA), cash flows (CF), or sales 43.9% of the time (225 of 513 reports). In this paper, we treat these approaches as separate methods to be as conservative as possible. But they are delevered versions of $\mathbb{E}[\text{EPS}] \times \text{TrailingPE}$. Analysts often use $\mathbb{E}[\text{EBITDA}] \times \text{TrailingEVtoEBITDA}$ in situations where a company's EPS has been negative in recent years.

Most analysts used multiples analysis to set price targets

	2004	2011	2019	All Am	Total
Any Multiple	85.7% 78	91.4% 85	96.8% 150	98.9% 172	94.5% 485
P/E ratio	79.1% 72	83.9% 78	80.0% 124	69.0% 120	76.8% 394
EBITDA, CF, Sales	27.1% 25	31.9% 30	50.6% 82	50.6% 88	43.9% 225
Book Value	7.7% 7	16.1% 15	7.7% 12	3.4% 6	7.8% 40
P/E-to-Growth	8.8% 8	9.7% 9	40.7% 18	11.6% 18	10.3% 53
Dividend Yield	8.8% 8	2.2% 2	5.2% 8	8.6% 15	6.4% 33
# Reports	91	93	155	174	513

Table 4. *"Any Multiple": report used at least one multiple to calculate the price target. "P/E Ratio": report used a firm's price-to-earnings ratio (P/E). "EBITDA, CF, Sales": report set a price target based on a multiple of EBITDA, cash flow, or sales. "Book Value": report used a multiple of the book value of a firm's assets. "P/E-to-Growth": report used the ratio of a company's P/E to its EPS growth rate. "Dividend Yield": report used a firm's dividend yield when setting a price target. Top number in each cell is the percent relative to the total for the column. e.g.,* 78 *of* 91 *reports in 2004 described using some form of multiples analysis,* 78/91 = 85.7%.

There is nothing inherently wrong with multiples analysis. It is the kind of multiple that conflicts with textbook theory. Most sell-side analysts use a trailing value. Table 5 shows analysts looked at a firm's own trailing multiple in 63.5% of our sample (326 of 513 reports). They looked at the recent pricing of the firm's peer group in 74.1% of our sample (380 reports), and they made both kinds of comparisons in over half of the reports in our sample (260 out of 513 reports; 50.7%). Sell-side analysts typically choose a multiple based on where a firm and others like it have been trading at in recent years.

The popularity of peer-group comparisons in our data is likely driven in part by the fact that coverage-initiation reports make up a third of our sample (174 of 513 reports; 33.9%). Many of these firms recently went public. In these cases, there is often not enough trailing data to compute an average.

Analysts pick multiples based on past realizations

	2004	2011	2019	All Am	Total
Own Past Pricing	50.5% 46	50.5% 47	54.8% 85	85.1% 148	63.5% 326
Pricing of Peers	69.2% 63	60.2% 56	59.4% 92	97.1% 169	74.1% 380
Both Comparisons	38.5% 35	31.2% 29	31.6% 49	84.5% 147	50.7% 260
# Reports	91	93	155	174	513

Table 5. *"Own past pricing": analyst computed a multiple that reflects a firm's own past pricing in recent years. "Pricing of peers": analyst computed a multiple that reflects the past pricing of a company's peer group. "Both comparisons": analyst made both comparisons. Top number in each cell is the percent relative to the total for the column. e.g.,* 46 *of* 91 *reports in 2004 described using a multiple based on a company's own past pricing,* $46/91 = 50.5\%$.

Analysts are fully capable of calculating a forward-looking multiple if they want to. They regularly perform calculations that are far more involved than $\mathbb{E}[\text{EPS}] \times \left(\frac{1}{r-g}\right)$. Table 6 shows that they performed "sum of the parts (SOTP)" analysis in 9.6% of our sample (49 of 513 reports). For example, Figure 7 shows an October 2019 earnings report about Amazon from Wolfe Research. The analyst who wrote this report, Chris Bottiglieri, used a different multiple to value each line of Amazon's business.

In fact, sell-side analysts set a price target based on multiple multiples in 38.8% of our sample (199 of 513 reports). For example, Figure 8 shows an April 2010 coverage-initiation report about Avis Budget written by an All-American analyst, Himanshu Patel. It is a thorough report by a high-quality analyst, and the price target has nothing to do with an expected discounted payoff.

We are not arguing that all analysts always use the exact same formula. Analysts are smart people who are capable of nuance. Our findings suggest they start with $\text{PriceTarget} = \mathbb{E}[\text{EPS}] \times \text{TrailingPE}$ and then make adjustments as needed. By contrast, all asset-pricing theory currently stems from a single assumption: price equals expected discounted payoff. Given their importance to the literature, it is a problem that most analysts take a different approach.

WOLFE RESEARCH

HARDLINES & INTERNET RETAIL

Internet Retail – Market Overweight

October 24, 2019

AMAZON.COM, INC.

(AMZN – $1780.78 – Outperform)

Trading and Fundamental Data	
Target Price YE '20	$1,918

(a) Top of first page

Investment Conclusion

AMZN shares are up 17% year to date but traded off ~7% in post-market trading after its earnings release earlier today. AMZN is underperforming the S&P 500, which is up 20% YTD. AMZN outperformed in 2018, increasing 28% vs. the S&P 500's return of -6%.

We are cutting our 2020 and 2021 estimates by 30% and 25%, respectively. Our 2020 and 2021 EBITDA estimates for AMZN are 30% and 25% below prior Consensus, which we expect to get revised downward.

We arrive at our $1,918 CY 20 price target (was $2,234) using a sum-of-the-parts valuation framework. We apply a blended 20.2x NTM EV/EBITDA multiple to our 2021E EBITDA estimate. Our 20.2x EV/EBITDA multiple uses 17.5x EV/EBITDA on North America, 1.5x EV/Sales on International, and 17.5x EV/EBITDA on AWS. A 20.2x EV/EBITDA multiple is above where shares are currently trading, but roughly in-line with AMZN's 1 and 3yr averages.

Exhibit 1: Sum of the Parts Valuation Framework

Sum-of-the-Parts Analysis			
FY21E Amazon North America EBITDA	$17,420	Amazon North America	$304,856
Times: Estimated EV/EBITDA Multiple	17.5x	Amazon International	$147,954
Enterprise Value	**$304,856**	Amazon AWS	$445,571
		Amazon Unallocated D&A	$89,694
FY21E Amazon International Sales	$98,636	**Total Enterprise Value**	**$988,074**
Times: Estimated Sales Multiple	1.5x	Less: Total Debt	$56,224
Enterprise Value	**$147,954**	Plus: Cash & Cash Equivalents	$37,465
Note: FY21E International EBITDA	$963	**Equity Value**	**$969,315**
FY21E Amazon AWS EBITDA	$25,461	Divide: Shares Outstanding	505
Times: Estimated EBITDA Multiple	17.5x		
Enterprise Value	**$445,571**	**Implied Value / Share (CYE '2**	$1,918
FY21 Amazon Unallocated EBITDA	$5,125	Consolidated EBITDA (after SB	$48,970
Times: Estimated EBITDA Multiple	17.5x	Implied Consolidated Multiple	20.2x
Enterprise Value	**$89,694**		

Source: Wolfe Research Estimates, Company Filings

(b) Methods section

Figure 7. *Earning report about Amazon, which was published on October 24th 2019 by Wolfe Research. The lead analyst on this report was Chris Bottiglieri, and he computed a different multiple to value each of Amazon's four lines of business. This represents an example of sum of the parts (SOTP) analysis.*

Analysts average the price targets implied by different methods

	2004	2011	2019	All Am	Total
Used 2+ Multiples	30.8% 28	36.6% 34	43.9% 68	39.7% 69	38.8% 199
Sum of the Parts (SOTP)	4.4% 4	5.4% 5	16.8% 26	8.0% 14	9.6% 49
# Reports	91	93	155	174	513

Table 6. *"Used 2+ Multiples": report described calculating a firm's price target using a blend of two or more multiples. "Sum of the Parts (SOTP)": report described calculating a firm's price target by taking a weighted average of industry-specific values of the same multiple with weights that reflect the importance of each line of business. Top number in each cell is the percent relative to the total for the column. e.g.,* 28 *of* 91 *reports in 2004 described using multiple multiples,* 28/91 = 30.8%.

1.3 Price Seldom Equals Expected Discounted Payoff

Table 7 shows that sell-side analysts mention a discounted cash-flow (DCF) or dividend discount model in just 30.2% of reports (155 of 513). This statistic includes any report that mentions the terms "DCF" or "Discounted Cash Flow" in the methods section. Many reports talk about DCF modeling in boilerplate language without providing any specifics. In 9 out of 10 reports which mention a DCF model, "there is no recognizable DCF model provided in the report itself. (Green, Hand, and Zhang, 2016)"

Sell-side analysts rarely use a discount model in isolation (5.5% of the time; just 28 reports). 19 of these 28 discount-model-only reports were written by three analysts at Credit Suisse. In many ways, the 19 DCF-only reports are the worst pieces of research in our sample. The methods section in Figure 9 is not even a complete sentence. When we compare with Table 6, we see that analysts were more likely to use multiple multiples (38.8% of reports) than to use any sort of discounting model (30.2% of reports).

One might expect that analysts would be more likely to use a DCF model in coverage-initiation reports. After all, many of the 174 reports in our sample were written about newly public firms with little historical data. On top of this, the reports themselves tend to be longer and more thorough. However, the "All

J.P.Morgan

North America Equity Research
12 April 2010

Avis Budget and Dollar Thrifty

Initiate CAR at OW; Positively Inclined Toward Neutral-Rated DTG

Himanshu Patel, CFA[AC]
(1-212) 622-3906
himanshu.patel@jpmorgan.com

Ryan Brinkman
(1-212) 622-4137
ryan.brinkman@jpmorgan.com

CAR, CAR US
Overweight
$14.91

Price Target: $20.00

DTG, DTG US
Neutral
$36.21

Price Target: $39.00

(a) Top of first page

Valuation

We establish a Dec-2010 price target of $20 for CAR, based upon 13x our 2012E EPS of $1.55. We value Neutral-rated Hertz (HTZ) at 15x 2011E EPS. Relative to HTZ, we believe a discounted multiple for CAR is justified given CAR's lower margins and higher leverage on trailing EBITDA.

Historically, CAR traded at a premium to DTG on 12-month-forward P/E.

But, since 2001, CAR and DTG have been trading more synchronously.

Figure 17: Historical 12-Month-Forward P/E for CAR and DTG

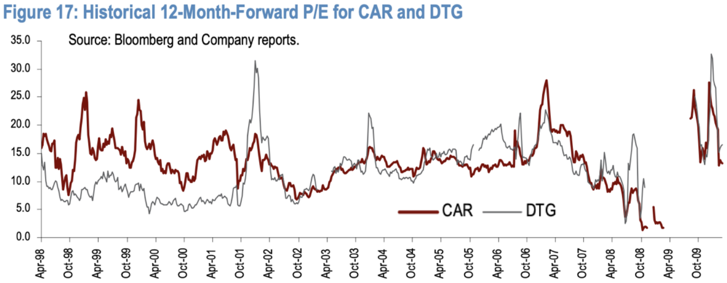

(b) Methods section

Figure 8. *Coverage-initiation report about Avis Budget (CAR), which was published on April 10th 2010 by JP Morgan. The lead analyst on this report was Himanshu Patel, a member of Institutional Investor magazine's All-America team.*

Am" column in Table 7 shows that DCF analysis is even less common in this subset of our data. Only one in five coverage-initiation reports makes use of a discount model in any capacity (34 of 174 reports; 19.5%).

The All-American analysts who are responsible for these reports often talk about DCF models as a second-best option. For example, Figure 10 shows a coverage-initiation report about Pacific Biosciences (PACB) from December 2010. In the methods section of his report, the lead analyst explains that while "multiple-based valuations (e.g., P/E and EV/EBITDA) are common in the life sci-

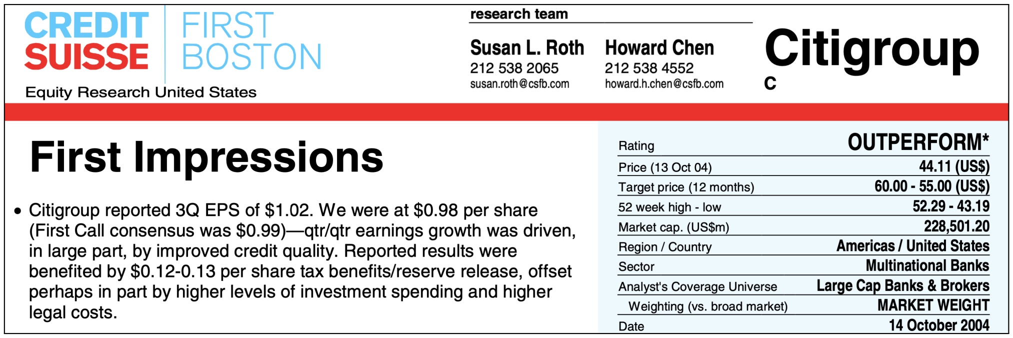

CREDIT SUISSE | FIRST BOSTON

Equity Research United States

research team

Susan L. Roth
212 538 2065
susan.roth@csfb.com

Howard Chen
212 538 4552
howard.h.chen@csfb.com

Citigroup
C

First Impressions

- Citigroup reported 3Q EPS of $1.02. We were at $0.98 per share (First Call consensus was $0.99)—qtr/qtr earnings growth was driven, in large part, by improved credit quality. Reported results were benefited by $0.12-0.13 per share tax benefits/reserve release, offset perhaps in part by higher levels of investment spending and higher legal costs.

Rating	OUTPERFORM*
Price (13 Oct 04)	44.11 (US$)
Target price (12 months)	60.00 - 55.00 (US$)
52 week high - low	52.29 - 43.19
Market cap. (US$m)	228,501.20
Region / Country	Americas / United States
Sector	Multinational Banks
Analyst's Coverage Universe	Large Cap Banks & Brokers
Weighting (vs. broad market)	MARKET WEIGHT
Date	14 October 2004

(a) Top of first page

Method: Discounted Cash Flow (DCF) Valuation

(b) Methods section

Figure 9. *Earning report about Citigroup, which was published on October 14th 2004 by Credit Suisse. The lead analyst on this report was Susan Roth.*

ence tools industry," he has "chosen to use a DCF methodology" out of necessity. "PACB is unprofitable (and yet lacks revenue)." Sell-side analysts are perfectly capable of doing these calculations. But they typically choose not to.

When analysts do use a discount model, they often implement it in a way that is inconsistent with present-value reasoning. They place no special emphasis on the forward-looking nature of the DCF model. They think about $\left(\frac{1}{r-g}\right)$ as just another trailing multiple (Mukhlynina and Nyborg, 2020). Table 7 indicates that analysts blend together the price targets implied by a DCF model and a trailing multiple in 24.6% of our sample (126 reports).

In most of these reports, the analyst literally just takes the average. For example, Figure 11 shows a December 2019 report about Citigroup where the lead analyst, Mike Mayo, takes a "simple average of six valuation techniques (PE, price-to-book, dividend discount model, PE/G ratio analysis and sum of the parts for both PE and PB)." This is not how theorists model asset prices, but it is not so different from how empiricists perform robustness checks. After presenting the main regression results, a researcher will often start adding in lots of unrelated control variables and then point out how the original coefficient of interest never changes much (Harbaugh, Maxwell, and Shue, 2016).

Analysts rarely focus solely on discount rates

	2004	2011	2019	All Am	Total
Discount Model	45.1% 41	32.3% 30	32.3% 50	19.5% 34	30.2% 155
Multiples Analysis	85.7% 78	91.4% 85	96.8% 150	98.9% 172	94.5% 485
Only Discounting	14.3% 13	8.6% 8	3.2% 5	1.1% 2	5.5% 28
Only Multiples	54.9% 50	67.7% 63	67.7% 105	80.5% 140	69.8% 358
Both Approaches	30.8% 28	23.7% 22	29.0% 45	18.4% 31	24.6% 126
# Reports	91	93	155	174	513

Table 7. *"Discount Model": report described using either a discounted cash-flow (DCF) or dividend discount model to calculate the price target. "Multiples Analysis": report calculated a price target using multiples analysis. "Only Discounting": report calculated a price target based solely on a discount model. "Only multiples": report calculated a price target based solely on multiples analysis. "Both approaches": report described using both a discount model and multiples analysis to calculate its price target. Top number in each cell is the percent relative to the total for the column. e.g.,* 41 *of* 91 *reports in 2004 described using either a DCF or dividend discount model to calculate the price target,* 41/91 = 45.1%.

1.4 DCF Models Are Mainly Used In Niche Industries

We study 513 reports written about large publicly traded companies—i.e., the firms researchers typically have in mind when writing down models. For this group of firms, we find that sell-side analysts do not typically set price equal to expected discounted payoff. However, they do regularly use DCF analysis in specific niche industries. Present-value logic is the norm when analysts value shipping companies, which are set up as master limited partnerships (MLPs) for tax reasons. Analysts also use DCF models when valuing real-estate investment trusts (REITs) and resource-extraction companies (oil, gas, mining, etc).

It is obvious when an analyst is thinking in present-value terms. Figure 12 shows a coverage-initiation report written by Michael Webber about GasLog Ltd in January 2014. Michael Webber clearly states that he is using a DCF model. He

J.P.Morgan

Pacific Biosciences Inc.

Third Generation Sequencing Comes of Age; Initiate at Overweight

Tycho W. Peterson[AC]
(1-212) 622-6568
tycho.peterson@jpmorgan.com

Evan Lodes
(1-212) 622-5650
evan.lodes@jpmorgan.com

North America Equity Research
06 December 2010

Initiation
Overweight
PACB, PACB US
Price: $12.97
Price Target: $17.00

Life Science Tools & Diagnostics

(a) Top of first page

Valuation
Multiple-based valuations (e.g. P/E and EV/EBITDA) are common in the life science tools industry, though since PACB is unprofitable (and as yet lacks revenue), we have chosen to use a DCF methodology.

(b) Methods section

Figure 10. *Coverage-initiation report about Pacific Biosciences published on December 6th 2010 by JP Morgan. The lead analyst on this report was Tycho Peterson, a member of Institutional Investor magazine's All-America team.*

gives us the precise numerical inputs needed to do the calculation. Asset-pricing researchers assume that every earnings report looks like this. If they did, we would clearly be able to recognize this fact. We show that, outside of a few special situations, this is just not how the world works.

1.5 Expected Returns Do Not Reflect (Exotic) Risks

Asset-pricing textbooks argue that a stock's expected return will be determined by how its payoffs are distributed across good and bad future states of the world. This follows from a state-contingent generalization of Equation (1)

$$\text{Price}_t = \mathbb{E}_t\left[\frac{\text{Dividend}_{s,t+1} + \text{Price}_{s,t+1}}{1+r_s}\right] \tag{5a}$$

$$= \mathbb{E}_t\left[\left(\frac{1}{1+r_s}\right) \times \{\,\text{Dividend}_{s,t+1} + \text{Price}_{s,t+1}\,\}\right] \tag{5b}$$

The realization of the stochastic discount factor (SDF) in a given state, $m_s \overset{\text{def}}{=} \left(\frac{1}{1+r_s}\right)$, is the current price of an asset that will pay $1 next year in that state.

December 19, 2019 | Equity Research

WELLS FARGO SECURITIES

Citigroup Inc.

C: New ValueAct Extension; Up Price Target

Overweight/$97

Large-Cap Banks

(a) Top of first page

Price Target: $97 from $85

We use a simple average of six valuation techniques (PE, price-to-book, discount dividend model, PE/G ratio analysis and sum of the parts for both PE and PB). This average yields a price target of $97.

(b) Methods section

Figure 11. *Earning report about Citigroup, which was published on December 19th 2019 by Wells Fargo. The lead analyst on this report was Mike Mayo.*

To keep things simple, suppose there are just two states, $s \in \{\text{good}, \text{bad}\}$. In this framework, investors would be willing to pay $\$1 \cdot m_{\text{bad}} = \left(\frac{\$1}{1+r_{\text{bad}}}\right)$ today to receive \$1 next year in the bad state. Researchers assume that $\$1 \cdot m_{\text{bad}} = \left(\frac{\$1}{1+r_{\text{bad}}}\right) > \left(\frac{\$1}{1+r_{\text{good}}}\right) = \$1 \cdot m_{\text{good}}$ since that is when they will really need the money. When an asset's expected returns are unusually low, researchers figure that most of its future payoffs must arrive in some sort of bad state of the world that investors do not discount very much. The only question is which one?

On the one hand, our results are consistent with Bordalo, Gennaioli, La Porta, and Shleifer (2024), which argues that differences in expected returns are not compensation for bearing exotic state-contingent risks. None of the 513 earnings reports in our sample described anything remotely similar to the above logic. To be clear, we do not expect market participants to use jargon like "stochastic discount factor". However, the internal logic of the SDF approach requires most investors to think in a particular way. The average investor needs to be willing to strategically bid up the price of any asset that offers insurance against bad times. For this story to explain difference in expected returns, researchers cannot be the only people who think along these lines.

On the other hand, we find very little evidence that expected returns are compensation for simple risk either. Instead, sell-side analysts often say that their twelve-month return forecast comes from their view about a company's

January 13, 2014

Outperform / V

Sector: Marine GP
Overweight

WELLS FARGO SECURITIES

Equity Research

GasLog Ltd.

GLOG: Initiating Coverage With An Outperform Rating
Potential MLP Spin Could Create Significant Value

- **Summary: Long-Term LNG Fundamentals, Potential MLP Spin Create Value In GLOG, In Our View.**

Michael Webber, CFA, Senior Analyst
(212) 214-8019 /
michael.webber@wellsfargo.com
Donald D. McLee, Associate Analyst
(212) 214-8029 /
donald.mclee@wellsfargo.com
Sameed Musvee, Associate Analyst
(212) 214-8040 /
sameed.musvee@wellsfargo.com

(a) Top of first page

Discounted Cash Flow (DCF). As noted in Figure 6, using a WACC of 6.75-7.25% and a terminal multiple estimate of 12.0-13.0x (about 2.0x higher than the group's average of about 11.0x which gives modest credit for its GP potential), we estimate GLOG's value on a DCF basis to be $20-24 per share. As noted in Figure 7, GLOG's 2-year beta is 1.2x (CAPM), driving a cost of equity of around 10%, while we estimate its current marginal cost of debt to be about 5.0%. Given a long-term net debt-to-capital ratio of 60%, we estimate GLNG's WACC at 7.1%.

(b) Methods section

Valuation

Figure 5. GLOG Valuation Summary

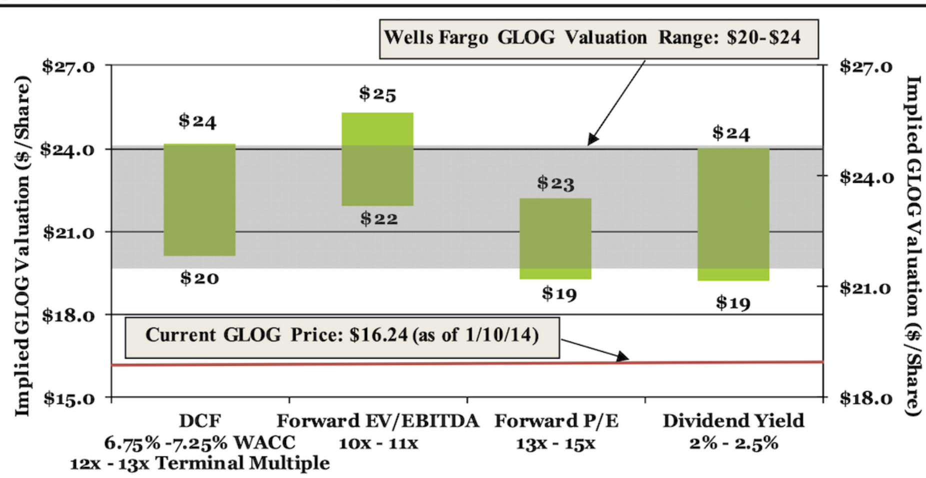

Source: Company filings, Wells Fargo Securities, LLC estimates

(c) Valuation summary

Figure 12. *Earning report about GasLog Ltd, which was published on January 13th 2014 by Wells Fargo. The lead analyst on this report was Michael Webber, a member of Institutional Investor magazine's All-America team.*

short-term earnings growth rate. For example, Figure 13 shows a coverage-initiation report written by Brian Tunick about Chico's FAS in May 2015. At the top of the first page, he predicts "mid-to high-teens total returns…comprised of 15% EPS CAGR (compound annual growth rate) from 2015–2017E and a ~2% dividend yield".

We note that this calculation is precisely what one would expect from an analyst who was in the habit of using a trailing P/E ratio to price a company's expected EPS. Suppose an analyst calculates $\mathbb{E}_t[\text{Price}_{t+1}] = \mathbb{E}_t[\text{EPS}_{t+2}] \times \text{TrailingPE}_t$ using a trailing twelve-month P/E, $\text{TrailingPE}_t = \text{Price}_t/\text{EPS}_t$. In that case, we could rewrite his one-year-ahead return forecasts as

$$\mathbb{E}_t[\text{Return}_{t+1}] = \left(\frac{\mathbb{E}_t[\text{Price}_{t+1}] - \text{Price}_t}{\text{Price}_t}\right) + \frac{\mathbb{E}_t[\text{Dividend}_{t+1}]}{\text{Price}_t} \tag{6a}$$

$$= \left(\frac{\mathbb{E}_t[\text{EPS}_{t+2}] \times \left(\frac{\text{Price}_t}{\text{EPS}_t}\right) - \text{Price}_t}{\text{Price}_t}\right) + \mathbb{E}_t[\text{DivYield}_{t+1}] \tag{6b}$$

$$= \left(\frac{\mathbb{E}_t[\text{EPS}_{t+2}]}{\text{EPS}_t} - 1\right) + \mathbb{E}_t[\text{DivYield}_{t+1}] \tag{6c}$$

This is extremely similar to the way that Brian Tunick made his return forecast for Chico's FAS. The main difference is that his "~2% dividend yield" was a trailing average rather than an expected value. More on this shortly.

1.6 Subjective Expectations Do Not Respect Identities

Asset-pricing researchers assume that Brian Tunick's price target for Chico's FAS (CHS) came from asking: "How much are the rights to CHS's expected future earnings worth in today's dollars?" Instead, he thought to himself: "If right now a company reported the earnings I expect CHS to generate in two years, how would this comparable firm be priced given recent multiples?" Given the logic behind his approach, there is no reason to believe Brian Tunick's price target would respect forward-looking accounting identities.

What does this mean? An asset's current price must satisfy the ex-post accounting identity $\text{Price}_t = (1 + \text{Return}_t) \cdot \text{Price}_{t-1} + \text{Dividend}_t$. But analysts do

EQUITY RESEARCH INITIATION

RBC Capital Markets

RBC Capital Markets, LLC
Brian Tunick (Analyst)
(212) 905-2926
brian.tunick@rbccm.com

Kate Fitzsimons (Associate)
(212) 428-6550
kate.fitzsimons@rbccm.com

Sector: Retailing/Department Stores & Specialty Softlines

May 4, 2015

Chico's FAS, Inc.

Initiating with Outperform; 2015–17 Self Help Puts $1.40 EPS Dream on the Table

Our view: After two years of earnings misses, we see Chico's as entering a period of positive estimate revisions in 2015 owing to company-specific actions to drive ticket gains and shore up profitability. As a result, we see CHS as a mid-to high-teens total returns story going forward, comprised of a 15% EPS CAGR from 2015–2017E and a ~2% dividend yield.

Outperform
NYSE: CHS; USD 17.09
Price Target USD 20.00
Scenario Analysis*

*Implied Total Returns

Downside Scenario	Current Price	Price Target	Upside Scenario
15.00	17.09	20.00	26.00
↓ 10%		↑ 19%	↑ 54%

(a) Top of first page

Valuation

Shares of CHS have traded at 14.3x for the last three years. Currently trading at 17.9x consensus FY2 P/E, we believe that CHS can see upside to the historical multiple given our expectation for at least accelerating mid-teens EPS growth, depressed margins, increased top- and bottom-line certainty owing to cost management in place, as well as generous use of the balance sheet. Our $20 price target applies roughly 20x to our 2016 EPS estimate of $0.98–a premium to our 15% 2015–2017E EPS CAGR owing to the above.

Trading at a 25% premium to three-year average of 14.3x

	FY 2 P/E		
	3 Year Avg	Today	Vs Avg
ANF	12.4x	19.3x	56%
LB	15.8x	21.8x	38%
ASNA	12.2x	16.7x	36%
CHS	**14.3x**	**17.9x**	**25%**
PLCE	14.2x	17.2x	21%
EXPR	11.4x	13.8x	21%
ANN	13.5x	16.1x	20%
AEO	14.0x	16.5x	18%
GPS	13.0x	13.0x	0%
LULU	27.9x	27.3x	-2%
URBN	17.4x	17.0x	-2%
FRAN	18.8x	17.3x	-8%
Group	**15.4x**	**17.8x**	**16%**

Depressed Margins, Increased Bottom-Line Certainty, and PE Carrot Can Support the Multiple

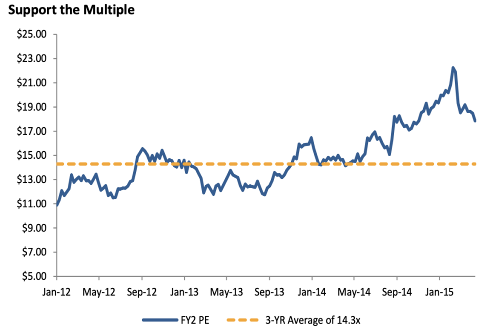

Source: Company reports and RBC Capital Markets estimates

(b) Methods section

Figure 13. *Coverage-initiation report about Chico's FAS, which was published on May 4th 2015 by RBC Capital Markets. The lead analyst on this report was Brian Tunick, a member of Institutional Investor magazine's All-America team.*

not have to use the ex-ante version of the same accounting identity when setting price targets, $\text{PriceTarget}_t \neq (1 + \mathbb{E}_t[\text{Return}_{t+1}]) \cdot \text{Price}_t + \mathbb{E}_t[\text{Dividend}_{t+1}]$. A price target is not a price. The calculation takes place entirely in an analyst's own head. They are free to run the numbers as they see fit.

This is an important theoretical point. Many researchers would like to believe that analysts' P/E ratios have to reflect $\left(\frac{1}{r-g}\right)$ in some way, shape, or form. This is not true. The Gordon model comes from replacing the Price_{t+1} on the right-hand side of Equation (1) with $\frac{\mathbb{E}_{t+1}[\text{Dividend}_{t+2}]+\mathbb{E}_{t+1}[\text{Price}_{t+2}]}{1+r}$ and then swapping out Price_{t+2} in the resulting expression with $\frac{\mathbb{E}_{t+2}[\text{Dividend}_{t+3}]+\mathbb{E}_{t+2}[\text{Price}_{t+3}]}{1+r}$ and so on... Researchers find this infinite recursion completely natural, but analysts' subjective beliefs do not have to satisfy such identities. We document that they do not, which breaks the connection with $\left(\frac{1}{r-g}\right)$.

The same critique applies to Campbell and Shiller (1988a), which provides a generalization of $\left(\frac{1}{r-g}\right)$ that allows for time-varying r and g

$$\frac{\text{Price}_t}{\text{Dividend}_t} \approx \left(\frac{1}{e^{\sum_{h=0}^{\infty} \rho^h \cdot \mathbb{E}_t[r_{(t+h)+1}] - \sum_{h=0}^{\infty} \rho^h \cdot \mathbb{E}_t[g_{(t+h)+1}]}}\right) \tag{7}$$

There is an entire strand of literature trying to understand the asset-pricing implications of biased subjective beliefs by plugging IBES data into this approximate present-value formula. But as John Campbell notes in his textbook, this exercise only make sense "for irrational expectations that respect identities. (Campbell, 2017)" Analysts' subjective beliefs do not exhibit this key property.

In fact, they suggest an entirely different thought process. To underscore this point, notice the tension between Andrea Teixeira's trading recommendation in Figure 14 and her backward-looking multiple. She gave Pepsi an "Overweight" rating, meaning that "[she] expected [the company to] outperform the average total return of the other stocks in [her] coverage universe. (JP Morgan, 2019b)" Yet, even though Ms Teixeira thought Pepsi's past price was too low, she still set her price target with a trailing 24× P/E.

The numbers in their reports frequently do not mean what researchers would guess. For example, Figure 14(c) shows a table of key metrics from an October 2019 earnings report written by Andrea Teixeira about Pepsi. The

J.P.Morgan

PepsiCo

Resilient Growth Continues to Drive PEP Higher; Reiterate OW

Andrea Teixeira, CFA AC
(1-212) 622-6735
andrea.f.teixeira@jpmorgan.com

North America Equity Research
03 October 2019

Overweight

PEP, PEP US
Price (03 Oct 19): $138.23
▲ **Price Target (Dec-20): $154.00**
Prior (Dec-20): $148.00

(a) Top of first page

Valuation

We rate PepsiCo Overweight. PEP is currently trading at ~24x our NTM EPS estimate, which is a 19% premium to the company's two-year average and a 17% premium to the five-year average. Our December 2020 price target moves to $154 (up from $148), based on 24x and our revised 2021 estimate. With the earnings re-base behind Pepsi by the end of this year and organic growth reaccelerating to the MSD range, we think the company will go back to be a growth compounder and maintain current valuation. We also still think Pepsi compares favorably to other large-cap multinational peers in our coverage universe because of the growth momentum in both developing and emerging markets.

(b) Methods section

Key Metrics (FYE Dec)

	FY18A	FY19E	FY20E	FY21E
Financial Estimates				
Revenue	64,662	66,871	69,252	72,026
Adj. EBITDA	13,019	13,081	14,068	15,092
Adj. EBIT	10,620	10,636	11,374	12,157
Adj. net income	8,065	7,739	8,285	8,833
Adj. EPS	5.66	5.50	5.95	6.41
Valuation				
EV/EBITDA	15.9	16.2	15.2	14.2
Adj. P/E	24.4	25.1	23.2	21.6

(c) Table of key metrics

Figure 14. *Report about Pepsi by Andrea Teixeira (JP Morgan, 2019b). The "Adj. EPS" row highlighted in red is Pepsi's announced (A) or expected (E) EPS in a given year. 2019 is marked as expected since Pepsi had not yet announced its Q4 numbers. The "Adj. P/E" row highlighted in blue is Ms Teixeira's own calculation for Pepsi's P/E ratio in that year.*

row highlighted in blue shows Pepsi's share price in October 2019, $\text{Price}_{\text{Oct‘19}} =$ \$138.23/sh, divided by its EPS in a given year

$$24.4\times = \frac{\$138.23/\text{sh}}{\$5.66/\text{sh}} = \frac{\text{Price}_{\text{Oct‘19}}}{\text{EPS}_{‘18}} \tag{FY18A}$$

$$25.1\times = \frac{\$138.23/\text{sh}}{\$5.52/\text{sh}} = \frac{\text{Price}_{\text{Oct‘19}}}{\mathbb{E}[\text{EPS}_{‘19}]} \tag{FY19E}$$

$$23.2\times = \frac{\$138.23/\text{sh}}{\$5.95/\text{sh}} = \frac{\text{Price}_{\text{Oct‘19}}}{\mathbb{E}[\text{EPS}_{‘20}]} \tag{FY20E}$$

$$21.6\times = \frac{\$138.23/\text{sh}}{\$6.41/\text{sh}} = \frac{\text{Price}_{\text{Oct‘19}}}{\mathbb{E}[\text{EPS}_{‘21}]} \tag{FY21E}$$

This is exactly the sort of P/E ratio one would expect from someone who is thinking about how a company's future earnings would be priced under current market conditions. No researcher would report these numbers as coming from the same variable in an academic paper. You probably would never even think to perform this calculation. And we think this is one reason why researchers have previously overlooked this glaring piece of evidence.

Skeptical? Let's run an experiment. Go back to page 3 in the introduction. In Figure 2, Chris Horvers calculated the P/E ratios in his valuation matrix just like Andrea Teixeira. Did you notice? Home Depot's closing price on December 11th 2019 was \$212.00, and the P/E ratios in Chris Horvers' valuation matrix correspond to $21.4\times = \frac{\text{Price}_{\text{Dec‘19}}}{\text{EPS}_{‘18}} = \frac{\$212.00/\text{sh}}{\$9.89/\text{sh}}$, $21.1\times = \frac{\text{Price}_{\text{Dec‘19}}}{\mathbb{E}[\text{EPS}_{‘19}]} = \frac{\$212.00/\text{sh}}{\$10.05/\text{sh}}$, $20.2\times = \frac{\text{Price}_{\text{Dec‘19}}}{\mathbb{E}[\text{EPS}_{‘20}]} = \frac{\$212.00/\text{sh}}{\$10.48/\text{sh}}$, and $18.4\times = \frac{\text{Price}_{\text{Dec‘19}}}{\mathbb{E}[\text{EPS}_{‘21}]} = \frac{\$212.00/\text{sh}}{\$11.50/\text{sh}}$. We would never have thought to look for this calculation prior to writing this paper. Our guess is that, before reading our paper, the thought had not crossed your mind either.

1.7 Analysts Price Expected EPS Not Expected Payoffs

Asset-pricing textbooks assume that analysts care about earnings because these cash flows allow a firm to pay dividends. Given this prediction, it is noteworthy how few of the earnings reports discuss a company's dividend payout rate. Table 4 shows that analysts mention a company's dividend yield in just 6.4% of all reports (33 of 513).

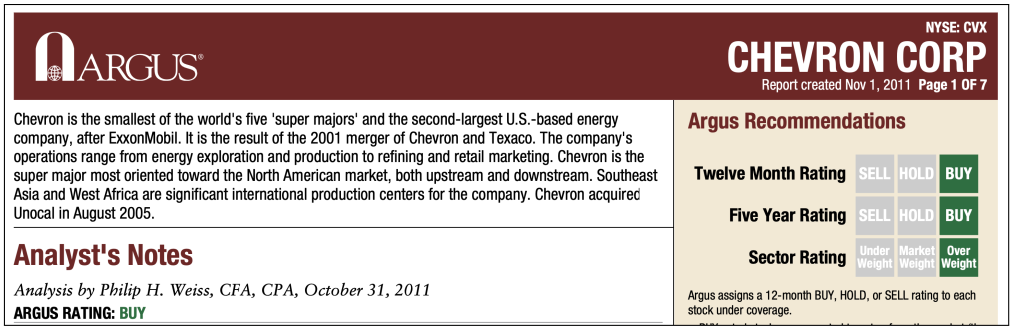
ARGUS

NYSE: CVX

CHEVRON CORP

Report created Nov 1, 2011 **Page 1 OF 7**

Chevron is the smallest of the world's five 'super majors' and the second-largest U.S.-based energy company, after ExxonMobil. It is the result of the 2001 merger of Chevron and Texaco. The company's operations range from energy exploration and production to refining and retail marketing. Chevron is the super major most oriented toward the North American market, both upstream and downstream. Southeast Asia and West Africa are significant international production centers for the company. Chevron acquired Unocal in August 2005.

Analyst's Notes

Analysis by Philip H. Weiss, CFA, CPA, October 31, 2011

ARGUS RATING: BUY

Argus Recommendations

Twelve Month Rating	SELL	HOLD	BUY
Five Year Rating	SELL	HOLD	BUY
Sector Rating	Under Weight	Market Weight	Over Weight

Argus assigns a 12-month BUY, HOLD, or SELL rating to each stock under coverage.

(a) Top of first page

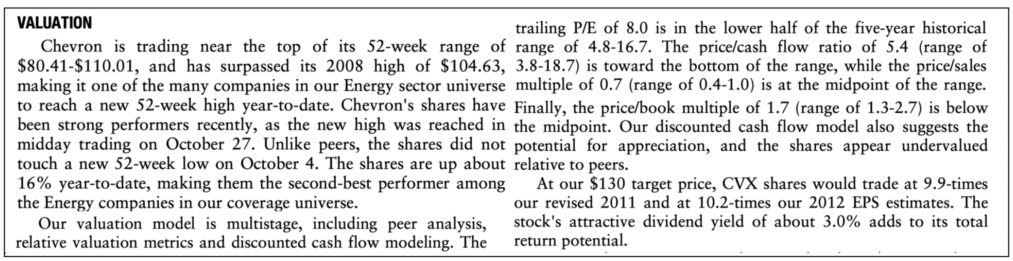
VALUATION

Chevron is trading near the top of its 52-week range of $80.41-$110.01, and has surpassed its 2008 high of $104.63, making it one of the many companies in our Energy sector universe to reach a new 52-week high year-to-date. Chevron's shares have been strong performers recently, as the new high was reached in midday trading on October 27. Unlike peers, the shares did not touch a new 52-week low on October 4. The shares are up about 16% year-to-date, making them the second-best performer among the Energy companies in our coverage universe.

Our valuation model is multistage, including peer analysis, relative valuation metrics and discounted cash flow modeling. The trailing P/E of 8.0 is in the lower half of the five-year historical range of 4.8-16.7. The price/cash flow ratio of 5.4 (range of 3.8-18.7) is toward the bottom of the range, while the price/sales multiple of 0.7 (range of 0.4-1.0) is at the midpoint of the range. Finally, the price/book multiple of 1.7 (range of 1.3-2.7) is below the midpoint. Our discounted cash flow model also suggests the potential for appreciation, and the shares appear undervalued relative to peers.

At our $130 target price, CVX shares would trade at 9.9-times our revised 2011 and at 10.2-times our 2012 EPS estimates. The stock's attractive dividend yield of about 3.0% adds to its total return potential.

(b) Methods section

Figure 15. *Earning report about Chevron Corp, which was published on November 1st 2011 by Argus Research. The lead analyst on this report was Philip Weiss.*

While most analysts do not use any sort of present-value model, those that do tend to compute the present discounted value of a company's cash flows not its dividend payouts to shareholders. Outside of a few special cases, analysts consistently ignore a company's plowback rate. Suppose that two firms have the same future earnings stream, but one pays out a much larger dividend. Most analysts would assign both firms the same price target. The December 2019 Wells Fargo report about Citigroup in Figure 11 is one of the few reports that specifically talks about using a dividend discount model.

When analysts do mention a company's dividend yield, they typically only use it to compute their return forecast. Dividends rarely play a role in setting price targets. In short, analysts "track capital gains and dividends as separate and largely independent variables. (Hartzmark and Solomon, 2019)"

Figure 15 shows a November 2011 report about Chevron Corp written by Philip Weiss. In the methods section of his report, Mr Weiss says he considered both trailing multiples analysis as well as a DCF model when setting $\text{PriceTarget}_t = \$130/\text{sh}$, which was 24% higher than Chevron's current price, \$105.05/sh. But he did not use Chevron's dividend yield to set his price target.

Chevron's dividend yield only showed up when Mr Weiss made his "buy" recommendation. Chevron had paid a dividend of \$3.12/sh per share to each shareholder in 2011. Mr Weiss argued that an investor should expect Chevron's returns to reflect both the 24% capital gain implied by his price target as well as the company's trailing-twelve-month dividend yield, $\frac{\$3.12/\text{sh}}{\$105.05/\text{sh}} \approx 3\%$,

$$\underbrace{\mathbb{E}_t[\text{Return}_{t+1}]}_{27\%} = \underbrace{\left(\frac{\text{PriceTarget}_t - \text{Price}_t}{\text{Price}_t}\right)}_{(\$130.00 - \$105.05)/\$105.05 \approx 24\%} + \underbrace{\left(\frac{\text{Dividend}_t}{\text{Price}_t}\right)}_{\$3.12/\$105.05 \approx 3\%} \tag{8}$$

It might at first seem like Mr Weiss was following textbook logic, but the nature of these two differences shows that he was not. A company's expected return should be equal to its expected capital gain plus its expected dividend yield, $\mathbb{E}_t[\text{Return}_{t+1}] = \left(\frac{\mathbb{E}_t[\text{Price}_{t+1}] - \text{Price}_t}{\text{Price}_t}\right) + \left(\frac{\mathbb{E}_t[\text{Dividend}_{t+1}]}{\text{Price}_t}\right)$, but this is not what Mr Weiss calculated in Equation (8). His price target was based on trailing multiples, and he used Chevron's trailing twelve-month (TTM) dividend yield rather than its expected dividend yield next year, $\text{Dividend}_t \neq \mathbb{E}_t[\text{Dividend}_{t+1}]$.

Mr Weiss did not just deviate from textbook logic. He did the precise opposite. On the first page of his report, he predicts that Chevron's dividend yield will grow 8.80% over the next twelve months. Yet he still used a trailing dividend yield to calculate a 27% expected return. That is not a mistake. It is a choice.

1.8 Do Analysts Give Credible Descriptions?

We have talked to a large number of analysts. Our general sense is that the methods sections of their reports contain brief honest accounts of how their price targets were calculated. Researchers are clearly comfortable using analysts' numerical forecast values. If these numbers represent a credible data source, we see no reason to discard the data about how they were calculated. Why should "4" be any more worthy of study than "two times two"?

Moreover, even if you think analysts do not put much effort into writing the methods section of their reports, this fact should not push them towards using a trailing P/E rather than a DCF model. It is just as easy to give a brief account of either. Many DCF-only reports have a one-line methods section (see Figure 9).

It is true that analysts are more likely to include a price target in an earnings report when they are optimistic about a company's future prospects (Brav and Lehavy, 2003). However, while this fact introduces an upward bias into analysts' price targets, it has no implications for the way that analysts describe their approach. It is just as easy to plug a small r into $(\frac{1}{r-g})$ as it is to cherry-pick a favorable trailing window when calculating a P/E.

Unlike an active investor with a profitable trading rule, a sell-side analyst has no incentive to hide their pricing rule. If anything, their incentives point in the opposite direction. Sell-side analysts are in the business of writing research articles that advertise how thoroughly they understand a company's fundamentals and future prospects. Misleading their readership about which pricing rule they are using does not help them accomplish this goal.

1.9 Is Our Data Sample Representative?

Researchers currently assume that analysts' price targets reflect the discounted value of their subjective earnings forecasts. We have more than enough statistical power to reject this hypothesis. Analysts apply a trailing multiple in 485 out of 513 reports (94.5% of our sample). By contrast, only 155 reports (30.2%) mention a DCF model. There are just 28 reports (5.5%) that exclusively rely on DCF analysis.

The majority of the reports in our sample use some version of $\text{PriceTarget} = \mathbb{E}[\text{EPS}] \times \text{TrailingPE}$. There are good reasons to think that this is the norm. According to the CFA Institute (2017), "historical average valuation multiples are frequently used in equity analysis as a reference point or as justification of a target multiple at which the shares are expected to trade in the future." We also directly test external validity in Subsection 3.2. The simplest version of this formula explains $R^2 = 91\%$ of the price-target variation in IBES.

There are two recent working papers, which seem to arrive at the opposite conclusion when studying much larger samples. It turns out that both studies highlight how infrequently DCF models get used.

Décaire and Graham (2024) applies natural-language processing (NLP) techniques to analyze the discount rates found in 78.5k analyst reports downloaded from Refinitiv. When describing their data, the authors acknowledge that only "40% of all reports" include a DCF model. Thus, the 78.5k reports in that study are analogous to the subset of 155 reports in our data (30.2% of sample) where the analyst mentions a DCF model. This is a non-representative subsample. Analysts do not use DCF at random (see Figure 10).

Gormsen and Huber (2024) employs a team of research assistants to analyze what managers said to sell-side analysts in 74k quarterly earnings calls. Just like before, the authors find that most transcripts do not make any reference to present-value logic. More than 60% of S&P 500 companies have never quoted a specific discount rate in *any* conference call over the past two decades. The same ballpark 30%-40% value shows up again.

1.10 Do Other Investors Think Like Analysts?

Suppose that, for the sake of argument, sell-side analysts were the only ones using trailing P/E ratios. Even if all other investors set price equal to expected discounted payoff, researchers do not get to observe these other investors' subjective beliefs. Much of what researchers think they know about discount rates comes from studying analysts' earnings forecasts. Our findings show that, for the most part, this particular group of market participants does not use one.

There are also good reasons to believe that sell-side analysts are not the only ones using trailing P/Es. It is called *sell*-side research. Presumably there are other investors interested in buying this research output. Sell-side analysts have been around in something resembling their current form since the 1970s. It seems implausible that no one uses the output of their calculations. Apple was founded in 1976. Given how long the company has lasted, it would be odd if no one had ever seen someone using a MacBook.

Regulatory filings tend to use multiples analysis for valuations

		# Reports	Discount Model	Multiples Analysis	Both Approaches
All Public Firms	8-K	628,446	17.3%	93.2%	10.5%
Firms Going Private	SC 13E3	5,410	75.1%	93.4%	68.5%
Public Acquirers	SC TO-T	4,953	19.9%	91.7%	11.6%
M&A Targets	SC 14D9	4,084	59.7%	90.3%	50.0%
Activist Shareholders	SC 13D	9,674	17.3%	90.4%	7.8%
Passive Blockholders	SC 13G	9,562	1.7%	98.3%	0.0%
Fund Managers	NPORT-P	36,520	39.9%	88.2%	28.1%
Total (w/o 8-Ks)		70,203	34.0%	90.6%	24.7%
Total		698,649	19.0%	92.9%	11.9%

Table 8. *Valuation method used in regulatory filings submitted to the Securities and Exchange Commission (SEC) from January 2001 through November 2023. "# Reports": number of reports with an explicit price calculation. "Discount Model": percent that used either a DCF or dividend discount model to do this calculation. "Multiples Analysis": percent that used multiples analysis. "Both Approaches": percent of documents that referenced a discount model and multiples analysis.*

We examine price calculations in seven different kinds of SEC regulatory filings from January 2001 through November 2023: (1) 8-K; a public company must submit one of these "current report" forms any time a major event takes place. (2) SC 13E3; a public company must file this form when going private. (3) SC TO-T; a public company must file this form when it makes a tender offer for another company's shares as part of a takeover bid. (4) SC 14D9; the target of this takeover bid must file its response to the tender offer using this form. (5) SC 13D; an investor must file this "beneficial ownership" form within 10 days of acquiring ownership of $\geq$ 5% of a company's stock. (6) SC 13G; this is an abbreviated version of form SC 13D, which is often used by large passive investors. (7) NPORT-P; 1940-Act funds use this form to report holdings, performance, assets under management, etc on a quarterly basis.

The last row of Table 8 shows that only 19.0% of all valuation-related forms in our sample included any of the following terms: "DCF", "discounted cash", "beta", "WACC", or "present value". By contrast, we find that 92.9% of these forms included the term “multiples” or “comparables”. 8-K filings make up 628k/698k = 90% of all valuation-related filings in our sample. So you might worry our results are being skewed by this one particular kind of form. But the second-to-last row of Table 8 should allay this concern. When we look at the remaining 70k observations, only 34.0% mention any sort of discount model while 90.6% talk about multiples analysis.

1.11 The Disconnect Between Theory And Practice

The day-to-day business of being an asset-pricing theorist involves writing down models of expected returns. Textbook models say that investors equate a company’s share price with its expected discounted payoff to each shareholder. Researchers assume that investors use the model-implied expected return as their discount rate when setting price levels. They see this discount rate as the most important part of the problem.

It is hard to escape the conclusion that researchers are modeling a problem that does not matter to sell-side analysts in the real world. Sell-side analysts describe their price-forecasting problem in an entirely different way than an academic researcher would. The sell-side analysts in our sample focus all their attention on predicting a company’s earnings and then pick a trailing P/E almost as an afterthought.

While the behavioral-finance literature has primarily studied biased EPS forecasts, at least analysts are trying hard to get those numbers right. In contrast, they do not even attempt to calculate the expected discounted payoff at the heart of every standard asset-pricing model.

We appreciate that every profession does some things on autopilot. For example, Petersen (2008) pointed out that, in the past, researchers often did not put too much thought into how they clustered their standard errors. The surprising thing is that analysts so pay little attention to the thing (the “P” in the P/E ratio) asset-pricing researchers obsess over. That is noteworthy. Even if analysts are not the marginal investor, this fact changes how we interpret decades of previous research.

2 A Simple Model

In this section, we build the simplest possible model that reflects how sell-side analysts say they price their own subjective cash-flow expectations. Then we explore the implications that follow. In principle, this modeling exercise might yield no fruit. A trailing P/E is a significant departure from present-value logic. Massive changes like this often render the resulting analysis completely intractable. This is why behavioral researchers typically study portable one-step extensions of an existing fully rational benchmark (Rabin, 2013).

Yet, we find that our simple model neatly rationalizes the use of trailing P/E ratios. It makes sense to set price targets using a trailing P/E because prices in our model turn out to be mostly backward-looking. The only forward-looking input is analysts' short-term EPS forecast. We show that this fact also generates sharp testable predictions. It implies that a piece of news can only affect returns by causing analysts to revise their choice of $\mathbb{E}[\text{EPS}]$. Because the trailing P/E is set in stone, any effect must operate through this one narrow channel.

2.1 Market Setting

In textbook models, investors care about earnings only insofar as these earnings translate into future payouts. However, as noted in the previous section, sell-side analysts price earnings for earnings' sake. So we put earnings per share (EPS) at the center of our model. We study a single company with earnings over the past twelve months, EPS_t, that are governed by the following law of motion

$$\left(\frac{\text{EPS}_{t+1} - \text{EPS}_t}{\text{EPS}_t}\right) = X_t + \epsilon_{t+1} \tag{9}$$

$X_t \approx \mathbb{E}_t[\Delta \log \text{EPS}_{t+1}]$ is the expected rate at which the company's earnings will grow over the next year, and $\epsilon_{t+1} \overset{\text{IID}}{\sim} \text{Normal}(0, \sigma^2)$ is a noise term.

Think about a firm that had earnings of $\text{EPS}_t = \$1.00/\text{sh}$ over the past year. Over the next twelve months, its earnings are expected to grow by $X_t = 5\%$ on average. But investors would not be surprised to see growth that was $\sigma = 2\%\text{pt}$

higher or lower. Given these assumptions, investors expect the company to generate earnings of $\mathbb{E}_t[\text{EPS}_{t+1}] = \$1.05(\pm\$0.02)/\text{sh}$ for each shareholder over the next year and $\mathbb{E}_t[\text{EPS}_{t+2}] = \$1.11(\pm\$0.04)/\text{sh}$ the year after.

Let Price_t denote the company's current price level. At each time t, the analysts in our model set a one-year-ahead price target

$$\text{PriceTarget}_t = \mathbb{E}_t[\text{EPS}_{t+2}] \times \text{TrailingPE}_t \tag{10}$$

For simplicity, we will assume analysts calculate the firm's trailing P/E ratio using the past twelve months of data, $\text{TrailingPE}_t \overset{\text{def}}{=} \text{Price}_t \,/\, \text{EPS}_t$.

Sell-side analysts make explicit trading recommendations by comparing their price target to a company's current price. These recommendations can be worth acting on (Birru, Gokkaya, Liu, and Stulz, 2022), and they focus on the relative price difference. For example, in an October 2019 report, Kaumil Gajrawala describes how he "[rated] PepsiCo *underperform* based on its expected return relative to our target price. (Credit Suisse, 2019)"

In our model, investors compare analysts' price target to the current price and adjust their demand proportionally

$$\left(\frac{\text{Demand}_{t+1} - \text{Demand}_t}{\text{Demand}_t}\right) = \mu \cdot \left(\frac{\text{PriceTarget}_t - \text{Price}_t}{\text{Price}_t}\right) \tag{11}$$

$\mu > 0$ is a positive constant, which is known as a demand "multiplier". When the price target is higher, $\frac{\text{PriceTarget}_t - \text{Price}_t}{\text{Price}_t} > 0\%$, they tell their broker to buy shares over the next year. When $\frac{\text{PriceTarget}_t - \text{Price}_t}{\text{Price}_t} < 0\%$, they say to sell.

To make things concrete, suppose the company is currently trading at $\text{Price}_t = \$100/\text{sh}$ and the demand multiplier is $\mu = 1$. If analysts set a one-year-ahead price target of $\text{PriceTarget}_t = \$103/\text{sh}$, then investors would respond by increasing their holdings $1 \cdot \left(\frac{\$103/\text{sh} - \$100/\text{sh}}{\$100/\text{sh}}\right) = 3\%$ over the next year. If investors currently hold $\text{Demand}_t = 300{,}000$ shares. Then, a year from now, they would like to own $\text{Demand}_{t+1} = 309{,}000$ shares in this example.

We saw in the previous section that sell-side analysts spend most of their time fine-tuning their EPS forecast. Then, when it comes time to capitalize these expected earnings into a price target, they use a trailing P/E. In other words, real-world analysts set price targets by asking themselves: "What would the firm's price be at current multiples if it had realized earnings of $\mathbb{E}_t[\text{EPS}_{t+2}]$ rather than EPS_t today?"

In our model, investors adjust their holdings based on the thing analysts actually care about. By plugging analysts' formula for creating price targets (Equation 10) into Equation (11), we see that

$$\left(\frac{\text{Demand}_{t+1} - \text{Demand}_t}{\text{Demand}_t}\right) = \mu \cdot \left(\frac{\mathbb{E}_t[\text{EPS}_{t+2}] - \text{EPS}_t}{\text{EPS}_t}\right) \tag{12}$$

Investors' demand will respond to changes in analysts' beliefs about a firm's short-term earnings growth.

To close the model, we need to make an assumption about how changes in investor demand affect asset prices. We take the simplest possible approach. We assume there exists a strictly positive constant, $\nu > 0$, such that

$$\left(\frac{\text{Price}_{t+1} - \text{Price}_t}{\text{Price}_t}\right) = \nu \cdot \left(\frac{\text{Demand}_{t+1} - \text{Demand}_t}{\text{Demand}_t}\right) + \varepsilon_{t+1} \tag{13}$$

If investors tell their broker to increase their positions by 1%pt over the upcoming year, then the company's share price will increase by ν%pt on average.

We are not claiming that analysts' short-term EPS forecasts explain every bump and jiggle in a firm's share price. The noise term $\varepsilon_{t+1} \overset{\text{IID}}{\sim} \text{Normal}(0, \varsigma^2)$ captures the many other reasons why a company's share price might increase or decrease over the next year. If a company's future returns are affected by analysts' price targets, our model should tell us what this effect will look like.

The goal is to build the simplest possible asset-pricing model that reflects what sell-side analysts say they do and then explore the implications that follow. We recognize this is different from the usual theoretical approach, which involves building a model from the ground up based on first principles that seem reasonable to us as researchers. However, the previous section showed that analysts typically set price targets in a way that does not follow from researchers' first principles. This is why we have taken a different tack.

2.2 Correct On Average

In a world where sell-side analysts do not apply present-value logic, there is no reason to expect their price targets to equal the present discounted value of a firm's expected future dividends. But this does not imply that their price targets are wrong. If the firm's equilibrium price also does not stem from present-value logic, then analysts' price targets might still be roughly correct.

It would make sense to use a trailing P/E in a world where prices are mostly backward-looking. It turns out that this is exactly what happens in our simple model. It is possible for a mostly backward-looking pricing rule to be correct on average because prices themselves are mostly backward-looking.

Proposition 2.2 (Correct On Average)**.** *Suppose that investors choose their demand according to Equation* (11) *and that realized price growth is governed by the law of motion in Equation* (13)*. If $\nu = 1/\mu$, then*

$$\hat{\mathbb{E}}_t[\mathrm{Price}_{t+1}] = \mathbb{E}_t[\mathrm{EPS}_{t+2}] \times \mathrm{TrailingPE}_t \tag{14}$$

where $\hat{\mathbb{E}}_t[\mathrm{Price}_{t+1}]$ is the average price next year observed by an econometrician.

Most analysts say they use a trailing P/E to set price targets à la Equation (10). They also explain how their trading recommendations come from comparing a company's price target for next year to its current price level as shown in Equation (11). Given these two starting points, it is not surprising that there exists some price path under which it makes sense to use trailing P/E ratios. The surprising thing is that the required price path in Equation (13) is so simple.

For example, Grossman and Stiglitz (1980) guessed that a risky asset's price would be a linear function of a signal about the asset's future payout and an aggregate supply shock, Price $= A + B \cdot$ Signal $- C \cdot$ Shock. The authors figured out what this price function "implied for risky asset demand, substituted that demand function into the market-clearing condition, and matched coefficients to verify their [initial] hypothesis (Veldkamp, 2011)" about the price function being linear. We do something similar, except that we match coefficients to verify that the law of motion for price growth is linear.

However, notice that the functional forms in Grossman and Stiglitz (1980) were dictated by theoretical considerations. The authors studied a CARA-normal setting because that would make it natural to expect a linear pricing rule. By contrast, our functional forms were not chosen to make the model more tractable. They were dictated by the descriptions offered by real-world market participants. It turns out the resulting asset-pricing model is tractable anyway.

Equations (11) and (13) are written down in percentage changes, so μ and ν can be seen as a demand multiplier and a price elasticity. This connects our work to the literature on demand-system asset pricing (Koijen and Yogo, 2019; Gabaix and Koijen, 2024). Under this interpretation, it would be natural to expect $\mu = 1/\nu$ as required by Proposition 2.2.

That being said, the two parameters play very different roles in the two sets of models. The demand-system framework cares about μ and ν because they play a pivotal role in how markets clear when investors solve a forward-looking portfolio problem. By contrast, in our model, μ and ν emerge from taking seriously how analysts describe their pricing rule.

Our model describes how a firm's price will change over the next year given its trailing P/E today. But where does the first TrailingPE come from? Our analysis of coverage-initiation reports in the previous section sheds light on the answer. Sometimes, when a firm goes public, analysts look at the trailing P/E ratios of similar firms. Other times, they rely on a revenue multiple, such as EV/EBITDA. That being said, a company's trailing P/E when it first went public twenty years ago need not have any effect on how an analyst sets her price target today. A boundary condition need not affect interior solutions.

2.3 Exclusion Restriction

The internal logic of our model also suggests a simple way to test it. Analysts in our model only affect the price path through their beliefs about short-term EPS growth. So, according to our model, if a piece of news causes analysts to revise their beliefs, the subsequent price response should only reflect the revision in their short-term EPS forecast.

Because the firm's trailing P/E is already determined, our model naturally leads to a novel exclusion restriction—i.e., a "claim that an instrument operates through a single known channel. (Angrist and Pischke, 2009)" Unlike in the Gordon model, news that changes analysts' beliefs about the firm's long-run EPS growth rate, g, will have no pricing implications.

Let News_t be a piece of information revealed about a firm at time t. In our model, if News_t predicts the company's future return, then it must be correlated with changes in the firm's short-term EPS growth. It does not help to be correlated with other future outcomes that matter in textbook models.

Proposition 2.3 (Exclusion Restriction). *If* News_t *is uncorrelated with a firm's short-run EPS growth,* $\widehat{\text{Corr}}(X_t, \text{News}_t) = 0$*, then analysts' forecast revisions in response to this information release will not affect the company's subsequent returns*

$$\widehat{\text{Corr}}\big(\text{Return}_{t+1}, \text{News}_t\big) = 0 \tag{15}$$

This is true even if News_t *is correlated with expected EPS growth farther in the future,* $\widehat{\text{Corr}}(X_{t+h}, \text{News}_t) \neq 0$ *for* $h \geq 1$*, or with the discount rate that a forward-looking present-value investor would use,* $\widehat{\text{Corr}}(r_{t+h}, \text{News}_t) \neq 0$*.*

Researchers typically focus on things that *should* affect prices. An asset-pricing model's key predictions usually come from digging into the economic forces that determine the key parameters. Think about Grossman and Stiglitz (1980). The main predictions in that paper came from understanding the coefficient B in the pricing rule Price $= A + B \cdot \text{Signal} - C \cdot \text{Shock}$. The authors showed that, if more investors were to buy the private signal and become informed, the B coefficient would get larger, resulting in a negative feedback loop.

By contrast, the interesting thing about our model is all the things that *should not* affect prices. News cannot change the past. The firm's trailing P/E ratio is what it is. The only way a piece of news can alter investors' demand (and thus the equilibrium price) in our model is by changing the one forward-looking component: expected short-term EPS growth. In the following section, we will take this idea to the data by examining how market prices respond to a particular kind of news, earnings surprises.

BTIG

U.S. Equity Research
Telecommunications

AT&T Inc.

May 20, 2010

Initiating coverage with Neutral rating

- We expect EPS to fall 14% in 2012 primarily based on our belief that a mid-year 2011 introduction of an iPhone by an another carrier will lead to the loss of 4.4 million iPhone customers in the subsequent 6 quarters and a $6 billion reduction in revenue, reflecting 10% decline in post-paid service revenue.

Walter Piecyk
(212) 527-3524
wpiecyk@btig.com

Joseph Galone
(212) 527-3523
jgalone@btig.com

(a) Top of first page

AT&T has been a global pioneer in the transition to integrated phones because of its exclusive contract with Apple. We also think we constructively framed up the risk to AT&T's earnings from the loss of the iPhone in 2012. But that's just it. The EPS disaster we foretell is in 2012 and in the meantime we think AT&T will deliver upside to the near term consensus EPS estimates. It was also fairly challenging to construct a valuation target that would generate enough downside to merit a Sell rating. However, we realized that investors might notice that glaring fall-off in EPS growth that begins in 2011 and gets nasty in 2012.

(b) Methods section

Figure 16. *Earning report about AT&T published on May 20th 2010 by BTIG. The lead analyst on this report was Walter Piecyk, a member of Institutional Investor magazine's All-America team.*

But before we get there, it is worth asking: Is it even reasonable to think that analysts ignore information about a company's earnings three years from now? Yes. We found examples of this when reading our sample of 513 analyst reports. Figure 16 shows a May 2010 coverage-initiation report about AT&T written by Walter Piecyk, which describes this exact reasoning. Walter Piecyk recognizes that AT&T's earnings will plummet in three years when the company loses its exclusive contract for iPhones. So he concludes: "But that's just it. The EPS disaster we foretell is in 2012…[making it] fairly challenging to construct a valuation target that would generate enough downside to merit a Sell rating." AT&T's fiscal year 2012 was two years after Mr Piecyk's target date at the time he wrote his report in May 2010—i.e., $(t+3)$ in model time.

3 Econometric Analysis

In this section, we present two main sets of empirical results, each with a different purpose in mind. After describing our data, we first show that $\text{PriceTarget} = \mathbb{E}[\text{EPS}] \times \text{TrailingPE}$ explains over $R^2 = 91\%$ of the price-target variation in IBES data when using a simple trailing twelve-month P/E. This first set of results confirms that our conclusions from reading 513 reports extend to the broader data sample that researchers typically study. In the second part, we then show that market prices respond to earnings surprises in a way that is consistent with our model's novel exclusion restriction, which we derived in the previous section. This evidence suggests that analysts' use of trailing P/Es helps explain equilibrium prices, not just their own subjective beliefs.

3.1 Data description

We use data from IBES and the merged CRSP/Compustat daily file. We restrict our sample to common stocks (share codes 10 and 11) traded on NYES, Nasdaq, or AmEx during the period from 2003 to 2022. For the reasons discussed above, we exclude firms in the following six Fama-French industries: real estate, coal, steel, mines, oil, and gold.

Analysts forecast a company's price level at the end of the upcoming fiscal year. We refer to this future date as the "target date" and denote it with $(\tau + 1)$. For example, Chris Horvers wrote a report in December 2019 that set a price target of \$241/sh for Home Depot in December 2020 (target date).

We distinguish between trading days t and target dates τ because an analyst can revise his/her forecast for the same target date on successive trading days. For each analyst a tracking a particular firm n, we record their most recent price target, $\text{PriceTarget}_{n,t}^{a} = \mathbb{E}_t^a[\text{Price}_{n,\tau+1}]$, from 18 months to 6 months prior to each target date $(\tau + 1)$.

We write the analyst's corresponding EPS forecast as $\mathbb{E}_t^a[\text{EPS}_n]$. We use the two-year-ahead EPS forecast when available in IBES, $\mathbb{E}_t^a[\text{EPS}_{n,\tau+2}]$, otherwise we use the one-year-ahead value, $\mathbb{E}_t^a[\text{EPS}_{n,\tau+1}]$. We restrict our sample to include

Summary Statistics

	#	Avg	Sd	Min	Med	Max
	(1)	(2)	(3)	(4)	(5)	(6)
$\text{PriceTarget}^a_{n,t}$	2,394,531	\$67.63	\$147.53	\$1.00	\$38.00	\$5,500.00
$\mathbb{E}^a_t[\text{EPS}_{n,\tau+1}]$	2,004,937	\$3.46	\$5.50	\$0.01	\$2.20	\$253.30
$\mathbb{E}^a_t[\text{EPS}_{n,\tau+2}]$	1,302,001	\$4.22	\$6.91	\$0.01	\$2.65	\$387.61
$\mathbb{E}^a_t[\text{EPS}_n]$	2,061,108	\$3.73	\$6.16	\$0.01	\$2.33	\$387.61
$\text{ImpliedPE}^a_{n,t}$	1,900,758	18.4×	8.3×	5.0×	16.4×	50.0×
$\text{TrailingPE}_{n,t}$	1,745,571	19.7×	8.8×	5.0×	17.9×	50.0×

Table 9. *Summary statistics at the firm-analyst-month level from 2003 to 2022.* $\text{PriceTarget}^a_{n,t}$*: price forecast set for the end of a firm's upcoming fiscal year, roughly twelve months in the future.* $\mathbb{E}^a_t[\text{EPS}_{n,\tau+1}]$*: analyst's EPS forecast for the twelve-month period ending on the date of their price target.* $\mathbb{E}^a_t[\text{EPS}_{n,\tau+2}]$*: analyst's EPS forecast for the twelve-month period following the date of their price target.* $\mathbb{E}^a_t[\text{EPS}_n]$*: an analyst's two-year-ahead EPS forecast when available; else, the reported one-year-ahead forecast value.* $\text{ImpliedPE}^a_{n,t}$*: the analyst's price target divided by their EPS forecast.* $\text{TrailingPE}_{n,t}$*: trailing twelve-month P/E ratio, calculated as a company's previous closing price divided by the sum of its last four quarterly EPS announcements.*

observations with a positive EPS forecast, $\mathbb{E}^a_t[\text{EPS}_n] \geq \$0.01$. We also require firms to have a price target greater than \$1/sh and less than \$10,000/sh.

The resulting panel data set is organized by firm × analyst × target date. Figures 3 and 6 show what this panel looks like Chris Horvers' coverage of Home Depot and Andrea Teixeira's coverage of Coca-Cola (KO). Figures B1(a)-B1(n) in Appendix B provide additional examples.

Let $\text{TrailingPE}_{n,t}$ denote a company's trailing twelve-month P/E ratio—i.e., its closing price on the previous trading day divided by the sum of its last four quarterly EPS announcements. We write the implied P/E ratio as follows

$$\text{ImpliedPE}^a_{n,t} \overset{\text{def}}{=} \frac{\text{PriceTarget}^a_{n,t}}{\mathbb{E}^a_t[\text{EPS}_n]} \tag{16}$$

Figure 17 shows the distribution of both $\text{ImpliedPE}^a_{n,t}$ and $\text{TrailingPE}_{n,t}$ across firms. We restrict our sample to only include observations where both

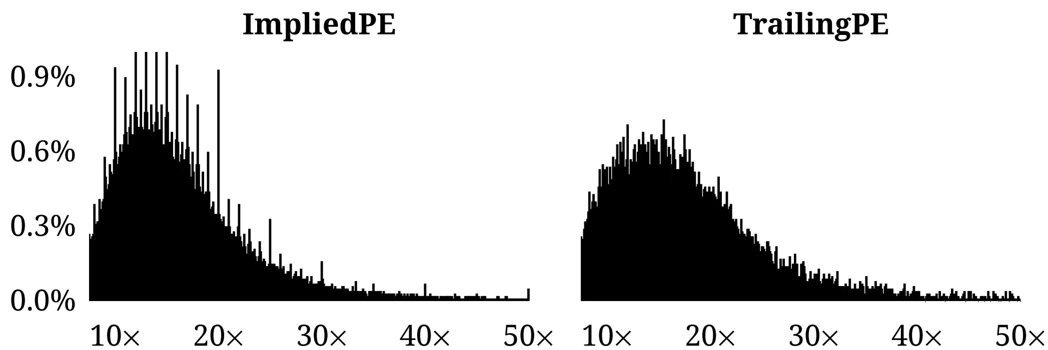

Figure 17. *Histograms showing the distribution of* $\mathrm{ImpliedPE}^a_{n,t}$ *(left panel) and* $\mathrm{TrailingPE}_{n,t}$ *(right panel) for all sell-side analyst reports in our sample with* $\mathbb{E}^a_t[\mathrm{EPS}_n] \geq \$1.00/\mathrm{sh}$ *from 2003 to 2022. x-axis denotes the P/E ratio in increments of* $0.1\times$. *y-axis represents the share of all observations that belong to that bin.*

these P/E ratios are between 5× and 50×. This choice is motivated by practical considerations: market participants often see P/E ratios outside of this range as extreme. In such situations, analysts usually apply an alternative valuation method, such as EV/EBITDA. However, we show in Appendix B Figures B2(a)-B2(e) that our findings extend outside this range.

3.2 Analysts' Price Targets

In the first part of our analysis, we read the text of 513 analyst reports and found that most analysts described using the formula, PriceTarget = $\mathbb{E}$[EPS] × TrailingPE. We now show that the simplest possible version of this approach explains the majority of the price-target variation in the full IBES sample.

We start by fitting the regression specification below to IBES data on days when analyst a updated their price target for the nth firm

$$\begin{aligned}\log(\mathrm{PriceTarget}^a_{n,t}) \overset{\text{OLS}}{\sim} \hat{\alpha} + \hat{\beta} \cdot \log(\mathbb{E}^a_t[\mathrm{EPS}_n]) \\ + \hat{\gamma} \cdot \log(\mathrm{TrailingPE}_{n,t})\end{aligned} \tag{17}$$

For example, when looking at Andrea Teixeira's coverage of Coca-Cola, we used the trading days with black diamonds in Figure 6. $\log(\mathrm{PriceTarget}^a_{n,t})$ is the log

Dep variable:	$\log(\text{PriceTarget}^a_{n,t})$			
	(1)	(2)	(3)	(4)
$\log(\mathbb{E}^a_t[\text{EPS}_n])$	0.93***	0.87***	0.91***	0.93***
	(0.01)	(0.01)	(0.01)	(0.01)
$\log(\text{TrailingPE}_{n,t})$	0.63***	0.47***	0.64***	0.57***
	(0.01)	(0.01)	(0.01)	(0.01)
Firm FE		Y		
Analyst FE			Y	
Month FE				Y
Adj. R^2	91.0%	93.6%	91.4%	92.4%
# Obs	1,666,655	1,666,587	1,666,655	1,666,449

Table 10. *Each column reports the results of a separate regression of the form found in Equation* (17)*. All regressions use the same underlying panel data set. Each panel represents a sequence of price targets and earnings forecasts made by analyst a about firm n prior to target date* $(\tau + 1)$*. We study the time window between* 18 *and* 6 *months prior to the end of a firm's fiscal year. We do not report the intercept or fixed-effect coefficients. Numbers in parentheses are standard errors clustered three ways by firm, analyst, and month. Sample: 2003 to 2022.*

of the analyst's price target, $\log(\mathbb{E}^a_t[\text{EPS}_n])$ is the log of the analyst's earnings forecast, and $\log(\text{TrailingPE}_{n,t})$ is the log of the firm's P/E ratio during the twelve months prior to day t when the analyst's report was published.

If sell-side analysts exclusively used the formula $\text{PriceTarget} = \mathbb{E}[\text{EPS}] \times \text{TrailingPE}$ and always calculated a trailing twelve-month P/E ratio, then we would estimate coefficients of $\beta = 1$ and $\gamma = 1$ with an $R^2 = 100\%$. Column (1) in Table 10 shows that this is a good first approximation to reality. We estimate $\hat{\beta} = 0.93(\pm 0.01)$ and $\hat{\gamma} = 0.63(\pm 0.01)$. We get minuscule standard errors even though we cluster in three different ways: by firm, by analyst, and by month.

Our simple trailing P/E formula using just the last twelve months of data generates an adjusted $R^2 = 91.0\%$. It explains all but 9% of the data without requiring additional fine-tuning. Columns (2)-(4) in Table 10 show that firm, analyst, and month fixed-effects do not add much.

We have found that binned scatterplots do a much better job of conveying the tight fit between theory and data. Asset-pricing researchers are used to seeing R^2s in the low single digits (Campbell and Thompson, 2008; Welch and

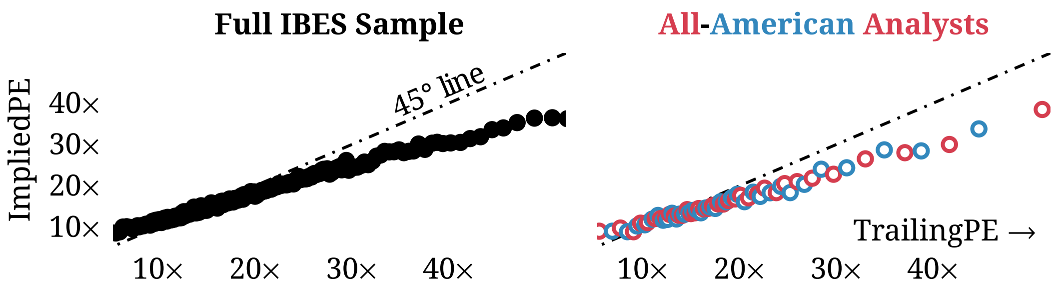

Figure 18. *(Left) Binned scatterplot using data from the full sample of IBES reports. x-axis shows the firm's trailing twelve-month P/E,* $\text{TrailingPE}_{n,t} = \text{Price}_{n,t} \,/\, \text{EPS}_{n,t}$. *y-axis shows the P/E ratio implied by the analyst's price target and EPS forecast,* $\text{ImpliedPE}_{n,t}^{a} \overset{\text{def}}{=} \text{PriceTarget}_{n,t}^{a} \,/\, \mathbb{E}_t^a[\text{EPS}_n]$. *(Right) Analogous binned scatterplot using data from the* 28 *analysts in Table 3 who have been named to Institutional Investor magazine's All-America research team. Sample: 2003 to 2022.*

Goyal, 2008). Many have a hard time appreciating what $R^2 = 91.0\%$ really means. At the very least, we know of two asset-pricing researchers whose first instinct was to ask questions about the remaining 9% our story does not explain.

The left panel of Figure 18 depicts the relationship between the P/E ratio implied by an analyst's price target and EPS forecast ($\text{ImpliedPE}_{n,t}^{a}$; y-axis) and a company's trailing twelve-month P/E ratio ($\text{TrailingPE}_{n,t}$; x-axis). If sell-side analysts set price targets using nothing but $\text{PriceTarget} = \mathbb{E}[\text{EPS}] \times \text{TrailingPE}$ and only ever calculated a trailing twelve-month P/E, then all the dots should sit up on the 45° line. The empirical best-fit line is a bit flatter, but there is no mistaking that it is a line. This is what it looks like when a simple linear model explains most of the observed variation in the data.

The right panel of Figure 18 performs the same analysis using reports written by the 28 analysts in Table 3 who were named to *Institutional Investor* magazine's All-America research team. The only thing separating the results in the left and right panels is the color scheme. Figures B2(a)-B2(e) in Appendix B show similar binned scatterplots using the data on 100 large publicly traded companies. We find that the same linear relationship holds when fitting a separate regression to data on each individual company. It is possible to count the number of exceptions on one hand.

Dep variable:	$\text{ImpliedPE}^a_{n,t}$			
	(1)	(2)	(3)	(4)
$\text{TrailingPE}_{n,t}$	0.58^{***} (0.01)	0.43^{***} (0.01)	0.58^{***} (0.01)	0.52^{***} (0.01)
Firm FE		Y		
Analyst FE			Y	
Month FE				Y
Adj. R^2	54.5%	67.7%	55.8%	61.5%
# Obs	1,646,279	1,646,207	1,646,279	1,646,077

Table 11. *Each column reports the results of a separate regression of the form found in Equation* (18)*. All regressions use the same underlying panel data set. Each panel represents a sequence of price targets and earnings forecasts made by analyst a about firm n prior to target date* $(\tau + 1)$*. We study the time window between* 18 *and* 6 *months prior to the end of a firm's fiscal year. We do not report the intercept or fixed-effect coefficients. Numbers in parentheses are standard errors clustered three ways by firm, analyst, and month. Sample: 2003 to 2022.*

We quantify the relationship between $\text{ImpliedPE}^a_{n,t}$ and $\text{TrailingPE}_{n,t}$ using regressions in Table 11. Just like before, each column shows the results of estimating a variation on the same underlying regression specification

$$\text{ImpliedPE}^a_{n,t} \overset{\text{OLS}}{\sim} \hat{\eta} + \hat{\theta} \cdot \text{TrailingPE}_{n,t} \tag{18}$$

If sell-side analysts were exclusively using trailing twelve-month P/Es to set price targets based on the formula $\text{PriceTarget} = \mathbb{E}[\text{EPS}] \times \text{TrailingPE}$, then we should estimate a coefficient of $\theta = 1$. Instead, in column (1) we estimate a value of $\hat{\theta} = 0.58(\pm 0.01)$ with an adjusted $R^2 = 54.5\%$. The best-fit line may be a bitter flatter than predicted, but it still explains more than half of the variation.

Why is the fit not perfect? We can think of a few reasons. First, analysts are not automatons. They elaborate on $\text{PriceTarget} = \mathbb{E}[\text{EPS}] \times \text{TrailingPE}$ when it seems like other information might be relevant. This simple formula is a starting point. Analysts often add ingredients when a company's trailing P/E is particularly extreme in either direction.

Second, analysts often set price targets based on round P/E ratios. Notice all the spikes in the left panel of Figure 17, showing the cross-sectional distribution of $\text{ImpliedPE}^{a}_{n,t}$. When a company's current price is 19.9× its earnings over the past twelve months, an analyst will likely use a 20× trailing P/E.

Third, not every analyst calculates a firm's trailing P/E in the same way. The $\text{TrailingPE}_{n,t}$ variable in our regressions is the firm's P/E ratio over the past twelve months. But some analysts use a longer trailing window. For example, we saw in Figure 2 that Chris Horvers used a three-year trailing average P/E to set his price target for Home Depot in October 2019.

We can see from Figure 19 that the most recent four quarters of earnings have the largest effect on implied P/E ratios. But there are also significant coefficients at longer lags. We are able to explain $R^2 = 91\%$ of the variation in analysts' price targets even before incorporating the effects of longer-term lags.

We created this figure by regressing an analyst's implied P/E ratio on the company's realized P/E in each of the last 20 quarters

$$\text{ImpliedPE}^{a}_{n,q} \overset{\text{OLS}}{\sim} \hat{\eta} + \sum_{\ell=1}^{20} \hat{\theta}_{\ell} \cdot \text{QuarterlyPE}_{n,q-\ell} \tag{19}$$

We use $\text{eps}_{n,q}$ to denote the nth stock's earnings in quarter q. The variable $\text{QuarterlyPE}_{n,q} \overset{\text{def}}{=} \text{Price}_{n,t} \,/\, (4 \cdot \text{eps}_{n,q})$ represents the company's closing price on the day before its earnings for the quarter were announced announcement divided by four times its realized EPS in the quarter.

The estimated coefficients for lags one through four sum to $\sum_{\ell=1}^{4} \hat{\theta}_{\ell} = 0.21 + 0.12 + 0.09 + 0.08 = 0.50$. This total is slightly less than the slope coefficient in column (1) of Table 11, $\hat{\theta} = 0.58$, which suggests that analysts incorporate trailing information from previous years when such information is available. We only require 4 quarters of trailing EPS data when estimating Table 11; whereas, Figure 19 requires 20 quarters of trailing EPS data.

The results in this subsection are not intended to rule out other stories. The goal is to show that our main takeaway from reading 513 analyst reports is also true in the broader IBES sample. Sell-side analysts typically say they set price targets using the formula, $\text{PriceTarget} = \mathbb{E}[\text{EPS}] \times \text{TrailingPE}$. Price targets in IBES are consistent with this claim.

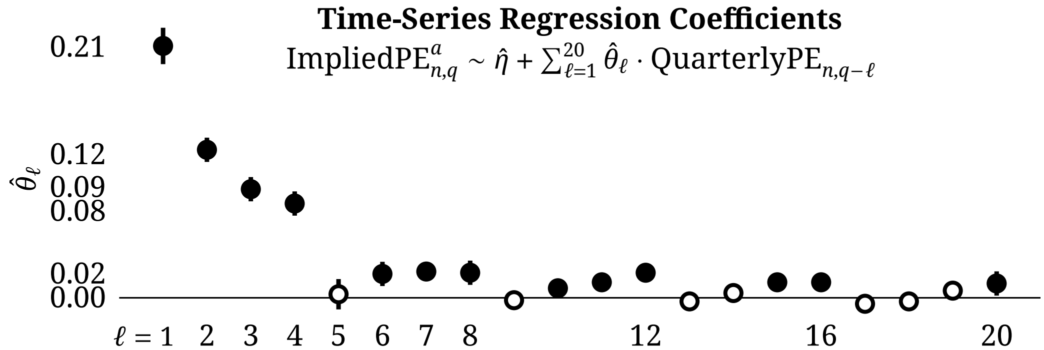

Figure 19. *Each dot denotes one of the* 20 *estimated slope coefficients,* $\{\hat{\theta}_\ell\}_{\ell=1}^{20}$, *from the regression specification in Equation* (19). $\text{ImpliedPE}^a_{n,q}$*: P/E ratio implied by an analyst's price target and EPS forecast.* $\text{QuarterlyPE}_{n,q}$*: company's closing price the day before the announcement divided by four times its realized EPS in quarter q. Vertical lines denote* 99% *confidence intervals using standard errors clustered three ways by firm, analyst, and month. White dots denote insignificant coefficient estimates. Sample: 2003q1 to 2022q4.*

3.3 Realized Price Changes

Researchers have spent decades learning about discount rates from IBES data. Our main finding shows that the analysts responsible for these numbers do not typically use a discount rate to price them. This finding would have important implications for asset-pricing researchers even in a world where analysts were completely infra-marginal and their price targets did not affect on market prices.

However, we now provide evidence that markets react to earnings surprises in a way that is consistent with our simple model. The exclusion restriction in Proposition 2.3 implies that earnings surprises should predict future returns through their effect on analysts' short-term EPS forecasts. These events cannot change a firm's trailing P/E ratio, which has already been determined.

Of course, researchers disagree about how to best model changes in analysts' subjective beliefs. There is an active literature on this topic (Bordalo, Gennaioli, La Porta, and Shleifer, 2019; Bouchaud, Krueger, Landier, and Thesmar, 2019; de Silva and Thesmar, 2024). So our challenge is to show that the market reacts to earnings surprises through changes in analysts' short-term EPS forecasts and not due to multiples expansion/contraction while making minimal assumptions about analysts' belief-formation process.

We do this by exploiting the linearity of our problem with a two-stage procedure similar to Fama and MacBeth (1973). Asset-pricing researchers are comfortable checking for priced risk without specifying what the price will be beforehand. We want to check whether market reactions to earnings surprise are driven by updates to $\mathbb{E}[\text{EPS}]$ without specifying analysts' updating rule.

Let $\$\text{surprise}_{n,q}$ denote the difference between the nth firm's realized earnings per share for quarter q and the consensus forecast

$$\$\text{surprise}_{n,q} \overset{\text{def}}{=} \text{eps}_{n,q} - \mathbb{E}_t[\text{eps}_{n,q}] \tag{20}$$

where $\mathbb{E}_t[\cdot]$ denotes analysts' consensus on the day before the firm announced its earnings for the quarter. When $\$\text{surprise}_{n,q} \neq \$0.00/\text{sh}$, we say that the nth firm has experienced an earnings surprise in quarter q.

It makes sense that analysts would revise their short-term EPS forecast following an earnings surprise. Suppose that the revision is proportional to the size of surprise for some $\lambda > 0$

$$\mathbb{E}_t[\text{EPS}_{n,\tau+2} \,|\, \$\text{surprise}_{n,q} = s] \approx \mathbb{E}_t[\text{EPS}_{n,\tau+2}] + \lambda \cdot s \tag{21}$$

For example, Equation (21) would hold exactly with $\lambda = \frac{1^2}{1^2 + \phi^2/\delta^2}$ if the company's short-term EPS obeyed $\text{EPS}_{n,\tau+2} \sim \text{Normal}(\mathbb{E}_t[\text{EPS}_{n,\tau+2}], 1^2)$ and analysts saw the earnings surprise as a noisy signal $\$\text{surprise}_{n,q} = \delta \cdot \{\text{EPS}_{n,\tau+2} - \mathbb{E}_t[\text{EPS}_{n,\tau+2}]\} + \text{Noise}_{n,q}$ with $\text{Noise}_{n,q} \overset{\text{IID}}{\sim} \text{Normal}(0, \phi^2)$ for $0 < \delta \leq 1$ and $\phi > 0$.

While there are many ways to justify this modeling choice, the economic content of Proposition 2.3 does not depend on precisely how analysts update their beliefs. The key point is that analysts set a price target using a trailing P/E ratio. When analysts revise their short-term EPS forecast, they should continue to capitalize their new beliefs into a price target with the same trailing P/E as before. An earnings surprise should only impact analysts' price target (and thus future prices) by altering analysts' $\mathbb{E}_t[\text{EPS}_{n,\tau+2}]$.

Notice that, given the reduced-form linear updating rule in Equation (21), differences in analysts' price targets across firms with the same size earnings surprise will be proportional to each firm's trailing P/E

$$\Delta\text{PriceTarget}_{n,t} = (\lambda \cdot s) \times \text{TrailingPE}_{n,t} \qquad \begin{array}{c}\text{among firms-quarters}\\ \text{where \$surprise}_{n,q} = s\end{array} \tag{22}$$

Of course, analysts' price forecasts are not the only thing that affects price growth. When analysts increase their price target by \$1/sh, a firm's share price will not typically increase by a full \$1 as well.

In our model, the slippage is governed by the demand multiplier $\mu > 0$ and the price elasticity $\nu > 0$. While these are free parameters, both are positive constants. So, whatever they are, we can include them in the constant of proportionality. This observation converts Equation (22) from being a claim about differences in price targets to being a claim about future price growth

$$\Delta\text{Price}_{n,q+1} = (\Lambda \cdot s) \times \text{TrailingPE}_{n,t} \qquad \begin{array}{c}\text{among firms-quarters}\\ \text{where \$surprise}_{n,q} = s\end{array} \tag{23}$$

We test this claim using a two-stage approach reminiscent of Fama and MacBeth (1973). First, we group stock-quarter observations into portfolios by the size of their surprise, $\text{\$surprise}_{n,q} = s \in \{-\$0.30/\text{sh}, \ldots, \$0.30/\text{sh}\}$. And, within each group, we run a separate first-stage regression

$$\Delta\text{Price}_{n,q+1} \overset{\text{OLS}}{\sim} \hat{a}_s + \hat{\beta}_s \cdot \text{TrailingPE}_{n,q} \qquad \begin{array}{c}\text{using data on firm-quarters}\\ \text{that have \$surprise}_{n,q} = s\end{array} \tag{24}$$

This gives us 60 different estimates of $\hat{\beta}_s = \Lambda \cdot s$, one for each group of firm-quarter observations with the same $\text{\$surprise}_{n,q} = s$.

To check if price changes following an earnings surprise are mainly driven by revisions to short-run EPS forecasts, we run a cross-sectional second-stage regression

$$\hat{\beta}_s \overset{\text{OLS}}{\sim} \bar{\alpha} + \bar{\Lambda} \cdot s \qquad \begin{array}{c}\text{using first-stage slope from each}\\ \text{bin } s \in \{-\$0.30/\text{sh}, \ldots, \$0.30/\text{sh}\}\end{array} \tag{25}$$

On the left-hand side is the first-stage slope coefficient associated with a specific dollar earnings surprise. On the right-hand side is the size of that surprise. If the exclusion restriction implied by our model is correct, then first-stage slope

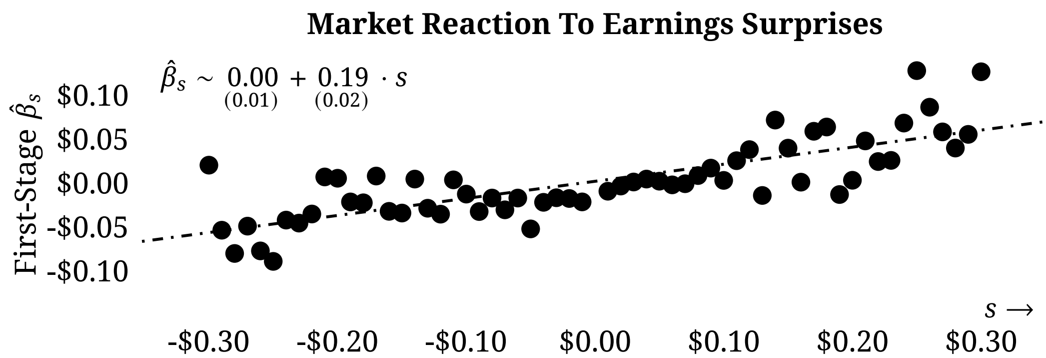

Figure 20. *The dashed line shows the best-fit OLS equation for the second-stage regression shown in Equation* (25). *The slope of this line is* $\bar{\Lambda} = 0.19(\pm 0.02)$. *The y-axis shows the estimated first-stage slope coefficients,* $\hat{\beta}_s$*, from the* 60 *separate regressions described by Equation* (24)*, each looking at a group of observations with the same size earnings surprise,* $\$\text{surprise}_{n,q} = s \in \{-\$0.30/\text{sh}, \ldots, \$0.30/\text{sh}\}$. *The x-axis shows the size of that earnings surprise, s, in* \$0.01/sh *bins. The highest bin is centered at* \$0.30/sh *while the lowest bin is centered at* −\$0.30/sh. *We omit the bin centered at* $s = \$0.00/\text{sh}$—*i.e., observations with no surprise.*

coefficients should be well-explained by a linear model with positive slope, $\Lambda > 0$, and zero intercept, $\alpha = 0$.

This is what Table 12 and Figure 20 show. We estimate $\bar{\Lambda} = 0.19$ and $\bar{\alpha} = 0.00$ with an Adj. $R^2 = 60.6\%$. The neat linear relationship suggests that trailing P/E ratios always have the same effect. No matter how large or small the earnings surprise, the subsequent price response is always proportional to the firm's trailing P/E at the time.

Consider two stocks, one with a 20× trailing P/E and the other with a 10× trailing P/E. Suppose both firms realize a \$0.10/sh earnings surprise. Our regression estimates imply that the price of the first stock will increase by $\$0.19/\text{sh} = 0.19 \cdot \{\$0.10/\text{sh} \times 20 - \$0.10/\text{sh} \times 10\}$ more than the price of the second over the following quarter. If both firms had realized a \$0.20/sh earnings surprise, then on average there would be a \$0.38/sh gap between their respective price changes the following quarter. A \$0.30/sh surprise will produce a difference of \$0.57/sh.

This is not how things would look in a world where earnings surprises lead to multiples expansion/contraction. Earnings surprises can contain information

Dep variable:	First-Stage $\hat{\beta}_s$		
Bin width:	\$0.01/sh	\$0.02/sh	\$0.05/sh
	(1)	(2)	(3)
Intercept, $\bar{\alpha}$	0.00 (0.01)	0.00 (0.01)	0.00 (0.01)
Slope, $\bar{\Lambda}$	0.19*** (0.02)	0.20*** (0.02)	0.16*** (0.03)
Adj. R^2	60.6%	76.0%	77.7%
# Bins	60	30	12

Table 12. *Each column reports the results of a separate second-stage regression as shown in Equation (25). The dependent variable in each regression is the estimated slope coefficient, $\hat{\beta}_s$, from a collection of first-stage regressions described by Equation (24), each looking at a group of observations with the same size earnings surprise. Column (1) reports results using* 60 *separate* \$0.01/sh *bins centered at* $\{-\$0.30/\text{sh}, \ldots, -\$0.01/\text{sh}, \$0.01/\text{sh}, \ldots, \$0.30/\text{sh}\}$. *These results match the dashed best-fit line in Figure 20. Column (2) shows results where we group observations into* 30 *separate* \$0.02/sh *bins centered at* $\{-\$0.30/\text{sh}, \ldots, -\$0.02/\text{sh}, \$0.02/\text{sh}, \ldots, \$0.30/\text{sh}\}$. *Column (3) shows a similar analysis using* 12 *bins that are* \$0.05/sh *wide,* $\{-\$0.30/\text{sh}, \ldots, -\$0.05/\text{sh}, \$0.05/\text{sh}, \ldots, \$0.30/\text{sh}\}$. *All three columns omit the bin centered at $s = \$0.00/\text{sh}$—i.e., stock-quarter observations where there was no earnings surprise.*

about a company's long-run EPS growth rate, g. In a Gordon model, news that increased this parameter would cause analysts to use a larger P/E, $\left(\frac{1}{r-g}\right)$. However, there is no evidence of systematic repricing via multiples expansion/contraction due to learning about g. The relationship between a firm's trailing P/E and its subsequent price growth does not change with the size of the surprise.

We know from cross-sectional asset pricing that a researcher's choice of test portfolios can affect how well a model appears to fit the data (Lewellen, Nagel, and Shanken, 2010). So, in Table 12, we show the results of analogous second-stage regressions where we group firm-quarter observations into bins that are \$0.02/sh wide and \$0.05/sh wide. We get quantitatively similar results no matter how finely we divide our portfolios. The intercept is always a precisely estimated zero. This straight line exists because investors are using a trailing P/E ratio to update a firm's price following earnings surprises.

It is also important to emphasize that we are analyzing dollar price changes, not quarterly returns as is standard in the literature. We are measuring the size of a company's earnings surprise in units of dollars per share, not as a percent of its share price. The key dependent variable is a first-stage estimate capturing the relationship between subsequent price changes and trailing P/Es among firm-quarters with the same size surprise. We do not know of any other papers linking this particular set of variables in these non-standard units.

Conclusion

Most market participants do not share their subjective payoff expectations with us. Sell-side analysts are the exception. As a result, their numerical forecasts have had a massive impact on the asset-pricing literature. However, analysts also share how they price their subjective payoff expectations.

In this paper, when we read the text of analyst reports and find that they do not usually discount anything. Instead, analysts typically rely on trailing P/E ratios. Agents in textbook models ask: "What is the present discounted value of a company's expected future earnings stream in today's dollars?" Real-world analysts ask themselves: "How would a comparable firm have been priced last year if it had announced similar earnings?"

We are not making a blanket claim about how every investor values every asset. We recognize that investors do sometimes set price equal to expected discounted payoff. For example, bond markets are largely governed by present-value logic. Investors also rely on DCF models in certain niche industries. This is how marine shipping MLPs, mining operations, and REITs typically get valued.

However, asset-pricing researchers currently take it for granted that investors always enforce present-value relationships. This is simply not true. "Asset-pricing theory all stems from one simple concept: price equals expected discounted payoff. (Cochrane, 2009, page 1)" That is clearly not a sensible way to model sell-side analysts. The next time you write down a model, you should ask if it makes sense for the agents in your application.

References

Adam, K., A. Marcet, and J. Beutel (2017). Stock price booms and expected capital gains. *American Economic Review 107*(8), 2352–2408.

Adam, K., D. Matveev, and S. Nagel (2021). Do survey expectations of stock returns reflect risk adjustments? *Journal of Monetary Economics 117*, 723–740.

Afrouzi, H., S. Kwon, A. Landier, Y. Ma, and D. Thesmar (2023). Overreaction in expectations: Evidence and theory. *Quarterly Journal of Economics 138*(3), 1713–1764.

Angrist, J. and J.-S. Pischke (2009). *Mostly harmless econometrics: An empiricist's companion*. Princeton University Press.

Argus Research (2011, Nov). Chrevron Corp [CVX]. Technical report, Philip Weiss.

Bartram, S. and M. Grinblatt (2018). Agnostic fundamental analysis works. *Journal of Financial Economics 128*(1), 125–147.

Basu, S. (1983). The relationship between earnings' yield, market value and return for nyse common stocks: Further evidence. *Journal of Financial Economics 12*(1), 129–156.

Ben-David, I. and A. Chinco (2024). Modeling managers as EPS maximizers. Working paper.

Bhojraj, S. and C. Lee (2002). Who is my peer? a valuation-based approach to the selection of comparable firms. *Journal of Accounting Research 40*(2), 407–439.

Birru, J., S. Gokkaya, X. Liu, and R. Stulz (2022). Are analyst short-term trade ideas valuable? *Journal of Finance 77*(3), 1829–1875.

Bordalo, P., N. Gennaioli, R. La Porta, and A. Shleifer (2019). Diagnostic expectations and stock returns. *Journal of Finance 74*(6), 2839–2874.

Bordalo, P., N. Gennaioli, R. La Porta, and A. Shleifer (2020). Expectations of fundamentals and stock market puzzles. Working paper.

Bordalo, P., N. Gennaioli, R. La Porta, and A. Shleifer (2024). Finance without (exotic) risk. Working paper.

Bordalo, P., N. Gennaioli, Y. Ma, and A. Shleifer (2020). Overreaction in macroeconomic expectations. *American Economic Review 110*(9), 2748–2782.

Bouchaud, J.-P., P. Krueger, A. Landier, and D. Thesmar (2019). Sticky expectations and the profitability anomaly. *Journal of Finance 74*(2), 639–674.

Brav, A. and R. Lehavy (2003). An empirical analysis of analysts' target prices: Short-term informativeness and long-term dynamics. *Journal of Finance 58*(5), 1933–1967.

BTIG (2010, May). AT&T [T]. Technical report, Walter Piecyk.

Campbell, J. (2017). *Financial decisions and markets: a course in asset pricing*. Princeton University Press.

Campbell, J. and R. Shiller (1988a). The dividend-price ratio and expectations of future dividends and discount factors. *Review of Financial Studies 1*(3), 195–228.

Campbell, J. and R. Shiller (1988b). Stock prices, earnings, and expected dividends. *Journal of Finance 43*(3), 661–676.

Campbell, J. and S. Thompson (2008). Predicting excess stock returns out of sample: Can anything beat the historical average? *Review of Financial Studies 21*(4), 1509–1531.

CFA Institute (2017). *CFA Program Curriculum 2018 Level II*. John Wiley & Sons.

Charles, C., C. Frydman, and M. Kilic (2024). Insensitive investors. *Journal of Finance 79*(4), 2473–2503.

Chinco, A., S. Hartzmark, and A. Sussman (2022). A new test of risk factor relevance. *Journal of Finance 77*(4), 2183–2238.

Cochrane, J. (2009). *Asset Pricing*. Princeton University Press.

Cochrane, J. (2011). Presidential address: Discount rates. *Journal of finance 66*(4), 1047–1108.

Coibion, O. and Y. Gorodnichenko (2015). Information rigidity and the expectations formation process: A simple framework and new facts. *American Economic Review 105*(8), 2644–2678.

Cooper, I. and N. Lambertides (2023). Optimal equity valuation using multiples: The number of comparable firms. *European Financial Management 29*(4), 1025–1053.

Cready, W. and U. Gurun (2010). Aggregate market reaction to earnings announcements. *Journal of Accounting Research 48*(2), 289–334.

Credit Suisse (2004, Oct). Citigroup [C]. Technical report, Susan Roth.

Credit Suisse (2019, Oct). PepsiCo [PEP]. Technical report, Kaumil Gajrawala.

Da, Z. and E. Schaumburg (2011). Relative valuation and analyst target price forecasts. *Journal of Financial Markets 14*(1), 161–192.

De la O, R. and S. Myers (2021). Subjective cash flow and discount rate expectations. *Journal of Finance 76*(3), 1339–1387.

de Silva, T. and D. Thesmar (2024). Noise in expectations: Evidence from analyst forecasts. *Review of Financial Studies 37*(5), 1494–1537.

Décaire, P. and J. Graham (2024). Valuation fundamentals. Working paper.

Fama, E. and J. MacBeth (1973). Risk, return, and equilibrium: Empirical tests. *Journal of Political Economy 81*(3), 607–636.

Gabaix, X. and R. Koijen (2024). In search of the origins of financial fluctuations: The inelastic markets hypothesis. Working paper.

Gebhardt, W., C. Lee, and B. Swaminathan (2001). Toward an implied cost of capital. *Journal of Accounting Research 39*(1), 135–176.

Giglio, S., M. Maggiori, J. Stroebel, and S. Utkus (2021). Five facts about beliefs and portfolios. *American Economic Review 111*(5), 1481–1522.

Gormsen, N. and K. Huber (2024). Corporate discount rates. Working paper.

Green, J., J. Hand, and F. Zhang (2016). Errors and questionable judgments in analysts' DCF models. *Review of Accounting Studies 21*, 596–632.

Greenwood, R. and A. Shleifer (2014). Expectations of returns and expected returns. *Review of Financial Studies 27*(3), 714–746.

Grossman, S. and J. Stiglitz (1980). On the impossibility of informationally efficient markets. *American Economic Review 70*(3), 393–408.

Harbaugh, R., J. Maxwell, and K. Shue (2016). Consistent good news and inconsistent bad news. Working paper.

Hartzmark, S. and D. Solomon (2019). The dividend disconnect. *Journal of Finance 74*(5), 2153–2199.

JP Morgan (2010a, Apr). Avis Budget [CAR]. Technical report, Himanshu Patel.

JP Morgan (2010b, Dec). Pacific Biosciences [PACB]. Technical report, Tycho Peterson.

JP Morgan (2019a, Dec). Home Depot [HD]. Technical report, Christopher Horvers.

JP Morgan (2019b, Oct). PepsiCo [PEP]. Technical report, Andrea Teixeira.

JP Morgan (2019c, Oct). Coca-Cola [KO]. Technical report, Andrea Teixeira.

Kim, M. and J. Ritter (1999). Valuing IPOs. *Journal of Financial Economics 53*(3), 409–437.

Koijen, R. and M. Yogo (2019). A demand-system approach to asset pricing. *Journal of Political Economy 127*(4), 1475–1515.

Kothari, S., J. Lewellen, and J. Warner (2006). Stock returns, aggregate earnings surprises, and behavioral finance. *Journal of Financial Economics 79*(3), 537–568.

Kothari, S., E. So, and R. Verdi (2016). Analysts' forecasts and asset pricing: A survey. *Annual Review of Financial Economics 8*, 197–219.

La Porta, R. (1996). Expectations and the cross-section of stock returns. *Journal of Finance 51*(5), 1715–1742.

Lamont, O. (1998). Earnings and expected returns. *Journal of Finance 53*(5), 1563–1587.

Lewellen, J. (2004). Predicting returns with financial ratios. *Journal of Financial Economics 74*(2), 209–235.

Lewellen, J., S. Nagel, and J. Shanken (2010). A skeptical appraisal of asset-pricing tests. *Journal of Financial Economics 96*(2), 175–194.

Liu, J., D. Nissim, and J. Thomas (2002). Equity valuation using multiples. *Journal of Accounting Research 40*(1), 135–172.

Malmendier, U. and S. Nagel (2011). Depression babies: Do macroeconomic experiences affect risk taking? *Quarterly Journal of Economics 126*(1), 373–416.

McNichols, M. and P. O'Brien (1997). Self-selection and analyst coverage. *Journal of Accounting Research 35*, 167–199.

Mukhlynina, L. and K. Nyborg (2020). The choice of valuation techniques in practice: Education versus profession. *Critical Finance Review 9*, 201–265.

Murfin, J. and R. Pratt (2019). Comparables pricing. *Review of Financial Studies 32*(2), 688–737.

Petersen, M. (2008). Estimating standard errors in finance panel data sets: Comparing approaches. *Review of Financial Studies 22*(1), 435–480.

Purnanandam, A. and B. Swaminathan (2004). Are IPOs really underpriced? *Review of Financial Studies 17*(3), 811–848.

Rabin, M. (2013). An approach to incorporating psychology into economics. *American Economic Review 103*(3), 617–622.

RBC Capital Markets (2015, May). Chico's FAS Inc [CHS]. Technical report, Brian Tunick.

So, E. (2013). A new approach to predicting analyst forecast errors: Do investors overweight analyst forecasts? *Journal of Financial Economics 108*(3), 615–640.

Stickel, S. (1992). Reputation and performance among security analysts. *Journal of Finance 47*(5), 1811–1836.

Veldkamp, L. (2011). *Information Choice in Macroeconomics and Finance*. Princeton University Press.

Welch, I. and A. Goyal (2008). A comprehensive look at the empirical performance of equity premium prediction. *Review of Financial Studies 21*(4), 1455–1508.

Wells Fargo (2014, Jan). GasLog Ltd [GLOG]. Technical report, Michael Webber.

Wells Fargo (2019, Dec). Citigroup [C]. Technical report, Mike Mayo.

Wolfe Research (2019, Oct). Amazon [AMZN]. Technical report, Chris Bottiglieri.

A Technical Appendix

Proof. **(Proposition 2.2)** *Suppose that year-over-year price growth is governed by the law of motion in Equation (13). Then, if we take expectations under the objectively correct distribution, we will get*

$$\frac{\hat{\mathbb{E}}_t[\text{Price}_{t+1}] - \text{Price}_t}{\text{Price}_t} = \nu \times \left(\frac{\text{Demand}_{t+1} - \text{Demand}_t}{\text{Demand}_t} \right) \tag{A.1}$$

Note that investors choose their demand for the upcoming year $(t+1)$ at time t, so Demand_{t+1} *is not a random variable.*

We use the fact that investors proportionally adjust their portfolio holdings in response to changes in analysts' near-term earnings forecasts to rewrite things as

$$\frac{\hat{\mathbb{E}}_t[\text{Price}_{t+1}] - \text{Price}_t}{\text{Price}_t} = (\nu \cdot \mu) \times \left(\frac{\mathbb{E}_t[\text{EPS}_{t+2}] - \text{EPS}_t}{\text{EPS}_t} \right) \tag{A.2}$$

We now have an equation linking analysts' subjective EPS forecast to the firm's average price under the physical density that researchers can observe in the data.

From here, we rearrange things to express the firm's average price next year as analysts' near-term earnings forecast times a trailing P/E ratio plus some additional terms

$$\frac{\hat{\mathbb{E}}_t[\text{Price}_{t+1}] - \text{Price}_t}{\text{Price}_t} = (\nu \cdot \mu) \times \left(\frac{\mathbb{E}_t[\text{EPS}_{t+2}] - \text{EPS}_t}{\text{EPS}_t} \right) \tag{A.3a}$$

$$\frac{\hat{\mathbb{E}}_t[\text{Price}_{t+1}]}{\text{Price}_t} = (\nu \cdot \mu) \times \left(\frac{\mathbb{E}_t[\text{EPS}_{t+2}]}{\text{EPS}_t} \right) + (1 - \nu \cdot \mu) \tag{A.3b}$$

$$\begin{aligned}\hat{\mathbb{E}}_t[\text{Price}_{t+1}] = (\nu \cdot \mu) \times \mathbb{E}_t[\text{EPS}_{t+2}] \times \left(\frac{\text{Price}_t}{\text{EPS}_t} \right) \\ + (1 - \nu \cdot \mu) \times \text{Price}_t\end{aligned} \tag{A.3c}$$

By inspection, it is clear that the unwanted terms disappear if $\mu = 1/\nu$. □

Proof. **(Proposition 2.3)** *The left-hand side of Equation (A.3a) from the proof to Proposition 2.2 above is* $\hat{\mathbb{E}}_t[\text{Return}_{t+1}]$. *The right-hand side is a function of the analysts' expectations about short-term EPS as defined in Equation (9). $X_t \approx \mathbb{E}_t[\Delta \log \text{EPS}_{t+1}]$ is the expected rate at which the company's earnings will grow over the next year, and $\epsilon_{t+1} \overset{\text{IID}}{\sim} \text{Normal}(0, \sigma^2)$ is a noise term. Hence, a signal that is uncorrelated with X_t cannot explain differences in* $\hat{\mathbb{E}}_t[\text{Return}_{t+1}]$. □

B Additional Results

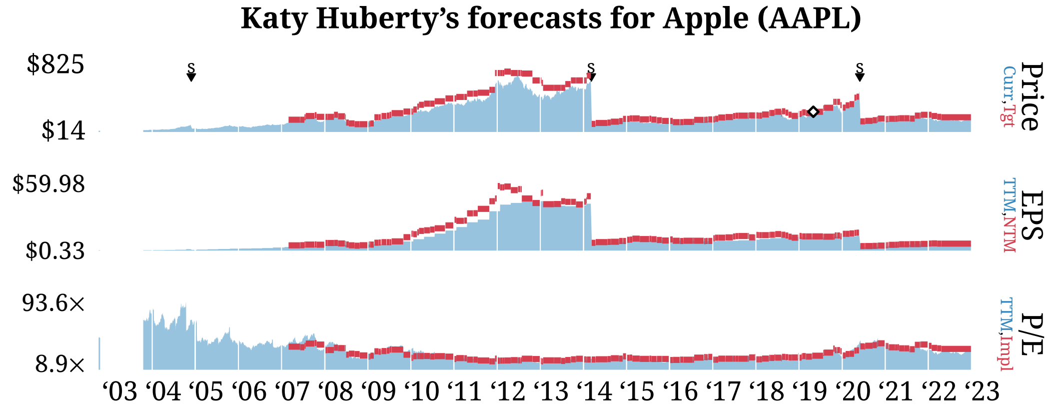

Figure B1(a). *y-axis shows min, median, and max. (Top) Blue ribbon is Apple's closing price on day t from CRSP,* Price_t*. Red line is Katy Huberty's price target,* $\text{PriceTarget}_t = \mathbb{E}_t[\text{Price}_{\tau+1}]$*, in IBES. (Middle) Blue is AAPL's trailing twelve-month (TTM) EPS on day t from IBES,* EPS_t*. Red is Katy Huberty's EPS forecast,* $\mathbb{E}_t[\text{EPS}]$*. (Bottom) Blue is AAPL's TTM P/E ratio,* $\text{TrailingPE}_t = \text{Price}_t \,/\, \text{EPS}_t$*. Red is the P/E implied by Katy Huberty's forecasts,* $\text{ImpliedPE}_t = \text{PriceTarget}_t \,/\, \mathbb{E}_t[\text{EPS}]$*. We flag split events with* $S_\blacktriangledown$ *pointers.*

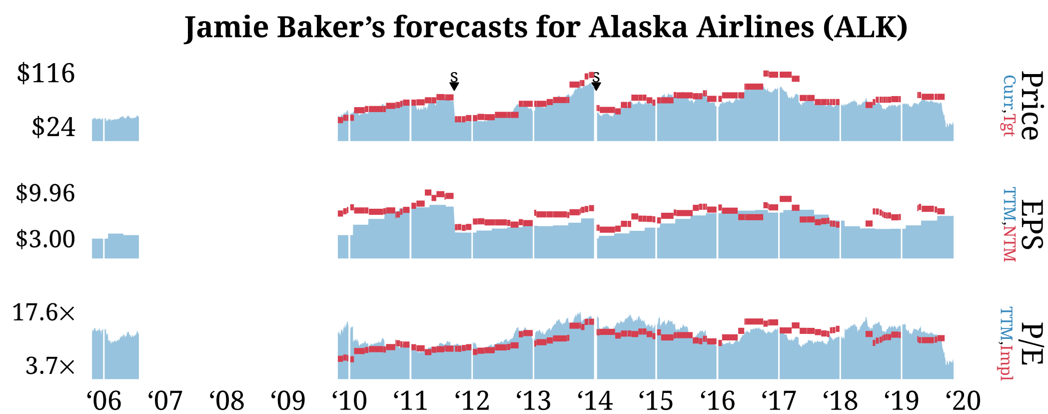

Figure B1(b). *y-axis shows min, median, and max. (Top) Blue ribbon is Alaska Airlines' closing price on day t from CRSP,* Price_t*. Red line is Jamie Baker's price target,* $\text{PriceTarget}_t = \mathbb{E}_t[\text{Price}_{\tau+1}]$*, in IBES. (Middle) Blue is ALK's trailing twelve-month (TTM) EPS on day t from IBES,* EPS_t*. Red is Jamie Baker's EPS forecast,* $\mathbb{E}_t[\text{EPS}]$*. (Bottom) Blue is ALK's TTM P/E ratio,* $\text{TrailingPE}_t = \text{Price}_t \,/\, \text{EPS}_t$*. Red is the P/E implied by Jamie Baker's forecasts,* $\text{ImpliedPE}_t = \text{PriceTarget}_t \,/\, \mathbb{E}_t[\text{EPS}]$*.*

CJ Muse's forecasts for Broadcom (AVGO)

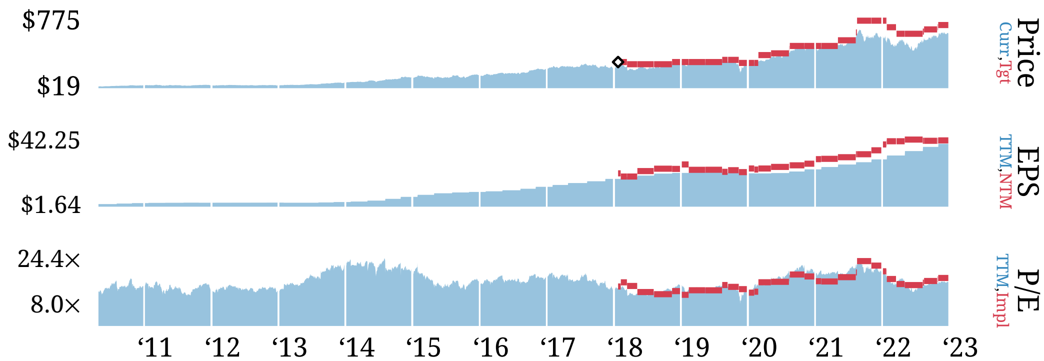

Figure B1(c). *y-axis shows min, median, and max. (Top) Blue ribbon is Broadcom's closing price on day t from CRSP,* Price_t. *Red line is CJ Muse's price target,* $\text{PriceTarget}_t = \mathbb{E}_t[\text{Price}_{\tau+1}]$, *in IBES. (Middle) Blue is Broadcom's trailing twelve-month (TTM) EPS on day t from IBES,* EPS_t. *Red is CJ Muse's EPS forecast,* $\mathbb{E}_t[\text{EPS}]$. *(Bottom) Blue is Broadcom's TTM P/E ratio,* $\text{TrailingPE}_t = \text{Price}_t \,/\, \text{EPS}_t$. *Red is the P/E implied by CJ Muse's forecasts,* $\text{ImpliedPE}_t = \text{PriceTarget}_t \,/\, \mathbb{E}_t[\text{EPS}]$. *We flag split events with* $S_\blacktriangledown$ *pointers.*

Vivek Juneja's forecasts for Citigroup (C)

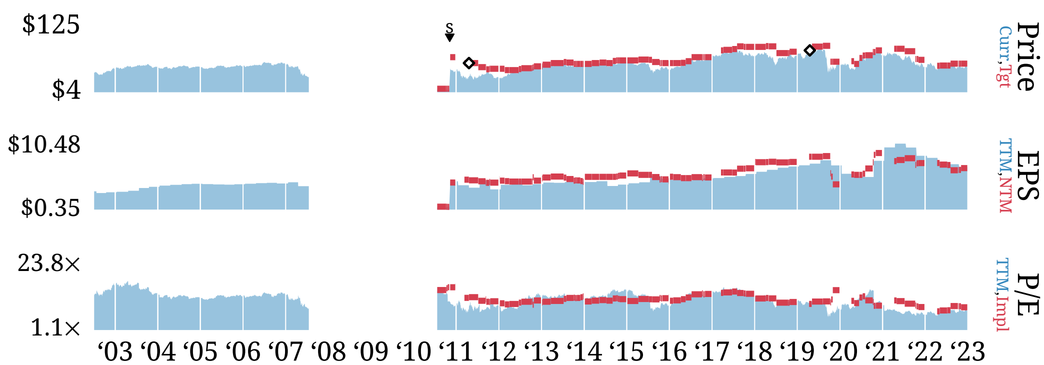

Figure B1(d). *y-axis shows min, median, and max. (Top) Blue ribbon is Citigroup's closing price on day t from CRSP,* Price_t. *Red line is Vivek Juneja's price target,* $\text{PriceTarget}_t = \mathbb{E}_t[\text{Price}_{\tau+1}]$, *in IBES. (Middle) Blue is Citi's trailing twelve-month (TTM) EPS on day t from IBES,* EPS_t. *Red is Vivek Juneja's EPS forecast,* $\mathbb{E}_t[\text{EPS}]$. *(Bottom) Blue is Citi's TTM P/E ratio,* $\text{TrailingPE}_t = \text{Price}_t \,/\, \text{EPS}_t$. *Red is the P/E implied by Vivek Juneja's forecasts,* $\text{ImpliedPE}_t = \text{PriceTarget}_t \,/\, \mathbb{E}_t[\text{EPS}]$.

Rod Hall's forecasts for Cisco Systems (CSCO)

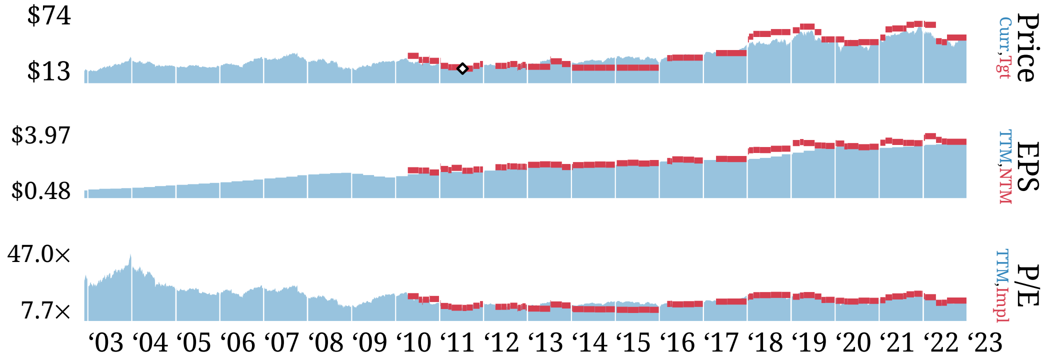

Figure B1(e). *y-axis shows min, median, and max. (Top) Blue ribbon is Cisco System's closing price on day t from CRSP,* Price_t*. Red line is Rod Hall's price target,* $\text{PriceTarget}_t = \mathbb{E}_t[\text{Price}_{\tau+1}]$*, in IBES. (Middle) Blue is Cisco's trailing twelve-month (TTM) EPS on day t from IBES,* EPS_t*. Red is Rod Hall's EPS forecast,* $\mathbb{E}_t[\text{EPS}]$*. (Bottom) Blue is Cisco's TTM P/E ratio,* $\text{TrailingPE}_t = \text{Price}_t \,/\, \text{EPS}_t$*. Red is the P/E implied by Rod Hall's forecasts,* $\text{ImpliedPE}_t = \text{PriceTarget}_t \,/\, \mathbb{E}_t[\text{EPS}]$*. We flag split events with* $S_\blacktriangledown$ *pointers.*

Heath Terry's forecasts for eBay (EBAY)

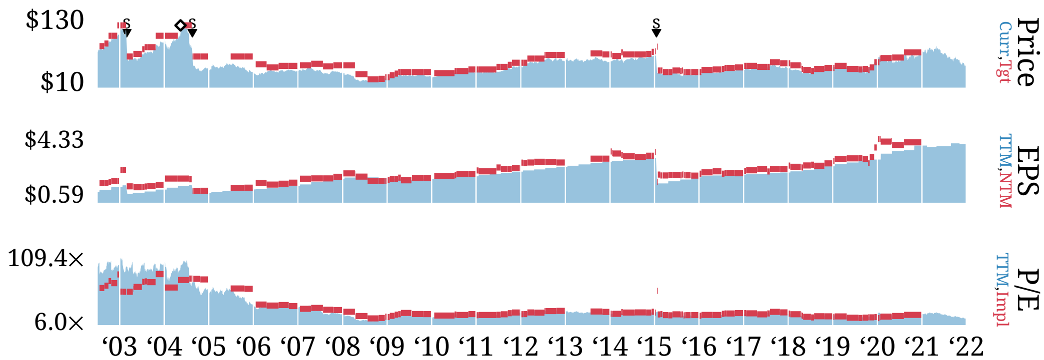

Figure B1(f). *y-axis shows min, median, and max. (Top) Blue ribbon is eBay's closing price on day t from CRSP,* Price_t*. Red line is Heath Terry's price target,* $\text{PriceTarget}_t = \mathbb{E}_t[\text{Price}_{\tau+1}]$*, in IBES. (Middle) Blue is eBay's trailing twelve-month (TTM) EPS on day t from IBES,* EPS_t*. Red is Heath Terry's EPS forecast,* $\mathbb{E}_t[\text{EPS}]$*. (Bottom) Blue is eBay's TTM P/E ratio,* $\text{TrailingPE}_t = \text{Price}_t \,/\, \text{EPS}_t$*. Red is the P/E implied by Heath Terry's forecasts,* $\text{ImpliedPE}_t = \text{PriceTarget}_t \,/\, \mathbb{E}_t[\text{EPS}]$.

Christophe Cherblanc's forecasts for Google (GOOGL)

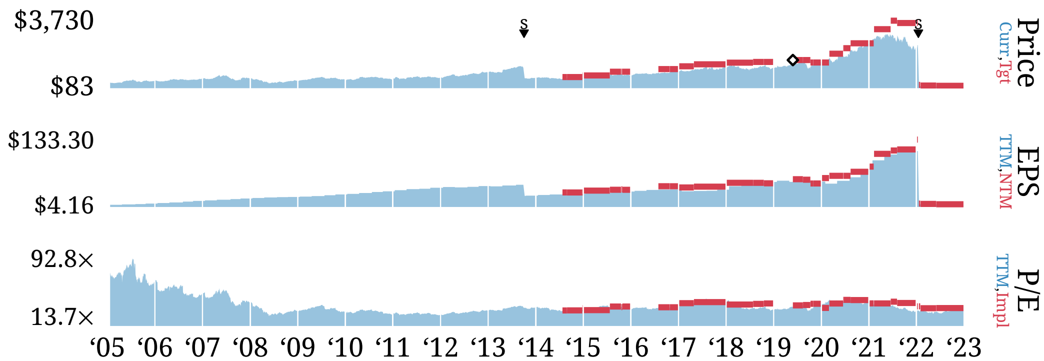

Figure B1(g). *y-axis shows min, median, and max. (Top) Blue ribbon is Google's closing price on day t from CRSP,* Price_t. *Red line is Christophe Cherblanc's price target,* $\text{PriceTarget}_t = \mathbb{E}_t[\text{Price}_{\tau+1}]$, *in IBES. (Middle) Blue is Google's trailing twelve-month (TTM) EPS on day t from IBES,* EPS_t. *Red is Cherblanc's EPS forecast,* $\mathbb{E}_t[\text{EPS}]$. *(Bottom) Blue is Google's TTM P/E ratio,* $\text{TrailingPE}_t = \text{Price}_t \,/\, \text{EPS}_t$. *Red is the P/E implied by Cherblanc's forecasts,* $\text{ImpliedPE}_t = \text{PriceTarget}_t \,/\, \mathbb{E}_t[\text{EPS}]$. *We flag split events with* $S_{\blacktriangledown}$ *pointers.*

Greg Badishkanian's forecasts for MarineMax (HZO)

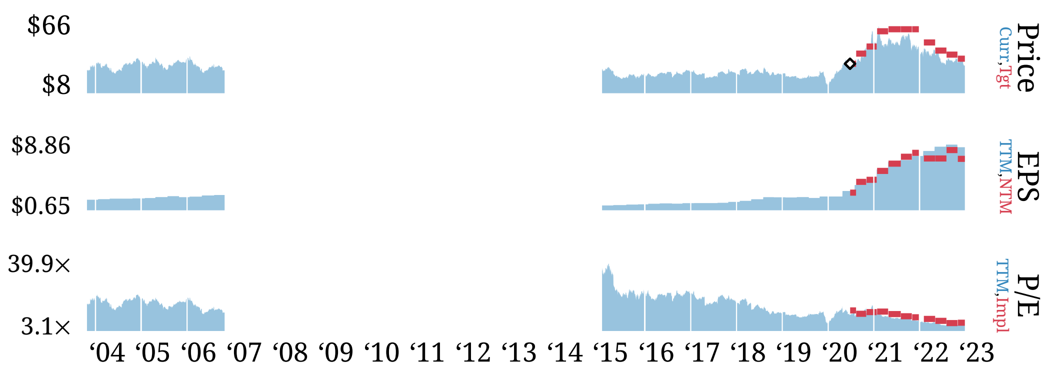

Figure B1(h). *y-axis shows min, median, and max. (Top) Blue ribbon is MarineMax's closing price on day t from CRSP,* Price_t. *Red line is Greg Badishkanian's price target,* $\text{PriceTarget}_t = \mathbb{E}_t[\text{Price}_{\tau+1}]$, *in IBES. (Middle) Blue is HZO's trailing twelve-month (TTM) EPS on day t from IBES,* EPS_t. *Red is Badishkanian's EPS forecast,* $\mathbb{E}_t[\text{EPS}]$. *(Bottom) Blue is HZO's TTM P/E ratio,* $\text{TrailingPE}_t = \text{Price}_t \,/\, \text{EPS}_t$. *Red is the P/E implied by Badishkanian's forecasts,* $\text{ImpliedPE}_t = \text{PriceTarget}_t \,/\, \mathbb{E}_t[\text{EPS}]$.

Moshe Orenbuch's forecasts for Mastercard (MA)

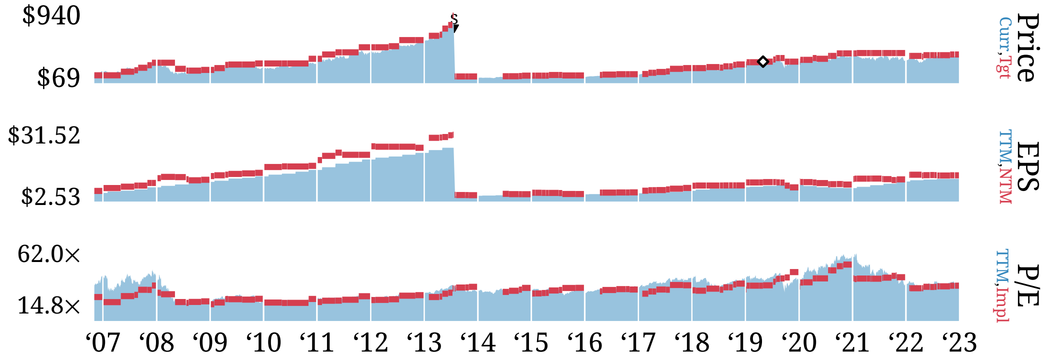

Figure B1(i). *y-axis shows min, median, and max. (Top) Blue ribbon is Mastercard's closing price on day t from CRSP,* Price_t*. Red line is Moshe Orenbuch's price target,* $\text{PriceTarget}_t = \mathbb{E}_t[\text{Price}_{\tau+1}]$*, in IBES. (Middle) Blue is MA's trailing twelve-month (TTM) EPS on day t from IBES,* EPS_t*. Red is Orenbuch's EPS forecast,* $\mathbb{E}_t[\text{EPS}]$*. (Bottom) Blue is MA's TTM P/E ratio,* $\text{TrailingPE}_t = \text{Price}_t \,/\, \text{EPS}_t$*. Red is the P/E implied by Orenbuch's forecasts,* $\text{ImpliedPE}_t = \text{PriceTarget}_t \,/\, \mathbb{E}_t[\text{EPS}]$*. We flag split events with* $S_\blacktriangledown$ *pointers.*

Tim Anderson's forecasts for Merck (MRK)

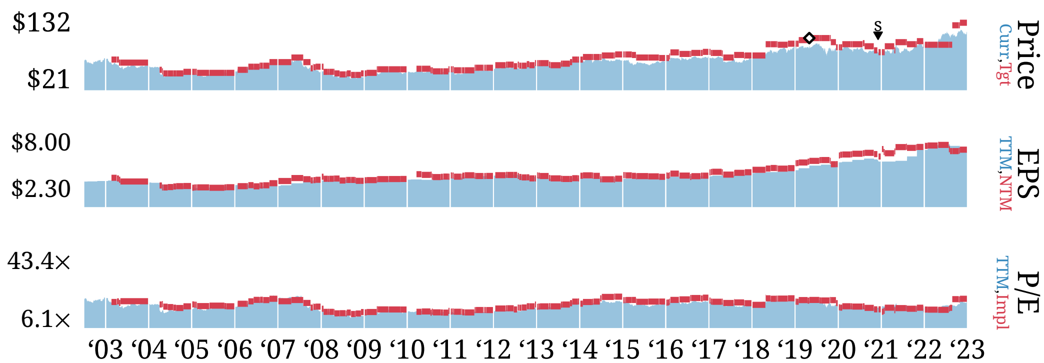

Figure B1(j). *y-axis shows min, median, and max. (Top) Blue ribbon is Merck's closing price on day t from CRSP,* Price_t*. Red line is Tim Anderson's price target,* $\text{PriceTarget}_t = \mathbb{E}_t[\text{Price}_{\tau+1}]$*, in IBES. (Middle) Blue is MRK's trailing twelve-month (TTM) EPS on day t from IBES,* EPS_t*. Red is Tim Anderson's EPS forecast,* $\mathbb{E}_t[\text{EPS}]$*. (Bottom) Blue is MRK's TTM P/E ratio,* $\text{TrailingPE}_t = \text{Price}_t \,/\, \text{EPS}_t$*. Red is the P/E implied by Tim Anderson's forecasts,* $\text{ImpliedPE}_t = \text{PriceTarget}_t \,/\, \mathbb{E}_t[\text{EPS}]$*.*

Lauren Lieberman's forecasts for Procter & Gamble (PG)

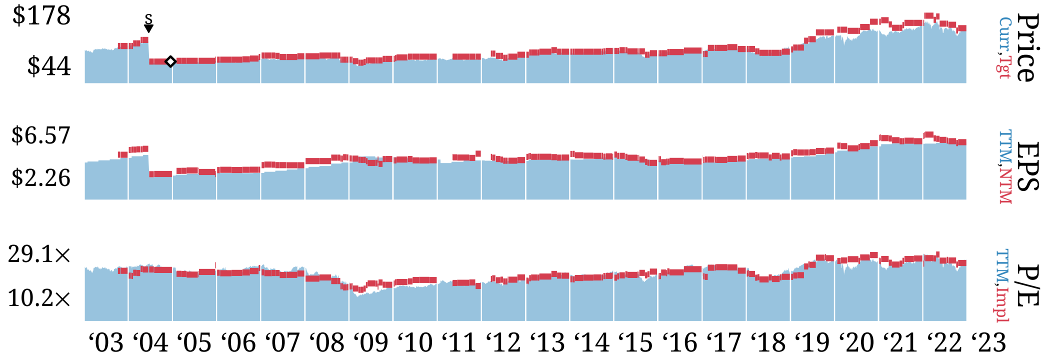

Figure B1(k). *y-axis shows min, median, and max. (Top) Blue ribbon is Procter & Gamble's closing price on day t from CRSP,* Price_t*. Red line is Lauren Lieberman's price target,* $\text{PriceTarget}_t = \mathbb{E}_t[\text{Price}_{\tau+1}]$*, in IBES. (Middle) Blue is PG's trailing twelve-month (TTM) EPS on day t from IBES,* EPS_t*. Red is Lieberman's EPS forecast,* $\mathbb{E}_t[\text{EPS}]$*. (Bottom) Blue is PG's TTM P/E ratio,* $\text{TrailingPE}_t = \text{Price}_t \,/\, \text{EPS}_t$*. Red is the P/E implied by Lieberman's forecasts,* $\text{ImpliedPE}_t = \text{PriceTarget}_t \,/\, \mathbb{E}_t[\text{EPS}]$*. We flag split events with* $S_\blacktriangledown$ *pointers.*

Brian Tunick's forecasts for Ross Stores (ROST)

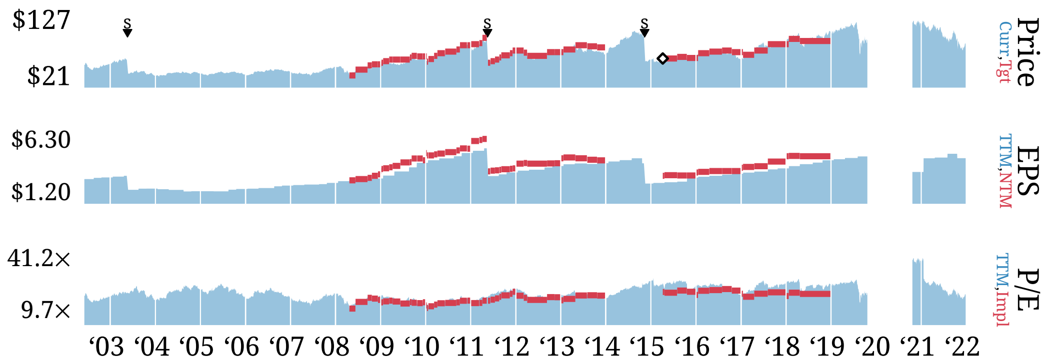

Figure B1(l). *y-axis shows min, median, and max. (Top) Blue ribbon is Ross' closing price on day t from CRSP,* Price_t*. Red line is Brian Tunick's price target,* $\text{PriceTarget}_t = \mathbb{E}_t[\text{Price}_{\tau+1}]$*, in IBES. (Middle) Blue is Ross' trailing twelve-month (TTM) EPS on day t from IBES,* EPS_t*. Red is Brian Tunick's EPS forecast,* $\mathbb{E}_t[\text{EPS}]$*. (Bottom) Blue is Ross' TTM P/E ratio,* $\text{TrailingPE}_t = \text{Price}_t \,/\, \text{EPS}_t$*. Red is the P/E implied by Brian Tunick's forecasts,* $\text{ImpliedPE}_t = \text{PriceTarget}_t \,/\, \mathbb{E}_t[\text{EPS}]$*.*

Justin Lake's forecasts for United Healthcare (UNH)

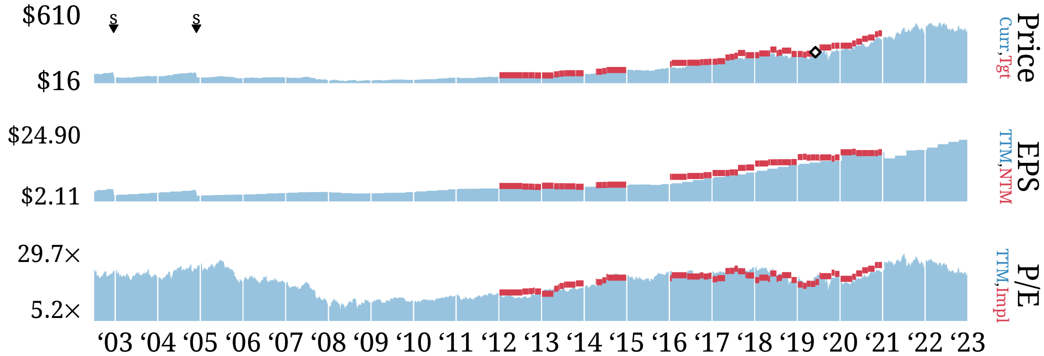

Figure B1(m). *y-axis shows min, median, and max. (Top) Blue ribbon is United Healthcare's closing price on day t from CRSP,* Price_t*. Red line is Justin Lake's price target,* $\text{PriceTarget}_t = \mathbb{E}_t[\text{Price}_{\tau+1}]$*, in IBES. (Middle) Blue is UNH's trailing twelve-month (TTM) EPS on day t from IBES,* EPS_t*. Red is Justin Lake's EPS forecast,* $\mathbb{E}_t[\text{EPS}]$*. (Bottom) Blue is UNH's TTM P/E ratio,* $\text{TrailingPE}_t = \text{Price}_t \,/\, \text{EPS}_t$*. Red is the P/E implied by Justin Lake's forecasts,* $\text{ImpliedPE}_t = \text{PriceTarget}_t \,/\, \mathbb{E}_t[\text{EPS}]$*. We flag split events with* $S_\blacktriangledown$ *pointers.*

Christopher Horvers' forecasts for Walmart (WMT)

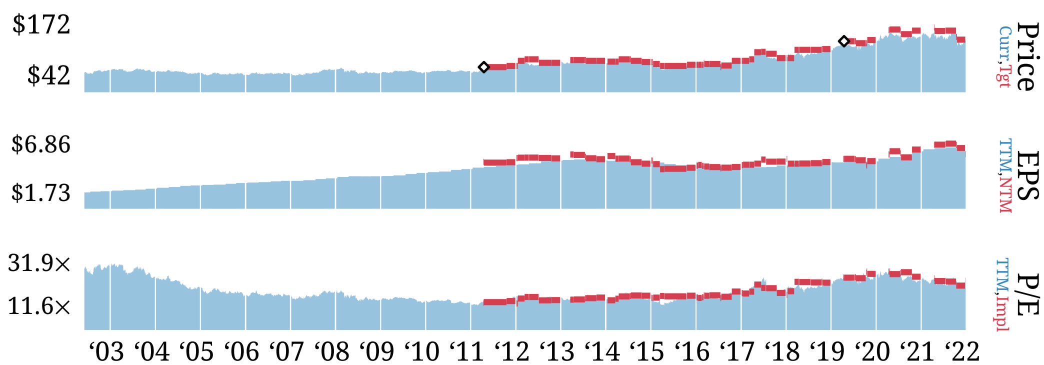

Figure B1(n). *y-axis shows min, median, and max. (Top) Blue ribbon is Walmart's closing price on day t from CRSP,* Price_t*. Red line is Chris Horvers' price target,* $\text{PriceTarget}_t = \mathbb{E}_t[\text{Price}_{\tau+1}]$*, in IBES. (Middle) Blue is WMT's trailing twelve-month (TTM) EPS on day t from IBES,* EPS_t*. Red is Chris Horvers' EPS forecast,* $\mathbb{E}_t[\text{EPS}]$*. (Bottom) Blue is WMT's TTM P/E ratio,* $\text{TrailingPE}_t = \text{Price}_t \,/\, \text{EPS}_t$*. Red is the P/E implied by Chris Horvers' forecasts,* $\text{ImpliedPE}_t = \text{PriceTarget}_t \,/\, \mathbb{E}_t[\text{EPS}]$*.*

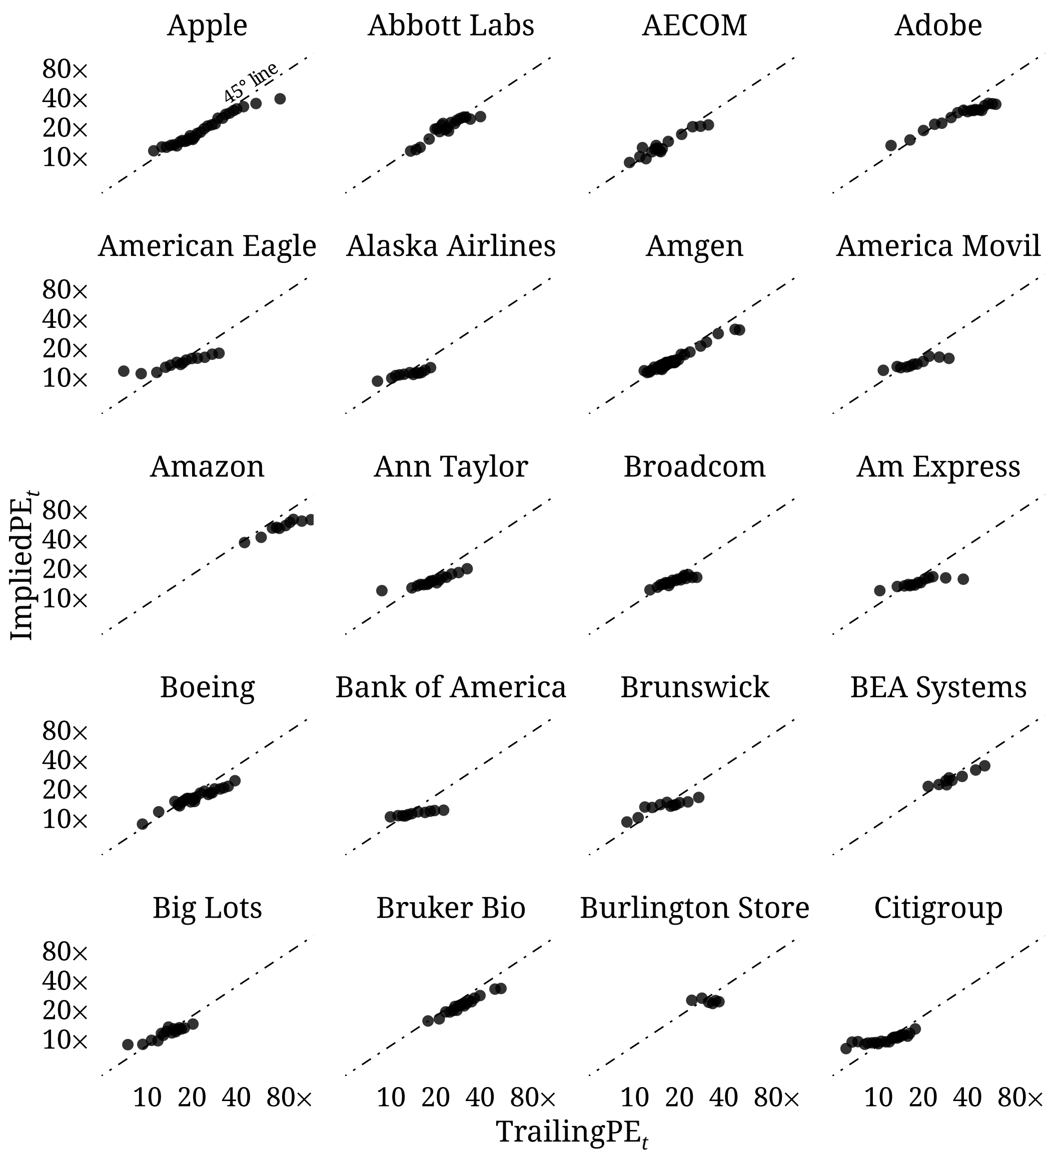

Figure B2(a). *Each panel shows a binned scatterplots using data from the full sample of IBES reports for a single firm. x-axis shows the firm's trailing twelve-month P/E,* $\text{TrailingPE}_{n,t} = \text{Price}_{n,t} / \text{EPS}_{n,t}$. *y-axis shows the P/E ratio implied by the analyst's price target and EPS forecast,* $\text{ImpliedPE}^a_{n,t} = \text{PriceTarget}^a_{n,t} / \mathbb{E}^a_t[\text{EPS}_n]$. *Sample: 2003 to 2022; 20 firms.*

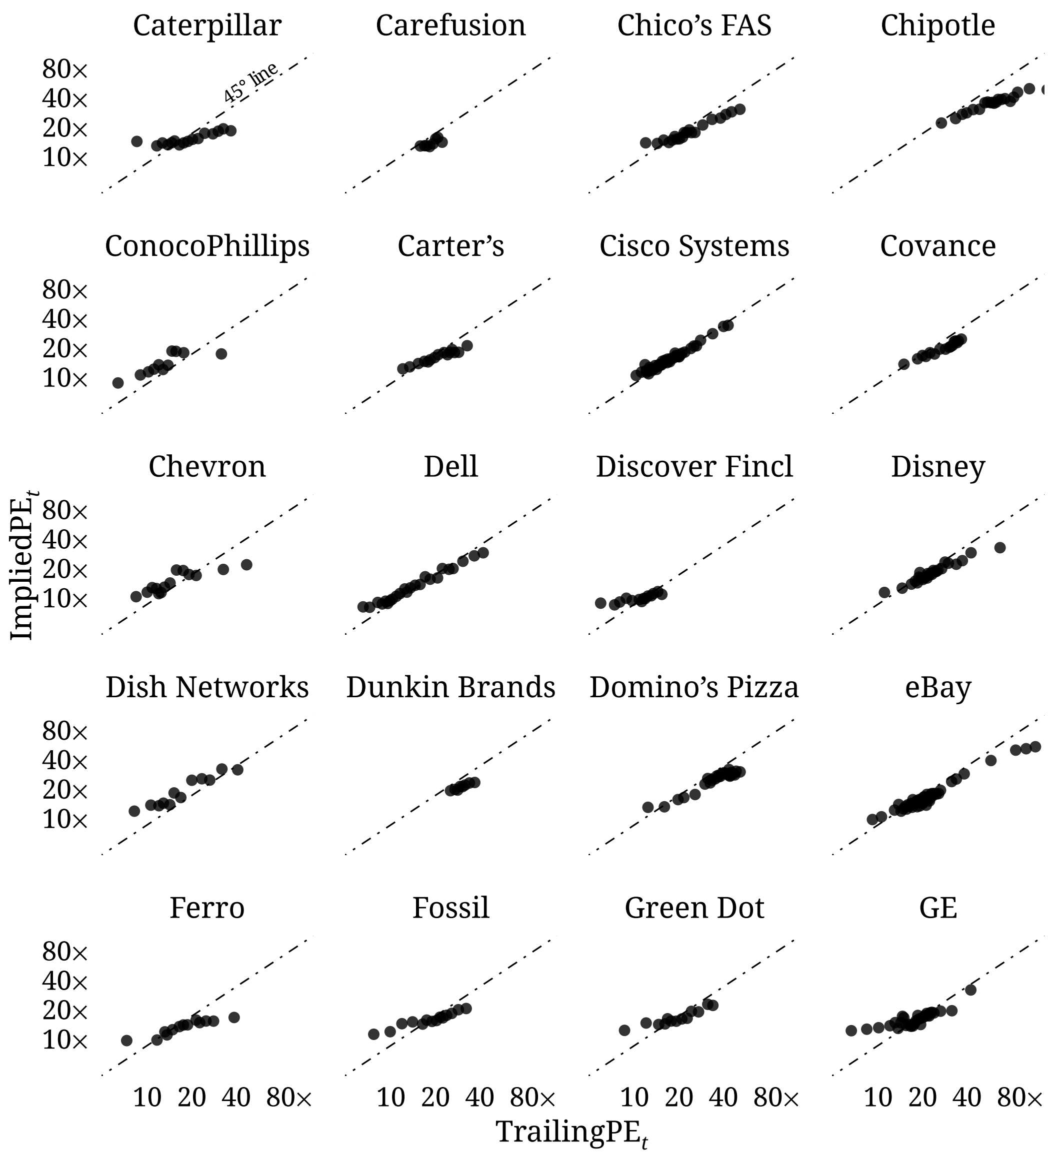

Figure B2(b). *Each panel shows a binned scatterplots using data from the full sample of IBES reports for a single firm. x-axis shows the firm's trailing twelve-month P/E,* $\text{TrailingPE}_{n,t} = \text{Price}_{n,t} / \text{EPS}_{n,t}$. *y-axis shows the P/E ratio implied by the analyst's price target and EPS forecast,* $\text{ImpliedPE}^a_{n,t} = \text{PriceTarget}^a_{n,t} / \mathbb{E}^a_t[\text{EPS}_n]$. *Sample: 2003 to 2022; 20 firms.*

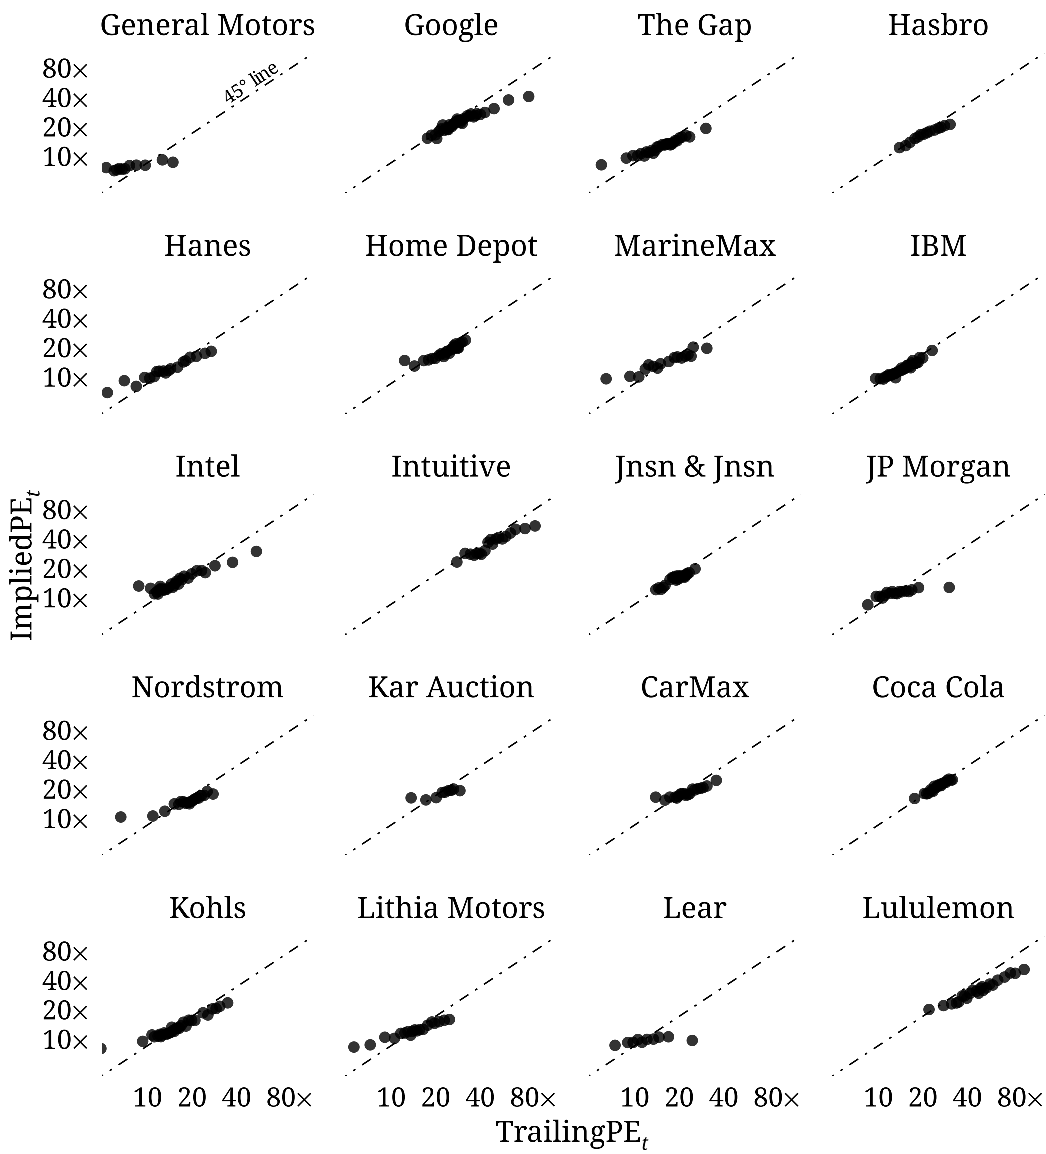

Figure B2(c). *Each panel shows a binned scatterplots using data from the full sample of IBES reports for a single firm. x-axis shows the firm's trailing twelve-month P/E,* $\text{TrailingPE}_{n,t} = \text{Price}_{n,t} / \text{EPS}_{n,t}$. *y-axis shows the P/E ratio implied by the analyst's price target and EPS forecast,* $\text{ImpliedPE}^a_{n,t} = \text{PriceTarget}^a_{n,t} / \mathbb{E}^a_t[\text{EPS}_n]$. *Sample: 2003 to 2022; 20 firms.*

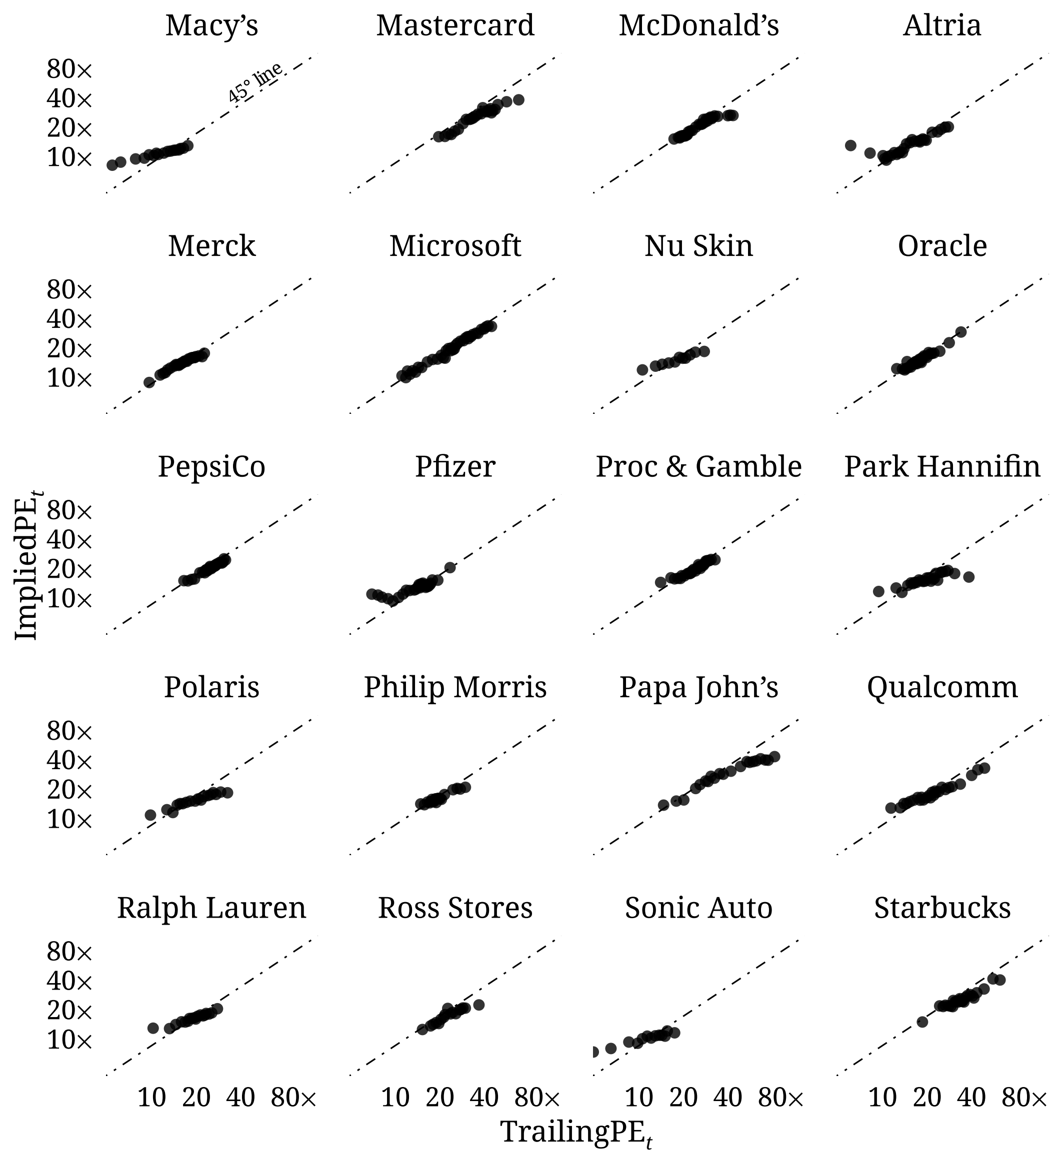

Figure B2(d). *Each panel shows a binned scatterplots using data from the full sample of IBES reports for a single firm. x-axis shows the firm's trailing twelve-month P/E,* $\text{TrailingPE}_{n,t} = \text{Price}_{n,t} / \text{EPS}_{n,t}$. *y-axis shows the P/E ratio implied by the analyst's price target and EPS forecast,* $\text{ImpliedPE}^a_{n,t} = \text{PriceTarget}^a_{n,t} / \mathbb{E}^a_t[\text{EPS}_n]$. *Sample: 2003 to 2022; 20 firms.*

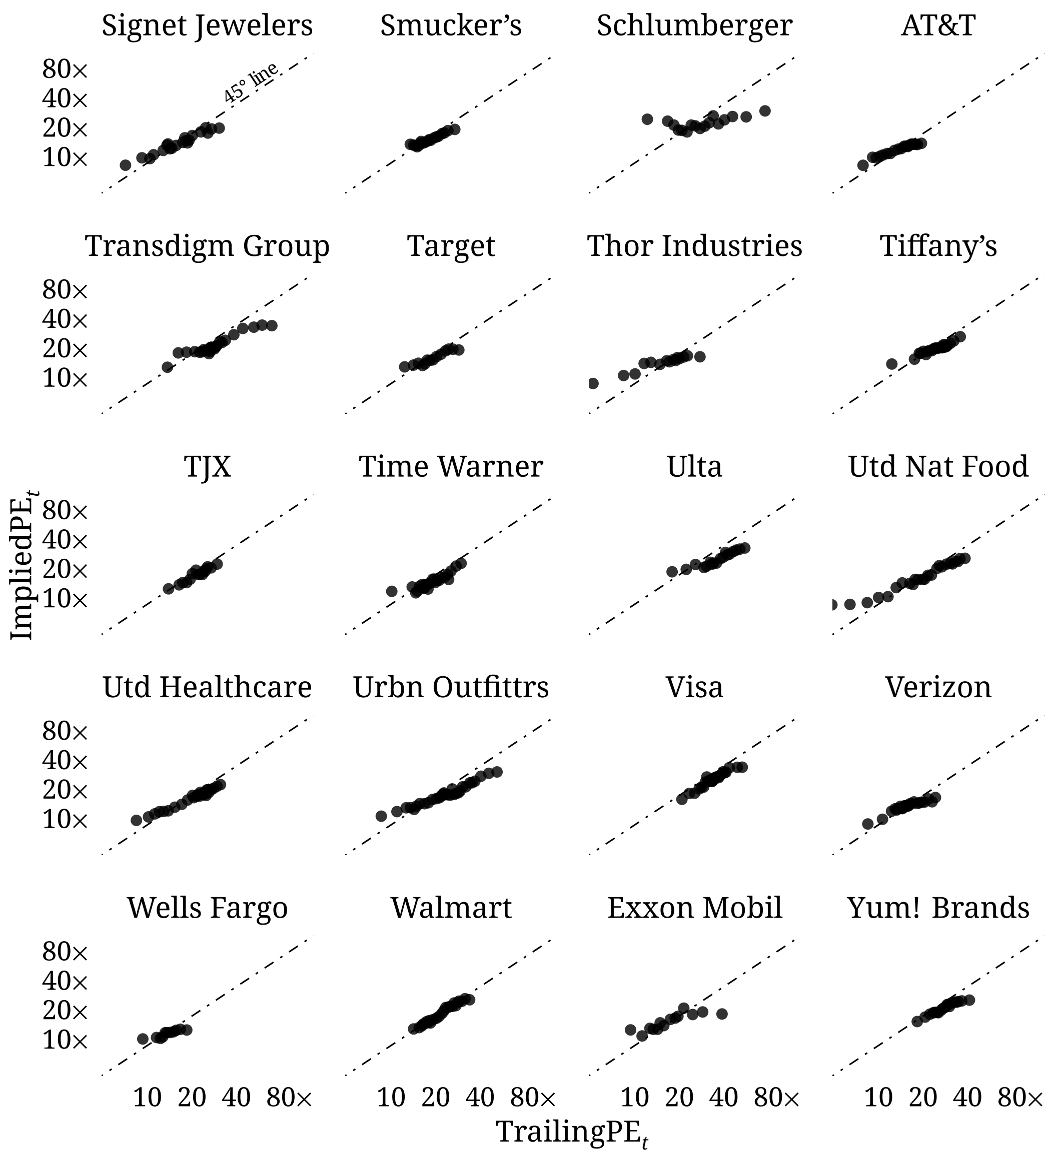

Figure B2(e). *Each panel shows a binned scatterplots using data from the full sample of IBES reports for a single firm. x-axis shows the firm's trailing twelve-month P/E,* $\text{TrailingPE}_{n,t} = \text{Price}_{n,t} / \text{EPS}_{n,t}$. *y-axis shows the P/E ratio implied by the analyst's price target and EPS forecast,* $\text{ImpliedPE}^a_{n,t} = \text{PriceTarget}^a_{n,t} / \mathbb{E}^a_t[\text{EPS}_n]$. *Sample: 2003 to 2022;* 20 *firms.*