NBER WORKING PAPER SERIES

EXPECTED EPS × TRAILING P/E

Itzhak Ben-David
Alex Chinco

Working Paper 32942
http://www.nber.org/papers/w32942

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2024

We would like to thank Xavier Gabaix, Sinan Gokkaya, Valentin Haddad, Jeff Meli, Stefan Nagel, Marco Sammon, Amir Sufi, Laura Veldkamp, and Jeff Wurgler for helpful comments. This paper has also benefited from feedback at the NBER SI Asset-Pricing meeting. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Expected EPS × Trailing P/E
Itzhak Ben-David and Alex Chinco
NBER Working Paper No. 32942
September 2024
JEL No. G12, G14, G32

ABSTRACT

All of asset-pricing theory currently stems from one key assumption: price equals expected discounted payoff. And much of what we think we know about discount rates comes from studying a particular kind of expected payoff: the earnings forecasts in analyst reports. Researchers typically access these numbers through an easy-to-use database and never read the underlying documents. This is unfortunate because the text of each report contains an explicit description of how the analyst priced their own earnings forecast. We study a sample of 513 reports and find that most analysts use a trailing P/E (price-to-earnings) ratio not a discount rate. Instead of computing the present value of a company's future earnings, they ask: "How would a firm with similar earnings have been priced last year?" Even if other investors do things differently, it does not make sense to put discount rates at the center of every asset-pricing model if market participants do not always use one. There are other options. Trailing twelve-month P/E ratios account for 91% of the variation in analysts' price targets.We construct a new kind of asset-pricing model around this fact and show that it explains the market response to earnings surprises.

Itzhak Ben-David
The Ohio State University
Fisher College of Business
606A Fisher Hall
Columbus, OH 43210-1144
and NBER
ben-david.1@osu.edu

Alex Chinco
Zicklin School of Business
Baruch College, CUNY
One Bernard Baruch Way
New York, NY 10010
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)" Researchers are willing to entertain the idea that investors might have biased subjective beliefs or hold non-standard preferences. But no one ever questions whether real-world investors apply present-value logic.

Suppose you buy a share of stock today and then sell it next year after collecting the dividend. Your total future payoff will be $\text{Dividend}_{t+1} + \text{Price}_{t+1}$. All standard models assume the current price will be given by something like

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

where $\mathbb{E}_t[\cdot]$ reflects investors' subjective beliefs about an asset's payoff next year and r is the discount rate they apply to these payoffs.

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 $\mathbb{E}_t[\text{Price}_{t+1}] = \$241$ 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 $\mathbb{E}_t[\text{EPS}_{t+1}] = \$10.48$ over the next twelve months (FY2020) and $\mathbb{E}_t[\text{EPS}_{t+2}] = \$11.50$ the year after (FY2021).

These are the sorts of numbers that researchers are used to downloading from the Institutional Brokers' Estimate System (IBES). Given that all asset-pricing models start with "price equals expected discounted payoff", researchers would love to have similar data on all investors' payoff expectations as well. Unfortunately, most market participants do not publicly announce their subjective beliefs. Sell-side analysts do, and as a result their forecasts have had a big impact on the asset-pricing literature (Kothari, So, and Verdi, 2016).

Researchers are primarily interested in discount rates, and much of what we think we know about them comes from studying IBES data. Researchers

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 improve 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.*

have spent a great deal of time trying to figure out how analysts discount their own subjective earnings forecasts. It is common to see researchers estimate an implied discount rate by replacing the numerator in Equation (1) with analysts' values of $\mathbb{E}_t[\text{EPS}_{t+1}] + \mathbb{E}_t[\text{Price}_{t+1}]$ and solving for r. Researchers frequently analyze longer-term effects by plugging analysts' forecasts into Campbell and Shiller (1988)'s log-linear approximation to Equation (1).

But researchers do not need to figure out how sell-side analysts price their own earnings forecasts. They describe their approach in the text of each report. We read what analysts write and find that most do not use a discount rate. They multiply their short-term earnings forecast times a trailing P/E (price-to-earnings) ratio. Rather than asking "How much is this company's expected earnings worth in today's dollars?", they ask "If a firm had reported similar earnings last year, how would it have been priced given recent multiples?"

Researchers cannot take it for granted that an asset's price always reflects its expected discounted payoff. Sometimes financial markets work this way. But in many important cases they do not. At the end of the day, it does not make sense to organize our entire field around the concept of discount rates if we cannot count on investors actually using one. The rest of the paper is about where we go from here. We model an alternative path forward that reflects how analysts describe their own pricing rule and show that this model can explain important patterns in real-world data.

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 *price target for Home Depot in his December 2019 earnings report for JP Morgan.*

Section 1 examines a sample of 513 sell-side analyst reports, which were written about large publicly traded companies from 2003 through 2022. These reports show that most analysts do not take a present-value approach to pricing their own subjective earnings expectations. Instead, they typically set a company's price target, $\text{PriceTarget}_t \overset{\text{def}}{=} \mathbb{E}_t[\text{Price}_{t+1}]$, equal to the firm's expected short-term earnings times its trailing P/E.

Sell-side research is more than just dry colorless lists of numbers. Analysts have been required to state their pricing rule in the text of each report since May 2002 when the SEC passed NASD Rule 2711. The current statute is FINRA Rule 2241, which says "any recommendation, rating, or price target [must be] accompanied by a clear explanation of any valuation method used."

Figure 2 shows the "clear explanation" that Chris Horvers gave in his report. He could have said that he carefully examined Home Depot's expected cash-flow stream from 2021 onward. He could have gone deep into the weeds, outlining precisely how he discounted these expected future cash flows back to December 2020. He was more than capable of doing this sort of analysis.

Chris Horvers chose not to. He plainly states that his \$241 price target for Home Depot was "based on ~ 21.0× our revised 2021E EPS, in line with its three-year average"

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

To forecast Home Depot's share price in December 2020 (end of period $t+1$), he multiplied his EPS forecast for 2021 (period $t+2$) times a 3-year trailing average P/E based on 2017, 2018, and 2019 (periods $t-2$, $t-1$, and t).

It is possible to interpret certain kinds of multiples analysis as special cases of present-value logic. For example, the classic Gordon 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 (E_i[Dividend_{i+1}] + E_i[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 resulting pricing rule in Equation (4) uses a forward-looking multiple, $\left(\frac{1}{r-g}\right) = \left(\sum_{h=1}^{\infty} \frac{\mathbb{E}_t[\text{Dividend}_{t+h}]}{(1+r)^h}\right) \Big/ \mathbb{E}_t[\text{Dividend}_{t+1}]$, as a shortcut for calculating the present value of an infinite stream of expected future dividend payments.

While the timing of Equation (2) strongly suggests Chris Horvers has seen the Gordon model, his underlying economic logic could not be more different. There is nothing forward-looking about the "Valuation Matrix" in Figure 2. Chris Horvers' arrived at his \$241 price target by asking himself: "How has the market typically priced each \$1 of Home Depot's earnings during the past few years?" Moreover, he did this in spite of the fact that he gave Home Depot an "overweight" rating. Chris Horvers felt the market had been undervaluing Home Depot and expected the company to experience multiples expansion going forward. He still used a trailing P/E to set his price target.

Our point is not that Chris Horvers is a bad analyst. Chris Horvers is an excellent analyst. He has been named to *Institutional Investor* magazine's All-American 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 actually read it.

There is so much more to sell-side research than the handful of numerical values found in IBES. Every report includes a detailed explanation showing how the analyst valued the company. Chris Horvers tells us that he is not setting price equal to expected discounted payoff, and most other sell-side analysts say the same thing. As a rule, they do not apply present-value logic when pricing their own subjective cash-flow expectations.

Frankly, it is surprising that so many papers have been written about biased earnings forecasts. Analysts put a lot of effort into getting those numbers right. Chris Horvers spent pages justifying his value of $\mathbb{E}_t[\text{EPS}_{t+2}] = \$11.50$. Then, when it came time to capitalize this expected cash flow into a share price, he did not even attempt to approximate $\sum_{h=1}^{\infty} \frac{\mathbb{E}_t[\text{EPS}_{t+h}]}{(1+r)^h}$. He chose not to apply the "one simple concept" that is the bedrock of our entire research program.

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 discount rates and perform Campbell and Shiller (1988) 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.

Section 2 outlines a productive path forward. In many applications, "price equals expected discounted payoff" is not a good foundation to build a model on. If you are interested in modeling how S&P 500 companies get priced and plan on using IBES data in your empirical work, then you need a theory that stems from some other "simple concept". We show what such a model might look like. Our theory provides a simple rationale for why analysts use trailing P/E ratios and also generates sharp testable predictions.

Here is how the model works. There is a single stock, and each year analysts set a one-year-ahead price target. If this price target is higher or lower than the firm's current share price, then investors proportionally adjust their holdings at rate $\mu > 0$. When the price target is 1% higher than a stock's current price, investors increase their holdings by μ% over the next year. The stock's subsequent price growth also proportionally responds to changes in investors' demand. Each 1% change in investor demand causes the company's expected return to change by ν% over the next year for some $\nu > 0$.

Our model suggests a plausible justification for using trailing P/E ratios. The equilibrium price in our model is mostly backward-looking. There is only one forward-looking input: analysts' short-term EPS forecast. It makes sense for analysts to use a mostly backward-looking pricing rule in a world where prices themselves are mostly backward-looking. In such a world, PriceTarget = $\mathbb{E}$[EPS] × TrailingPE can generate forecasts that are correct on average.

It also makes sharp testable predictions. Imagine that immediately after Chris Horvers published his December 2019 report, he saw something that led him to believe Home Depot's earnings would go up by $\Delta\mathbb{E}[\text{EPS}] = +\$0.05$ over the next year and then grow at a $\Delta g = +1\%$ faster rate thereafter. In our model, the change in Home Depot's long-run growth rate would be irrelevant. The company had been trading at 21.0×, so Chris Horvers' price target would go up by $\Delta\text{PriceTarget} = \$1.05 = 21.0 \times \$0.05$. By contrast, in the Gordon model, the +1% increase in Home Depot's long-run growth rate would cause Chris Horvers to use a larger multiple, leading to an even larger increase in his price target.

Section 3 demonstrates that trailing P/E ratios explain both analysts' price targets and realized price changes. We start with analysts' price targets. Chris Horvers said he used a trailing P/E ratio to set price targets for Home Depot. Figure 3 shows that IBES data support this claim.

The red line in the top panel shows Chris Horvers' most recent $\text{PriceTarget}_t \overset{\text{def}}{=} \mathbb{E}_t[\text{Price}_{\tau+1}]$ on each trading day t during the window from 18 months to 6 months prior to Home Depot's next fiscal year-end date, $(\tau + 1)$. The red line in the bottom panel shows the P/E ratio implied by Chris Horvers' price target and his earnings forecast for the following fiscal year, $\text{ImpliedPE}_t \overset{\text{def}}{=} \text{PriceTarget}_t \,/\, \mathbb{E}_t[\text{EPS}_{\tau+2}]$. This implied P/E closely tracks Home Depot's trailing twelve-month (TTM) P/E ratio shown in blue, $\text{TrailingPE}_t \overset{\text{def}}{=} \text{Price}_t \,/\, \text{EPS}_t$.

Christopher Horvers' forecasts for Home Depot (HD)

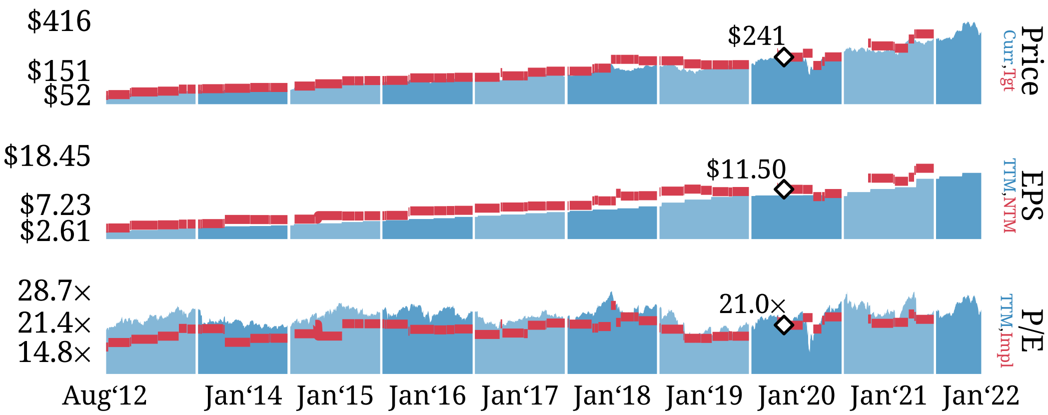

Figure 3. *y-axis labels correspond to the min, median, and max in each panel. (Top Panel) Blue ribbon is Home Depot's closing price on day t in CRSP, $Price_t$. Red line is Chris Horvers' price target $PriceTarget_t$ in IBES. (Middle Panel) Blue is the sum of HD's quarterly EPS in IBES over four quarters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

Of course, we know that Chris Horvers' 21.0× trailing P/E in December 2019 was computed over the past three years not just the most recent twelve months. These sorts of perturbations explain why the fit is not perfect. Nevertheless, the right panel of Figure 4 shows that our simple trailing twelve-month calculation is a good first approximation to how most analysts set price targets.

By capturing the essence of what analysts say they are doing, PriceTarget = $\mathbb{E}$[EPS] × TrailingPE, we are able to explain over 91% of the variation in their price targets. Our model's fit is not perfect. But it predicts investors' subjective return expectations far better than any other model we know of.

We next turn our attention to how the market responds to earnings surprises. When a company announces quarterly earnings that deviate from the consensus forecast, investors often learn important things about a firm that can have long-term ramifications. The larger the surprise, the more likely this is. In the Gordon model, a change in the long-run earnings growth rate would cause investors to price a firm's expected short-run earnings using a different multiple.

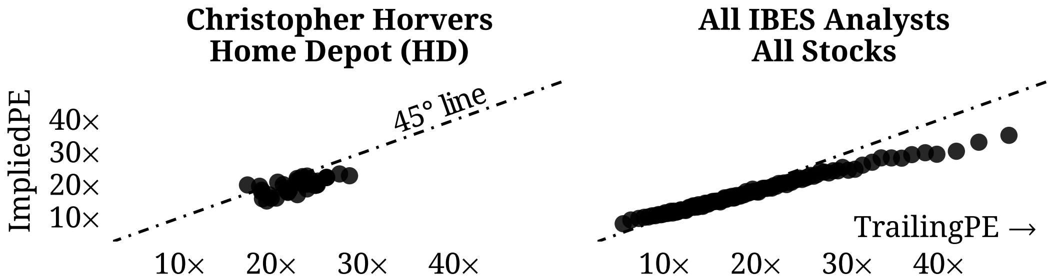

Figure 4. *(Left) Each dot denotes a day on which Chris Horvers updated his price target for Home Depot. x-axis is Home Depot's trailing twelve-month P/E, $TrailingPE_t = Price_t \,/\, EPS_t$. y-axis is the P/E ratio implied by Chris Horvers' forecast values, $ImpliedPE_t = PriceTarget_t \,/\, \mathbb{E}_t[EPS_{\tau+2}]$.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

By contrast, our model uses a trailing P/E that is set in stone. Analysts update their price targets to reflect changes in their short-term earnings forecast. The company's share price responds proportionally, $\Delta\text{Price} = \{\lambda \cdot (\text{EPS} - \mathbb{E}[\text{EPS}])\} \times \text{TrailingPE}$ for some $\lambda > 0$. Importantly, our model says that the same value of λ will apply no matter how large or small the earnings surprise was.

We test this prediction using an approach similar to Fama and MacBeth (1973). First, we group stock-quarter observations by the size of their earnings surprise, $s \in \{-\$0.30, \ldots, \$0.30\}$. Then, within each bucket, we run a separate regression to estimate the relationship between a stock's trailing P/E ratio and its subsequent price change, $\hat{\lambda}_s$. This gives us 60 distinct estimates for $\lambda \cdot s$, and we verify that they all imply a very similar value of $\hat{\lambda}_s/s = \lambda$.

There is a deep crack in the foundation of asset-pricing theory. Researchers currently assume that an asset's price must always reflect some sort of expected discounted payoff. They spend most of their time trying to understand how an asset's expected payoff is being discounted, and analysts' earnings forecasts have played a key role in this pursuit. However, when we read what analysts write in their reports, we find that most of them do not set price targets by discounting a company's expected EPS. Instead, they typically use a trailing P/E ratio. We point to a better way forward by constructing a new kind of model and showing that it explains key patterns in the data.

Related literature. This paper borrows from and builds on several strands of related literature. To start with, it 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.

We are not the first researchers to notice that market participants talk about short-term cash-flow multiples. There are numerous papers studying the accuracy of multiples analysis for pricing public equities (Bhojraj and Lee, 2002; Liu et al., 2002; Da and Schaumburg, 2011; Bartram and Grinblatt, 2018; Cooper and Lambertides, 2023), IPOs (Kim and Ritter, 1999; Purnanandam and Swaminathan, 2004), and syndicated loan deals (Murfin and Pratt, 2019). We think these papers are fascinating and deserve far more attention. But we also think they bury the lede. The fact that analysts use a trailing P/E is inconsistent with textbook present-value logic…no matter how accurate their forecasts are.

Our paper connects to the broader literature on belief formation (Malmendier and Nagel, 2011; Greenwood and Shleifer, 2014; Coibion and Gorodnichenko, 2015; Giglio, Maggiori, Stroebel, and Utkus, 2021; Bordalo, Gennaioli, Ma, and Shleifer, 2020). There is substantial evidence that sell-side analysts suffer from predictable biases when making earnings forecasts (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). We would also like to know how biases in analysts' subjective earnings forecasts affect asset prices. But, to answer this question, researchers need to correctly model how these subjective beliefs get priced.

Last but not least, this paper provides evidence against the discount-rate approach to asset pricing. Textbook models say that "asset markets are [supposed to be] in reality big insurance markets (Cochrane, 1999)." These models argue that a company's current share price should reflect 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). This paper proposes a simple alternative approach.

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 this section, we examine how analysts explain their own pricing rule and find that most do not compute any sort of expected discounted payoff. Instead, they typically rely on a trailing P/E ratio.

1.1 Data description

Analysts explicitly state how they price their own subjective cash-flow expectations in the text of each earnings report. We read these reports and ask: "How did the author of this report convert their earnings forecast, $\mathbb{E}_t[\text{EPS}]$, into a price target, PriceTarget_t?" Every report has an entire section where the analyst directly answers this question as shown in Figure 2.

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 sell-side analysts apply present-value reasoning to price their own subjective cash-flow expectations. However, these are run-of-the-mill reports written by average analysts. Perhaps the best analysts still use forward-looking present-value logic 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

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.*

repeatedly named to *Institutional Investor* magazine's All-American research team. These analysts are the best of the best (Stickel, 1992), and coverage-initiation reports are the ones that really matter. Analysts put a disproportionate amount of effort into writing them (McNichols and O'Brien, 1997), often laying out their general theory for pricing the firm. The average coverage-initiation report in our sample runs 29 pages. 20% are 40+ pages long.

Institutional Investor publishes their rankings in October. We read through these issues and recorded which analysts made the All-American research team each year. The magazine ranks analysts by GICS sector. For each sector, we created a list of the 10 analysts with the most years on the All-American team. The 174 documents in our second wave are coverage-initiation reports written by All-American analysts on our 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 firms also start out with negative earnings, making it difficult for analysts to apply the formula $\text{PriceTarget} = \mathbb{E}[\text{EPS}] \times \text{TrailingPE}$.

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-American 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.*

Just like in our first wave, our second wave of downloads only includes 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 economic research team and the single most-capped analyst on *Institutional Investor* magazine's All-American 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 in IBES. Moreover, the P/E ratios implied by these numbers (red lines; bottom panel) line up with the P/E ratios that analysts say they are using in their reports. We do not include the data from these additional reports in our main analysis; we only use them to ensure the accuracy of our raw numbers.

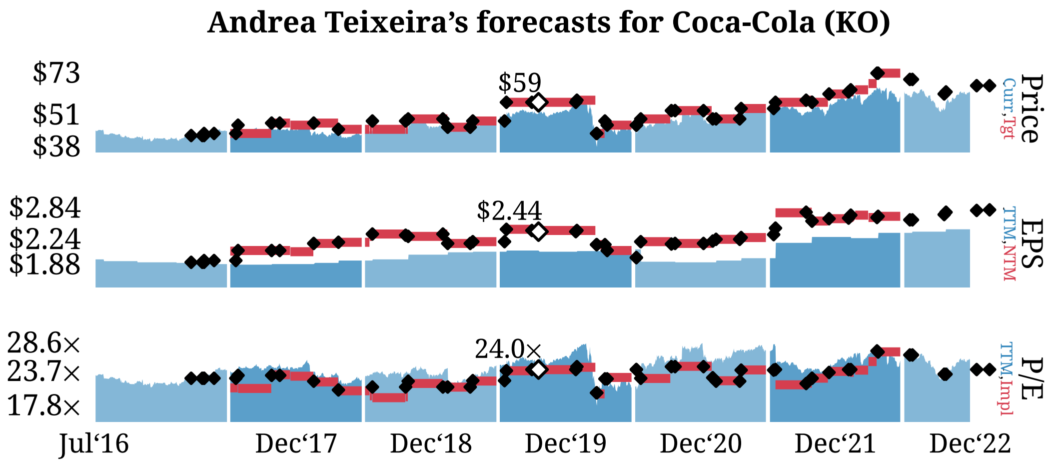

Figure 6. *y-axis shows min, median, and max. (Top) Blue ribbon is Coca-Cola's (KO)'s closing price on day t from CRSP, $Price_t$. Red line is Andrea Teixeira's price target, $PriceTarget_t = \mathbb{E}_t[Price_{\tau+1}]$, as reported in IBES. (Middle) Blue is KO's trailing twelve-month (TTM) earnings 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, $\mathbb{E}_t[EPS_{\tau+2}]$. (Bottom) Blue is KO's TTM price-to-earnings (P/E) ratio, $TrailingPE_t = Price_t \,/\, EPS_t$. Red is the P/E implied by Ms Teixeira's forecasts, $ImpliedPE_t = PriceTarget_t \,/\, \mathbb{E}_t[EPS_{\tau+2}]$. 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 norm. 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). These calculations 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. In this paper, we treat it as a separate method to be as conservative as possible.

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 (P/E) 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%.

Again, there is nothing inherently wrong with multiples analysis. The issue is the way that analysts calculate their multiple. Sell-side analysts typically choose their multiple based on where the firm and others like it have been trading at in recent years. Table 5 shows that analysts set a price target by looking 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%).

The popularity of peer-group comparisons in our data is driven by the fact that coverage-initiation reports make up a third of our sample (174 of 513 reports; 33.9%). As noted earlier, 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%.

There is every reason to believe that analysts are capable of calculating a forward-looking multiple. 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.

Likewise, analysts use 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 expected discounted payoff. Himanshu Patel uses a combination of backward-looking multiples.

Sell-side analysts have played a critical role in the asset-pricing literature, and "price equals expected discounted payoff" is the bedrock of every asset-pricing model. It is problematic that this particular group of market participants typically does not discount their own expected payoffs.

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). What's more, this statistic includes any report that mentions the terms "DCF" or "Discounted Cash Flow" in the methods section. Green, Hand, and Zhang (2016) notes that, in roughly 91% of reports that use these keywords, "there is no recognizable DCF model provided in the report itself." Many reports that talk about DCF modeling do so using boilerplate language without providing any specifics.

What's more, 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).

Prior to examining the data, we anticipated 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. However, the "All Am" column in Table 7 shows that DCF analysis is even

J.P.Morgan

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

North America Equity Research
12 April 2010

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

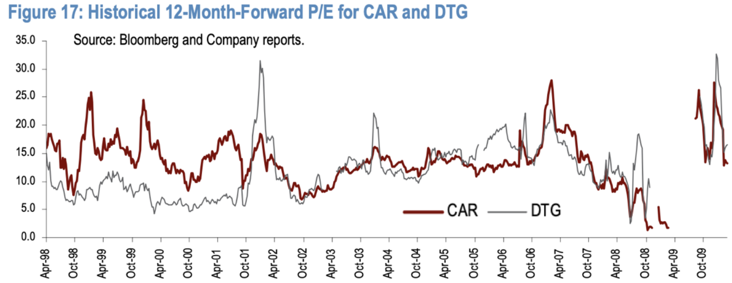

(b) Methods section

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-American team.*

less common in this subset of our data. Only one in five coverage-initiation reports make 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 science tools industry," he has "chosen to use a DCF methodology" out of necessity.

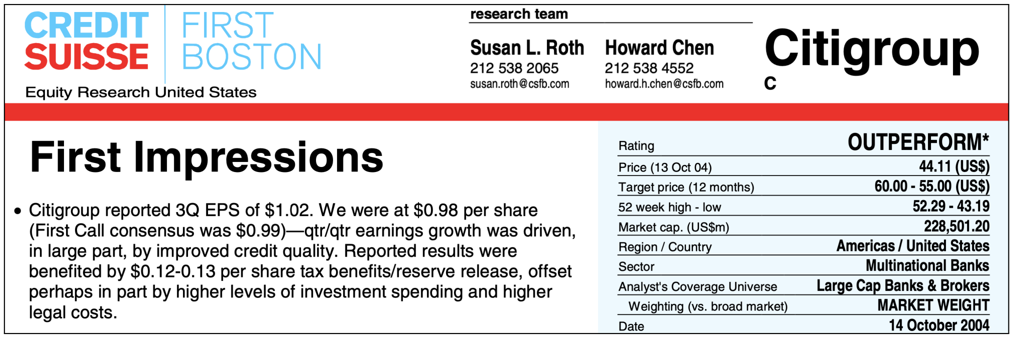

CREDIT SUISSE | FIRST BOSTON

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.*

"PACB is unprofitable (and yet lacks revenue)." Sell-side analysts are perfectly capable of doing these calculations. But they typically choose not to.

What's more, 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, describes computing the "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)."

1.4 DCF Models Are Mainly Used In Niche Industries

We study 513 earnings reports written about large publicly traded companies. These are the firms that researchers typically have in mind when modeling

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%.

the stock market. And, when valuing this group of firms, we find that sell-side analysts do not typically set price equal to expected discounted payoff.

However, analysts do regularly use DCF models 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. DCF is also used to price real-estate investment trusts (REITs) and resource-extraction companies (oil, gas, mining, etc).

It is obvious when a report uses present-value reasoning. 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 gives us the precise numerical inputs needed to do the calculation. Asset-pricing researchers assume that every earnings report looks like this. Outside of a few special situations, most do not.

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-American team.*

1.5 Expected Returns Do Not Reflect (Exotic) Risks

Asset-pricing textbooks argue that an asset's expected return is determined by how its payoffs are distributed across good and bad future states of the world. The logic 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 realized value of the stochastic discount factor (SDF) in state s, $m_s \overset{\text{def}}{=} (\frac{1}{1+r_s})$, reflects the current price of an asset that would pay \$1 next year in that state.

To keep things simple, suppose there are only two states, $s \in \{\text{good}, \text{bad}\}$. In this framework, investors would be willing to pay $\$1 \cdot m_{\text{bad}} = (\frac{\$1}{1+r_{\text{bad}}})$ today to receive \$1 next year in the bad state. Researchers assume that $\$1 \cdot m_{\text{bad}} = (\frac{\$1}{1+r_{\text{bad}}}) > (\frac{\$1}{1+r_{\text{good}}}) = \$1 \cdot m_{\text{good}}$ since that is when they will really need the money. When an asset's expected returns are unusually low, researchers figure

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.*

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 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. We do not expect market participants to use terms like "stochastic discount factor", but investors must be able to formulate certain basic ideas for the SDF approach to work. Investors can buy a stock that offers insurance against bad times without knowing it. But textbook models predict that this stock will have low expected returns because investors bid up its current price on account of this insurance value. Investors cannot do this by accident.

But, on the other hand, we also find very little evidence that expected returns are compensation for simple risk either. Instead, sell-side analysts often say that their next-twelve-month return forecast is derived from their beliefs about a company's 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".

January 13, 2014

Outperform / V

Equity Research

Sector: Marine GP
Overweight

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

WELLS FARGO SECURITIES

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.**

(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

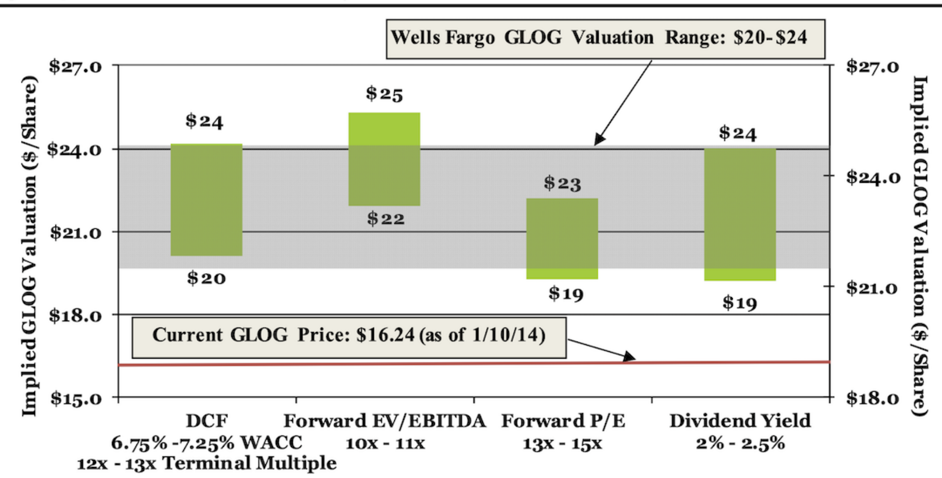

Source: Company filings, Wells Fargo Securities, LLC estimates

(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-American team.*

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

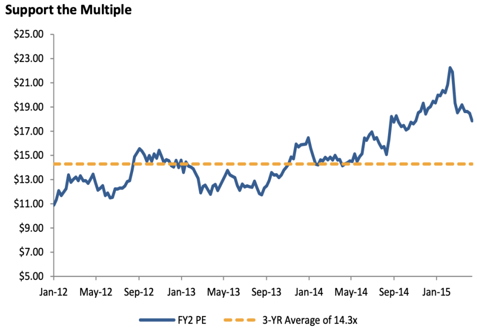

Source: Company reports and RBC Capital Markets estimates

(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-American team.*

This is precisely what we would expect from an analyst who thought about a company's price mainly in terms of trailing P/E ratios. Suppose an analyst calculates his price targets, $\mathbb{E}_t[\text{Price}_{t+1}] = \mathbb{E}_t[\text{EPS}_{t+2}] \times \text{TrailingPE}_t$, 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

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

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

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

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 his approach, there is no reason why Brian Tunick's price target would respect forward-looking accounting identities.

Researchers think it is completely natural to replace 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 swap out the 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... This is where the dividend discount model in Equation (3) comes from. But, if analysts do not set price equal to expected discounted payoff, there is no reason for them to enforce consistency between an asset's current price and the discounted value of their price target.

Asset-pricing researchers often try to understand the asset-pricing implications of biased subjective beliefs by plugging analyst earnings forecasts from IBES into Campbell and Shiller (1988)'s log-linear approximation for $\left(\frac{1}{r-g}\right)$

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

In his own textbook, John Campbell explains how Equation (7) should "hold ex ante, not only for rational expectations but also for irrational expectations that respect identities. (Campbell, 2017)" If sell-side analysts' price targets do not respect identities, it is not clear what we learn from this exercise.

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 used a backward-looking 24× P/E ratio to create her price target.

The numbers in earnings reports often do not mean what researchers think they mean. In many cases, they are computed in a way that suggests the analysis must be approaching the problem differently. For example, Figure 14(c) shows a table of key metrics from an October 2019 earnings report written by Andrea Teixeira about Pepsi. The row highlighted in blue shows Pepsi's share price in October 2019, $\text{Price}_{\text{Oct'19}} = \$138.23$, divided by its EPS in a given year

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

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

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

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

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

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 rebase 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.*

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 $21.4\times = \frac{\text{Price}_{\text{Dec'19}}}{\text{EPS}_{\text{'18}}} = \frac{\$212.00}{\$9.89}$, $21.1\times = \frac{\text{Price}_{\text{Dec'19}}}{\mathbb{E}[\text{EPS}_{\text{'19}}]} = \frac{\$212.00}{\$10.05}$, $20.2\times = \frac{\text{Price}_{\text{Dec'19}}}{\mathbb{E}[\text{EPS}_{\text{'20}}]} = \frac{\$212.00}{\$10.48}$, and $18.4\times = \frac{\text{Price}_{\text{Dec'19}}}{\mathbb{E}[\text{EPS}_{\text{'21}}]} = \frac{\$212.00}{\$11.50}$. 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).

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)"

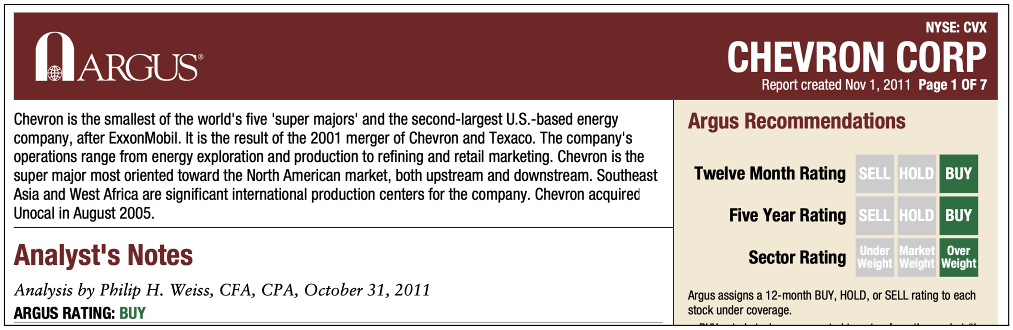

ARGUS

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

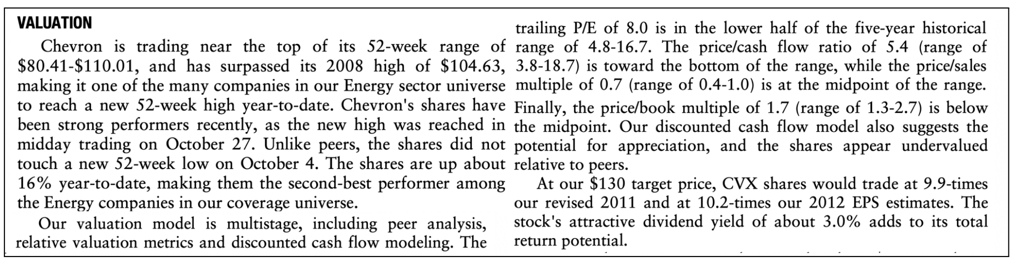

VALUATION

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.*

Figure 15 shows a November 2011 report about Chevron Corp written by Philip Weiss. In the methods section of his report, Mr Weiss says that he considered both trailing multiples analysis as well as a DCF model when setting $\text{PriceTarget}_t = \$130$, which was 24% higher than Chevron's current price, \$105.05. But he did not make any reference to Chevron's dividend yield when justifying 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 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, $\frac{\$3.12}{\$105.05} \approx 3\%$,

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

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 number of sell-side analysts. Based on these conversations, our general sense is that the methods section contains a brief honest account of how they actually compute their price targets. 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 analysts calculated them. Why should "4" be any more worthy of study than "two times two"?

Moreover, even if an analyst does not put in much effort when writing the methods section of their report, this fact should not push them towards descriptions of tailing P/E ratios rather than discounted cash flows. It is just as easy to give a brief account of either approach. 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 $\left(\frac{1}{r-g}\right)$ 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?

A recent working paper, Décaire and Graham (2024), analyzes a much larger data set and finds a similar point estimate. This paper uses machine-learning techniques to analyze the discount rates reported in a collection of 78.5k earnings reports. When describing their sample, the authors acknowledge that only "40% of all reports available on Refinitiv since 2009" include a DCF model. Thus, the 78.5k earnings reports in that study are analogous to the subset of 155 reports in our sample where the analyst talks about using a DCF model in some capacity. This is a highly non-representative sample.

Another recent working paper, Gormsen and Huber (2024), employed a team of research assistants to analyze what managers said to sell-side analysts in 74k quarterly earnings conference calls. Just like before, the authors find that most conference calls do not make any reference to present-value logic. Select $1 at random from the total value of the US stock market. There is a greater than 60% chance that this $1 came from a firm that has never mentioned a discount rate in *any* conference call over the past two decades. The same ballpark 40% value shows up again. These papers show how infrequently DCF models are used.

1.10 Do Other Investors Think Like Analysts?

There are good reasons to believe that sell-side analysts are not the only ones using trailing P/Es. Again, it is called *sell*-side research for a reason. Presumably there are other investors interested in buying this research output. Sell-side analysts have been around in something resembling their current form since

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.*

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.

We also 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 Where Do Researchers Go From Here?

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, analysts do not use one.

Sell-side analysts describe their price-forecasting problem in an entirely different way than an academic researcher would. The day-to-day business of being an asset-pricing theorist involves writing down models of expected returns. They assume that investors use the model-implied expected return as their discount rate when setting price levels.

Of the 155 reports in our sample that mentioned a DCF or dividend discount model, the majority never specify which discount rate was used. Philip Weiss said he used a DCF model to set a price target for Chevron in his November 2011 report. This document contained nearly 200 exact numerical values. For each trailing multiple he listed in Figure 15(b), Mr Weiss gave both the precise value as well as its historical range. However, at no point in his report did Philip Weiss bother to say which discount rate he used to calculate $\left(\frac{1}{r-g}\right)$.

With $N \geq 3$ observations, it is always possible to estimate a cross-sectional regression, $Y_n \sim \hat{\alpha} + \hat{\beta} \cdot X_n + \epsilon_n$. But every researcher recognizes that $\hat{\beta} \neq 0$ does not imply a causal relationship between X and Y. Likewise, there is nothing stopping a researcher from estimating the best-fit discount rate implied by analysts' earnings forecast and a company's current share price. But the existence of this fitted value does not imply that analysts are using it to set target prices or form expectations. It is clear from reading analysts' reports that they are not. This is not a solid enough foundation to build our entire field around.

Researchers have found that analyst-implied discount rates are consistent with the intertemporal capital asset-pricing model (ICAPM; Pástor, Sinha, and Swaminathan, 2008), international asset-pricing models (Lee, Ng, and Swaminathan, 2009), and the pricing of default risk (Chava and Purnanandam, 2010). But if analysts are not actually using these implied discount rates to price assets, where is the close empirical fit coming from? It is possible to accidentally buy insurance against a future risk, but investors cannot accidentally assign the correct price to this insurance (Chinco, Hartzmark, and Sussman, 2022).

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. Textbook models say that investors equate a company's share price with its expected discounted payoff to each shareholder. Financial economists think about the discount rate embedded in this pricing rule as the most important part of the problem. By contrast, the analysts in our sample focus all their attention on predicting a company's earnings. They pick a recent P/E almost as an afterthought.

While the behavioral-finance literature has spent most of its time looking for biases in analysts' EPS forecasts, at least analysts are trying hard to get those numbers right. In contrast, they are not even trying to calculate the expected discounted payoff at the heart of every standard asset-pricing model. This seems like the bigger issue.

We appreciate that every profession does some things on autopilot. For example, financial economists often cluster their standard errors without thinking too carefully about how they do it (Petersen, 2008). 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.

2 A Simple Model

We have just seen that, in addition to telling us their subjective earnings forecasts, sell-side analysts also tell us how they price them. They do not set "price equal to expected discounted payoff". Our hope is that, the next time you sit down to write an asset-pricing model, you will consider building this model on some other "simple concept". In this section, we provide an example of what such a model might look like.

In Subsection 2.1, we start by outlining how a company's earnings change over time. Then, in Subsection 2.2, we take analysts at their word when they describe using a firm's trailing P/E ratio to set price targets. We model investors who proportionally adjust their holdings of a firm after looking at the relative difference between analysts' price target and the company's current price.

In Subsection 2.3, we characterize the resulting asset prices. We show that it is possible for $\text{PriceTarget} = \mathbb{E}[\text{EPS}] \times \text{TrailingPE}$ to be correct on average because prices themselves are mostly backward-looking. Finally, in Subsection 2.4, we highlight how our model implies a strong exclusion restriction. For a piece of news to affect a company's expected return, it must change analysts' short-term earnings expectations.

2.1 Earnings Growth

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. 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 $\text{EPS}_t = \$1.00$ per share 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 the firm's earnings growth rate by $\sigma = 2\%$pt higher or lower. Given these assumptions, investors expect the company to generate earnings of $\mathbb{E}_t[\text{EPS}_{t+1}] = \$1.05 \pm \$0.02$ for each shareholder over the next year and $\mathbb{E}_t[\text{EPS}_{t+2}] = \$1.11 \pm \$0.04$ the year after.

If we were writing down a traditional asset-pricing model, then the company's current share price would be determined by the expected discounted value of its entire future cash-flow stream. In that sort of model, it would be important for us to explain how X_t evolves over time as well. But we are not writing down that sort of model. All that matters is short-term earnings growth, so we have deliberately chosen not to specify the longer-term dynamics.

2.2 Investor Demand

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, $\text{TrailingPE}_t \overset{\text{def}}{=} \text{Price}_t \,/\, \text{EPS}_t$.

In standard asset-pricing models, investors make portfolio decisions to maximize their expected utility. However, sell-side analysts make explicit trading recommendations by comparing their price target to a company's current price. 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)" And analysts' recommendations can be worth acting on (Birru, Gokkaya, Liu, and Stulz, 2022).

In our model, investors compare analysts' price target to the company's current price. 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 the price target is lower,

$\frac{\text{PriceTarget}_t - \text{Price}_t}{\text{Price}_t} < 0\%$, they tell their broker to reduce their position over the upcoming twelve months. Moreover, in both scenarios, investors will send their broker instructions to buy in proportion to the relative price difference

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

$\mu > 0$ is a positive constant, which is known as a demand "multiplier".

To make things concrete, suppose that the company is currently trading at $\text{Price}_t = \$100$ and the demand multiplier is $\mu = 1$. If analysts set a one-year-ahead price target of $\text{PriceTarget}_t = \$103$, then investors would respond by increasing their holdings $1 \cdot \left(\frac{\$103 - \$100}{\$100}\right) = 3\%$ 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 near-term earnings 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.

Proposition 2.2 (Demand Schedule)**.** *If analysts use trailing P/E ratios to set price targets (Equation 10) and investors proportionally adjust their holdings over the next year in response (Equation 11), then changes in investors' demand will reflect changes in expected short-term earnings growth*

$$\left(\frac{Demand_{t+1} - Demand_t}{Demand_t}\right) = \mu \cdot \left(\frac{\mathbb{E}_t[EPS_{t+2}] - EPS_t}{EPS_t}\right) \tag{12}$$

"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)" Researchers are so used to thinking this way that many assume it is the only way—i.e., that investor demand must ultimately be derived from expected discounted payoffs. But it is not so.

2.3 Asset Prices

Obviously, in a world where sell-side analysts do apply present-value logic, there is no reason to expect their price targets to line up with the present discounted value of a firm's expected dividend stream. But this does not imply that their forecasts are wrong. If realized prices do not reflect present-value logic either, then analysts' price targets might still be roughly correct.

We now point to a simple assumption about how prices evolve over time that leads to exactly this sort of scenario. Specifically, we assume that there exists a strictly positive constant, $\nu > 0$, such that

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

This assumption says that if investors tell their broker to increase their positions by 1% over the upcoming year, then the company's share price with increase by ν% on average. The noise term $\varepsilon_{t+1} \overset{\text{IID}}{\sim} \text{Normal}(0, \varsigma^2)$ captures all other reasons why a company's share price might increase or decrease next year. Under this assumption, it is possible for analysts' price targets to be accurate in a world where price does not equal expected discounted payoff.

Proposition 2.3 (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*

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

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

This result follows a long tradition in theoretical asset pricing: guess that the price function is linear and then verify that the implications are consistent with some goal. 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 are doing something similar. But, instead of guessing that prices are linear, we match coefficients to verify that price growth is linear.

Analysts say that they set price targets by multiplying a company's expected EPS times a trailing P/E as shown in 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 price path in Equation (13) is so simple. The functional forms used by Grossman and Stiglitz (1980) were dictated by theoretical considerations. They chose to study a CARA-normal setting because, in that sort of model, it would be natural to expect prices to be linear. By contrast, the functional forms in our model are dictated by what real-world market participants say. Nevertheless, we are still able to outline a simple scenario in which $\text{PriceTarget} = \mathbb{E}[\text{EPS}] \times \text{TrailingPE}$ is correct on average.

It is also noteworthy that the law of motion for prices (Equation 13) involves the same sort of proportional thinking found in investors' demand rule (Equation 11). It makes sense to use $\mathbb{E}[\text{EPS}] \times \text{TrailingPE}$ when investors' demand and price growth both respond proportionally.

This connects our work to the literature on demand-system asset pricing (Koijen and Yogo, 2019; Gabaix and Koijen, 2024). 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. Under this interpretation, it would be natural to expect $\mu = 1/\nu$ as required by Proposition 2.3.

That being same, our model focuses on demand multipliers and price elasticities for a very different reason. The demand-system framework cares about μ and ν because they play a pivotal role in models where investors solve a forward-looking portfolio problem and appreciate the implications of market clearing. By contrast, our model does not enforce market clearing. μ and ν emerge from taking seriously how analysts describe their pricing rule.

2.4 Exclusion Restriction

Present-value relationships are valuable because they rule out lots of possible mechanisms. Let News_t be a piece of information revealed about a firm at time t. If this is how the world works, then News_t can only predict the company's future returns if it is correlated with changes in investors' cash-flow expectations or their discount rate. All price effects must operate through these two specific channels.

"In general, the claim that an instrument operates through a single known channel is called an *exclusion restriction*. (Angrist and Pischke, 2009)" Our framework implies its own very different kind of exclusion restriction: if News_t predicts the company's future return, then it must be correlated with the firm's near-term expected earnings growth. It does not help to be correlated with other future outcomes that would matter in textbook models.

Proposition 2.4 (Exclusion Restriction)**.** *If $News_t$ is uncorrelated with expected short-run earnings growth, $\widehat{\text{Corr}}(X_t, News_t) = 0$, it will not predict average returns*

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

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

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 $\text{Price} = A + B \cdot \text{Signal} - C \cdot \text{Shock}$. 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

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-American research team.*

thus the equilibrium price) in our model is by changing the one forward-looking component: expected earnings growth over the next two years.

Do analysts really ignore information about a company's earnings three years from now? Yes. 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 Explanatory Power

In this section, we show that trailing P/E ratios have a substantial amount of explanatory power. We describe our data in Subsection 3.1. Then, in Subsection 3.2 we document how PriceTarget = $\mathbb{E}[\text{EPS}] \times$ TrailingPE can explain over 91% of the variation in analysts' price targets. In Subsection 3.3 we show that this simple formula also explains how market prices respond to earnings surprises.

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

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.

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. PriceTarget$^a_{n,t}$: price forecast set for the end of a firm's upcoming fiscal year, roughly twelve months in the future. $\mathbb{E}^a_t[EPS_{n,\tau+1}]$: analyst's EPS forecast for the twelve-month period ending on the date of their price target. $\mathbb{E}^a_t[EPS_{n,\tau+2}]$: analyst's EPS forecast for the twelve-month period following the date of their price target. $\mathbb{E}^a_t[EPS_n]$: an analyst's two-year-ahead EPS forecast when available; else, the reported one-year-ahead forecast value. ImpliedPE$^a_{n,t}$: the analyst's price target divided by their EPS forecast. TrailingPE$_{n,t}$: a company's current price divided by its trailing twelve-month EPS.*

We define the P/E ratio implied by an analyst's price target and EPS forecast as follows

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

If sell-side analysts set price targets based solely on a company's trailing twelve-month (TTM) P/E ratio, then each time an analyst posted a new price target we would find that $\text{ImpliedPE}^a_{n,t} = \text{TrailingPE}_{n,t}$.

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 these P/E ratios are between 5× and 50×. This sample restriction is motivated by practical considerations. Market participants 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.

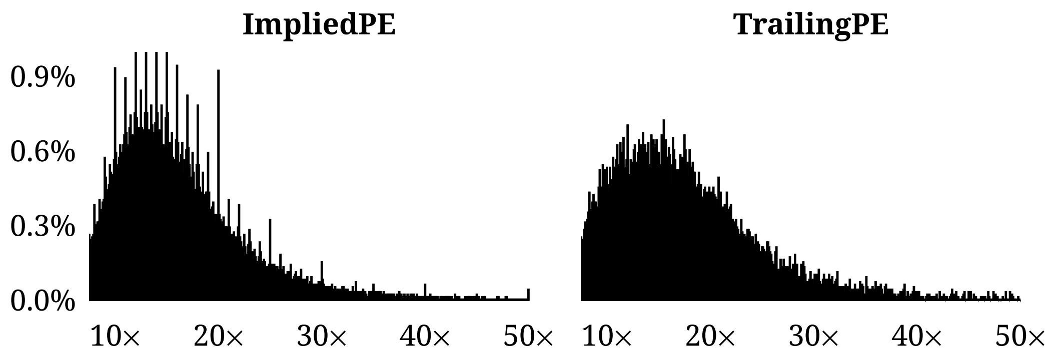

Figure 17. *Histograms showing the distribution of* $\text{ImpliedPE}^a_{n,t}$ *(left panel) and* $\text{TrailingPE}_{n,t}$ *(right panel) for all sell-side analyst reports in our sample with* $\mathbb{E}^a_t[EPS_n] \geq \$1.00$ *from 2003 to 2022. x-axis denotes the P/E ratio in increments of* $0.1\times$*.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

3.2 Analysts' Price Targets

Sell-side analysts typically use a trailing P/E ratio to convert their short-term earnings forecast into a price target. Table 10 shows the results of estimating the regression specification below

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

We fit the regression to data on days when analysts update their price target for the firm. i.e., for Andrea Teixeira's coverage of Coca-Cola, these dates correspond to the black diamonds in Figure 6. $\log(\text{PriceTarget}^a_{n,t})$ is the log 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 use the formula $\text{PriceTarget} = \mathbb{E}[\text{EPS}] \times \text{TrailingPE}$, then we should 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 three ways: firm, analyst, and month. What's more, our simple

Dep variable:	$\log(\text{PriceTarget}^a_{n,t})$			
	(1)	(2)	(3)	(4)
$\log(\mathbb{E}^a_t[\text{EPS}_n])$	0.93*** (0.01)	0.87*** (0.01)	0.91*** (0.01)	0.93*** (0.01)
$\log(\text{TrailingPE}_{n,t})$	0.63*** (0.01)	0.47*** (0.01)	0.64*** (0.01)	0.57*** (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.*

trailing P/E formula 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 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 ($\text{ImpliedPE}^a_{n,t}$; 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 PriceTarget = $\mathbb{E}[\text{EPS}] \times \text{TrailingPE}$, 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.

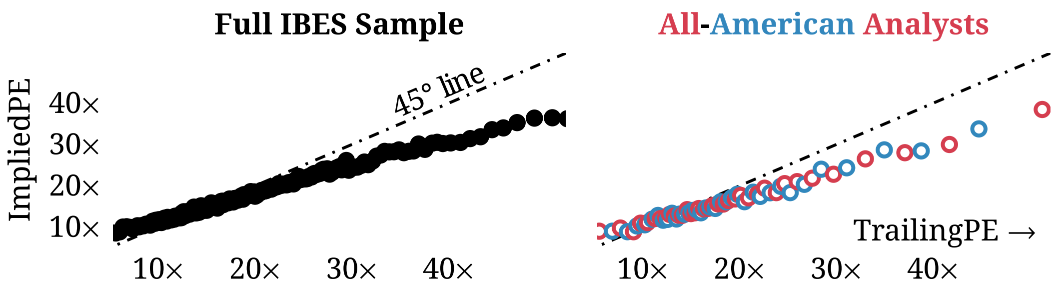

Figure 18. *(Left) Binned scatterplot using data from the full sample of IBES reports. x-axis shows the firm's trailing twelve-month P/E,* $TrailingPE_{n,t} = Price_{n,t} \,/\, EPS_{n,t}$*. y-axis shows the P/E ratio implied by the analyst's price target and EPS forecast,* $ImpliedPE^{a}_{n,t} \overset{\text{def}}{=}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

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-American 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 for each individual company. It is possible to count the number of exceptions on one hand.

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

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

If sell-side analysts were exclusively using $\text{PriceTarget} = \mathbb{E}[\text{EPS}] \times \text{TrailingPE}$ to set price targets, 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 $R^2 = 54.5\%$. In other words, the best-fit line may be a bitter flatter than predicted, but it still explains more than half of the variation in implied P/E ratios.

Why is the fit not perfect? We can think of a few reasons. First, analysts are not automatons. When they think other information might be relevant, they

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.*

elaborate on PriceTarget = $\mathbb{E}$[EPS] × TrailingPE. They treat this simple formula as a starting point, and they tend to add ingredients when a company's trailing P/E is 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 corresponds to 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 EPS realizations have the largest effect on an analyst's implied P/E ratio. But there are also significant coefficients at longer lags as well. We are able to explain 91% of the variation in analysts' price targets even before incorporating the effects of longer-term lags.

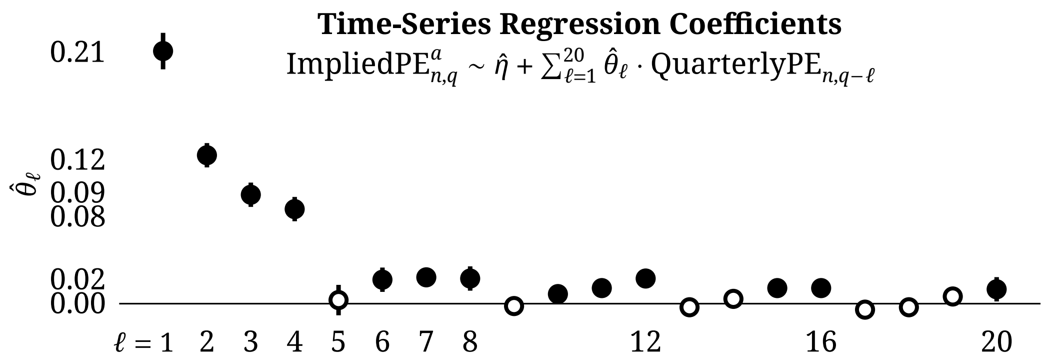

Figure 19. *Each dot denotes one of the* 20 *estimated slope coefficients,* $\{\hat{\theta}_\ell\}_{\ell=1}^{20}$, *from the regression specification in Equation* (19). *ImpliedPE*$^a_{n,q}$*: P/E ratio implied by an analyst's price target and EPS forecast. QuarterlyPE*$_{n,q}$*: company's closing price 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.*

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

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

The variable $\text{QuarterlyPE}_{n,q} \overset{\text{def}}{=} \text{Price}_{n,t_{\text{Ancmt}}-1} \,/\, (4 \cdot \text{eps}_{n,q})$ represents the company's closing price on the day before its earnings for the quarter are announced announcement divided by four times its realized EPS in the quarter. We use $\text{eps}_{n,q}$ to denote the nth stock's earnings in quarter q. The lowercase letters indicate that it is only 1/4th of the firm's earnings for the fiscal year.

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

3.3 Realized Price Changes

In the final part of our empirical analysis, we show that trailing P/E ratios do not just explain the future prices that analysts expect. They also explain the realized prices that we later observe in the data. To provide the cleanest possible empirical setting, we study how a company's share price changes following an earnings surprise.

To see the logic behind our test, imagine that a company just announced quarterly earnings that were +\$0.10 per share higher than analysts' expectations. It does not take a genius to predict that the firm's price will likely go up. But why? What are the possible mechanisms?

To start with, suppose we lived in a world were analysts used a Gordon model to price each company's expected future earnings stream

$$\text{Price}_t = \mathbb{E}_t[\text{EPS}_{t+1}] \times \left(\frac{1}{r-g}\right) \tag{20}$$

This is just Equation (4) but with EPS in place of dividends. In this scenario, there are two ways that a +\$0.10 per share earnings surprise could cause a company's price to increase. Short-term earnings expectations could rise—i.e., higher $\mathbb{E}_t[\text{EPS}_{t+1}]$. There could also be multiples expansion—i.e., larger $\left(\frac{1}{r-g}\right)$. Moreover, investors should be more likely to change their multiple following a large earnings surprise, which often signal important long-term changes.

By contrast, our model says that the price change should be proportional to the firm's trailing P/E ratio at the time of its +\$0.10 per share earnings surprise

$$\Delta\text{Price}_{t+1} = \{\lambda \cdot \$0.10\} \times \text{TrailingPE}_t \qquad \text{for some } \lambda > 0 \tag{21}$$

Good news cannot change the past. The company's trialing P/E is set in stone. A positive +\$0.10 per share earnings surprise this quarter will only cause analysts to increase their EPS forecast for the year by +\$0.10. This will push their price target up by $\Delta\text{PriceTarget}_t = \$0.10 \times \text{TrailingPE}_t$, and the company's share price will grow in response to the resulting increase in investor demand.

We cannot predict the value of λ based on first principles. But we can estimate this parameter using the available data. The same λ should apply regardless of whether the firm announces earnings that are +\$0.10 higher than expected (like we have considered so far) or +\$0.30 higher than expected. It also should not matter if the firm falls short of analysts' expectations by −\$0.20 per share. It should be the same constant of proportionality in all cases.

First, we group stock-quarter observations by the size of their earnings surprise, $s \in \{-\$0.30, \ldots, \$0.30\}$, which we define relative to the consensus forecast. Then, within each bucket, we run a separate regression to estimate the relationship between a stock's trailing P/E ratio and its subsequent price change

$$\Delta\text{Price}_{n,q+1} \sim \hat{\kappa}_s + \hat{\lambda}_s \cdot \text{TrailingPE}_{n,q} \qquad \begin{array}{c}\text{using data on firm-qtr}\\ \text{obs where } (n,q) \in N_s\end{array} \tag{22}$$

where $N_s = \{\, (n,q) \,:\, \text{eps}_{n,q} - \mathbb{E}_{t_{\text{Ancmt}}-1}[\text{eps}_{n,q}] = s \,\}$ is the set of stock-quarter observations where the firm announced had a s dollar per share earnings surprise. This gives us 60 distinct estimates for $\lambda \cdot s$. Finally, we check whether these estimated values all share a key property, $\hat{\lambda}_s / s = \lambda$.

Notice how this approach mirrors the standard logic behind textbook asset-pricing econometrics. For example, think about the Fama and MacBeth (1973) approach to testing the CAPM. First, you group stocks into portfolios. Then, you estimate each portfolio's market beta, β, by running a separate time-series regression. Finally, you check whether there is a linear relationship between each portfolio's excess returns and its estimated β.

Average excess returns and market betas do not line up neatly on the security market line as predicted by the CAPM. But Figure 20 shows that our estimated slope coefficients, $\{\hat{\lambda}_s\}$, do increase linearly in the size of the underlying earnings surprise, s. The black dots show the coefficients associated with 60 different levels of earnings surprise: $s \in \{-\$0.30, \ldots, -\$0.01, \$0.01, \ldots, \$0.30\}$ per share. We omit the $s = \$0.00$ bin where there is no surprise.

When a firm announces quarterly earnings that way above or below analysts' consensus expectation, it is often a sign that something important has changed about the firm's situation. Large surprises often signal a persistent change in the

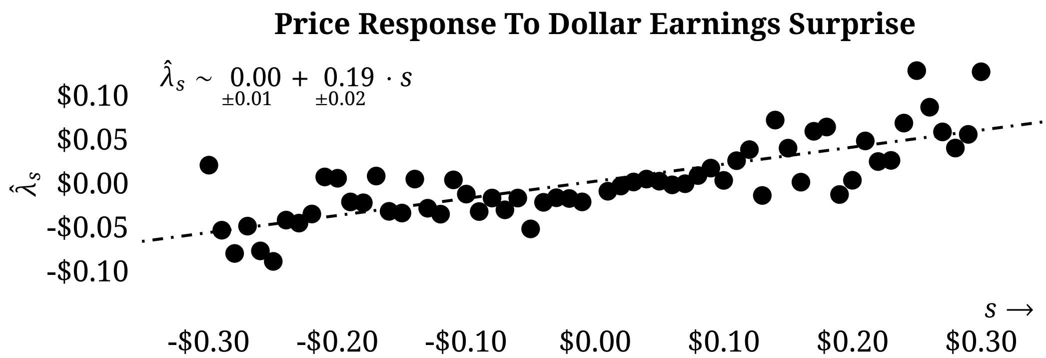

Figure 20. *Each dot denotes an estimated slope coefficient, $\hat{\lambda}_s$, from one of* 60 *separate regressions like the one shown in Equation* (22)*. The y-axis shows the first-stage slope coefficient, $\hat{\lambda}_s$, estimated using data on all firm-quarter observations that had the same dollar earnings surprise, s, relative to consensus. The x-axis shows the corresponding value of s in* \$0.01 *bins. The highest bin is centered at* \$0.30 *per share while the lowest bin is centered at* −\$0.30 *per share. This gives us* 60 *data points since we omit the* $s = \$0.00$ *bin containing observations with no earnings surprise. The dashed line is the best-fit OLS equation.*

company's future earnings. Hence, in a world where investors were following textbook logic, we would expect to see them use larger multiples following very positive earnings surprises and smaller multiples after very negative surprises.

Hence, in a textbook world, we would not find a neat linear relationship like the one we see in Figure 20. Instead, we should see an "S"-shaped pattern or even something that looks like a sine wave. If stock A has $\text{TrailingPE}_A = 20\times$ while stock B has $\text{TrailingPE}_A = 10\times$, then maybe stock A might have double the price reaction for small earnings surprises. But, for large surprises, forward-looking investors would revise their choice of multiple to reflect persistent changes, meaning that the difference between stock A and stock B should attenuate in the tails. We do not see that happen.

Our 0.19 second-stage slope coefficient implies that, following an $s = \$0.10$ per share earnings surprise, the share price of stock A with $\text{TrailingPE}_A = 20\times$ will increase by $\$0.19 = 0.19 \cdot \{\$0.10 \times 20 - \$0.10 \times 10\}$ more than the share price of stock B with $\text{TrailingPE}_B = 10\times$. If both firms had realized an $s = \$0.20$ per share earnings surprise, then stock A's share price would have gone up by \$0.38 more than stock B's over the subsequent quarter. An $s = \$0.30$ per share

Dep variable:	First-stage coefficient, $\hat{\lambda}_s$		
Bin width:	\$0.01	\$0.02	\$0.05
	(1)	(2)	(3)
Intercept	0.00	0.00	0.00
	(0.01)	(0.01)	(0.01)
Slope	0.19***	0.20***	0.16***
	(0.02)	(0.02)	(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. The left-hand-side variable in each regression is an estimated slope coefficient, $\hat{\lambda}_s$, for each earnings surprise bin like in Equation (22). Column (1) reports results using* 60 *separate* \$0.01 *bins centered at* $\{-\$0.30, \ldots, -\$0.01, \$0.01, \ldots, \$0.30\}$. *These results match the dashed best-fit line in Figure 20. Column (2) shows results where we group observations into* 30 *separate* \$0.02 *bins centered at* $\{-\$0.30, \ldots, -\$0.02, \$0.02, \ldots, \$0.30\}$. *Column (3) shows a similar analysis using* 12 *bins that are* \$0.05 *wide,* $\{-\$0.30, \ldots, -\$0.05, \$0.05, \ldots, \$0.30\}$. *All three columns omit the bin centered at s = \$0.00—i.e., stock-quarter observations where there was no earnings surprise.*

earnings surprise would lead to a \$0.57 difference between the two stocks' price growth. It is exactly as predicted by the exclusion restriction we outline in Proposition 2.4. The trailing P/E is set in stone. A piece of news changes a company's price by altering investors' short-term earnings expectations.

We know from cross-sectional asset pricing that a researcher's choice of test portfolios can effect 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 wide and \$0.05 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.

Conclusion

Investors 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. Every model should not start with price equals expected discounted payoff.

We show that sell-side analysts typically do not discount any expected future payoffs when valuing large publicly traded companies. Instead, they capitalize their short-term EPS expectations into a price target with a trailing P/E ratio. The analysts in textbook models ask themselves: "What is the present value of a company's expected future earnings stream in today's dollars?" Real-world analysts ask a different question: "How would a comparable firm have been price last year if it had announced similar earnings?"

Market participants do not usually share their subjective payoff expectations with us. Sell-side analysts are the exception. As a result, their earnings forecasts have had a massive impact on the asset-pricing literature. Researchers have spent decades plugging these numbers into various present-value relationships in an effort to understand analysts' discount rate. However, analysts tell us how they price their own subjective earnings expectations in the text of their report. More often than not, it has nothing to do with discount rates. As a rule, they generally use a trailing P/E.

"Price equals expected discounted payoff" is clearly not a sensible assumption to make when modeling sell-side analysts. The next time you write down an asset-pricing model, we want you to consider whether this "one simple concept" is a sensible assumption in your application. It might not. To show how productive alternative approaches can be, we build an equilibrium asset-pricing model based on $\text{PriceTarget} = \mathbb{E}[\text{EPS}] \times \text{TrailingPE}$. We take this model to the data and find that it explains both analysts' price targets and how market prices respond to earnings announcements.

References

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.

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

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.

Chava, S. and A. Purnanandam (2010). Is default risk negatively related to stock returns? *Review of Financial Studies 23*(6), 2523–2559.

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

Cochrane, J. (1999). Portfolio advice for a multifactor world. *Federal Reserve Bank of Chicago Economic Perspectives 23*(Q III).

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

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.

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.

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.

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.

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., 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.

Lee, C., D. Ng, and B. Swaminathan (2009). Testing international asset pricing models using implied costs of capital. *Journal of Financial and Quantitative Analysis 44*(2), 307–335.

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.

Pástor, L., M. Sinha, and B. Swaminathan (2008). Estimating the intertemporal risk-return tradeoff using the implied cost of capital. *Journal of Finance 63*(6), 2859–2897.

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.

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)** *This result follows from manipulating Equation* (10)

$$PriceTarget_t = \mathbb{E}_t[EPS_{t+2}] \times TrailingPE_t \tag{A.1a}$$

$$PriceTarget_t = \mathbb{E}_t[EPS_{t+2}] \times \left(\frac{Price_t}{EPS_t}\right) \tag{A.1b}$$

$$PriceTarget_t = \left(\frac{\mathbb{E}_t[EPS_{t+2}]}{EPS_t}\right) \times Price_t \tag{A.1c}$$

$$\frac{PriceTarget_t}{Price_t} = \frac{\mathbb{E}_t[EPS_{t+2}]}{EPS_t} \tag{A.1d}$$

$$\frac{PriceTarget_t - Price_t}{Price_t} = \frac{\mathbb{E}_t[EPS_{t+2}] - EPS_t}{EPS_t} \tag{A.1e}$$

Thus, if demand growth over the next year is proportional to $(\frac{PriceTarget_t - Price_t}{Price_t})$, *then it must also be proportional to* $(\frac{\mathbb{E}_t[EPS_{t+2}] - EPS_t}{EPS_t})$. □

Proof. **(Proposition 2.3)** *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*

$$\frac{\hat{\mathbb{E}}_t[Price_{t+1}] - Price_t}{Price_t} = \nu \times \left(\frac{Demand_{t+1} - Demand_t}{Demand_t}\right) \tag{A.2}$$

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

$$\frac{\hat{\mathbb{E}}_t[Price_{t+1}] - Price_t}{Price_t} = (\nu \cdot \mu) \times \left(\frac{\mathbb{E}_t[EPS_{t+2}] - EPS_t}{EPS_t} \right) \tag{A.3}$$

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 simply need to rearrange terms 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

$$\frac{\hat{\mathbb{E}}_t[Price_{t+1}] - Price_t}{Price_t} = (\nu \cdot \mu) \times \left(\frac{\mathbb{E}_t[EPS_{t+2}] - EPS_t}{EPS_t} \right) \tag{A.4a}$$

$$\frac{\hat{\mathbb{E}}_t[Price_{t+1}]}{Price_t} = (\nu \cdot \mu) \times \left(\frac{\mathbb{E}_t[EPS_{t+2}]}{EPS_t} \right) + (1 - \nu \cdot \mu) \tag{A.4b}$$

$$\begin{aligned} \hat{\mathbb{E}}_t[Price_{t+1}] = (\nu \cdot \mu) \times \mathbb{E}_t[EPS_{t+2}] \times \left(\frac{Price_t}{EPS_t} \right) \\ + (1 - \nu \cdot \mu) \times Price_t \end{aligned} \tag{A.4c}$$

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

Proof. **(Proposition 2.4)** *Equation (A.4a) from the proof to Proposition 2.3 above states that*

$$\frac{\hat{\mathbb{E}}_t[Price_{t+1}] - Price_t}{Price_t} = (\nu \cdot \mu) \times \left(\frac{\mathbb{E}_t[EPS_{t+2}] - EPS_t}{EPS_t} \right) \tag{A.4a}$$

The left-hand side of this equation is $\hat{\mathbb{E}}_t[Return_{t+1}]$. The right-hand side is a function of the analysts' expectations about short-term earnings as defined in Equation (9), which we have reproduced below

$$\left(\frac{EPS_{t+1} - EPS_t}{EPS_t} \right) = X_t + \epsilon_{t+1} \tag{9}$$

$X_t \approx \mathbb{E}_t[\Delta \log EPS_{t+1}]$ 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 $\hat{\mathbb{E}}_t[Return_{t+1}]$. □

B Additional Results

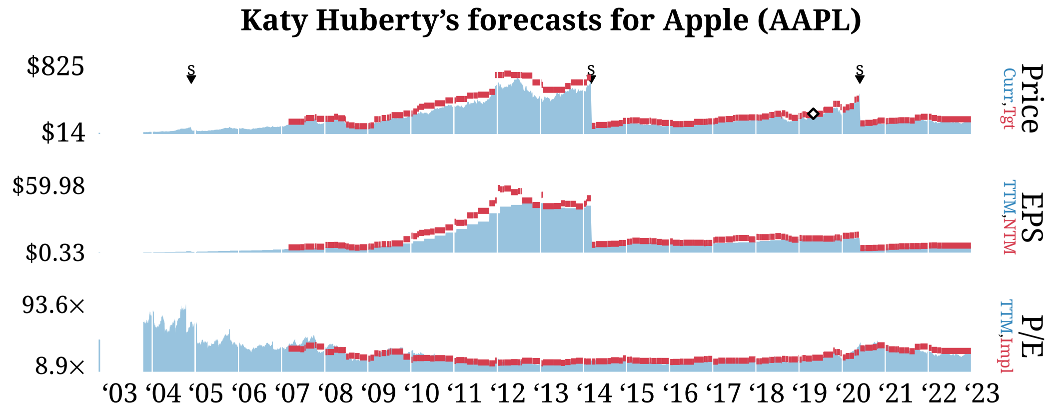

Figure B1(a). *y-axis shows min, median, and max. (Top) Blue ribbon is Apple's closing price on day t from CRSP, $Price_t$. Red line is Katy Huberty's price target, $PriceTarget_t = \mathbb{E}_t[Price_{\tau+1}]$, in IBES. (Middle) Blue is AAPL's trailing twelve-month (TTM) EPS on day t from IBES, EPS_t.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

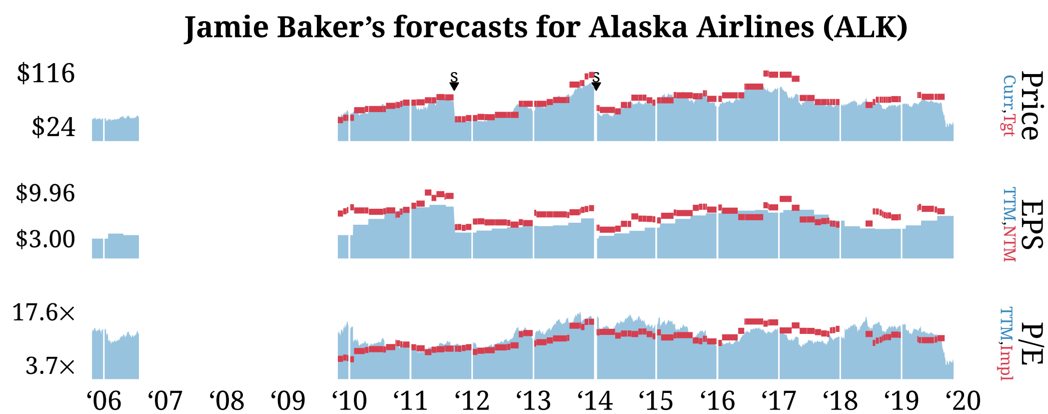

Figure B1(b). *y-axis shows min, median, and max. (Top) Blue ribbon is Alaska Airlines' closing price on day t from CRSP, $Price_t$. Red line is Jamie Baker's price target, $PriceTarget_t = \mathbb{E}_t[Price_{\tau+1}]$, in IBES. (Middle) Blue is ALK's trailing twelve-month (TTM) EPS on day t from IBES,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

CJ Muse's forecasts for Broadcom (AVGO)

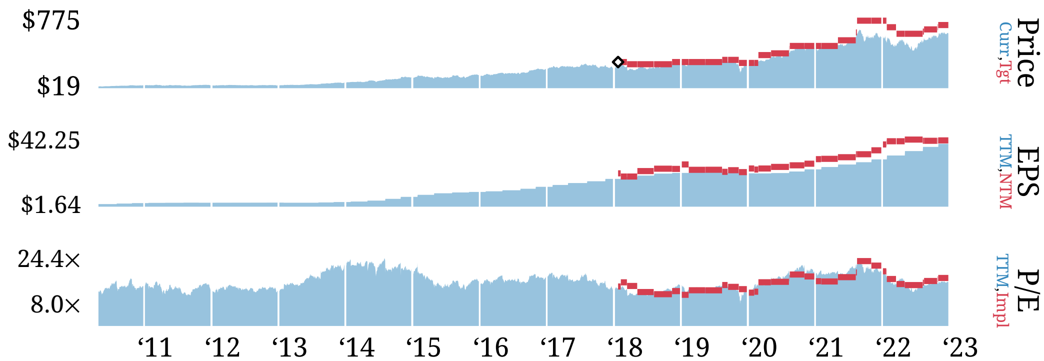

Figure B1(c). *y-axis shows min, median, and max. (Top) Blue ribbon is Broadcom's closing price on day t from CRSP, $Price_t$. Red line is CJ Muse's price target, $PriceTarget_t = \mathbb{E}_t[Price_{\tau+1}]$, in IBES. (Middle) Blue is Broadcom's trailing twelve-month (TTM) EPS on day t from IBES, EPS_t.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

Vivek Juneja's forecasts for Citigroup (C)

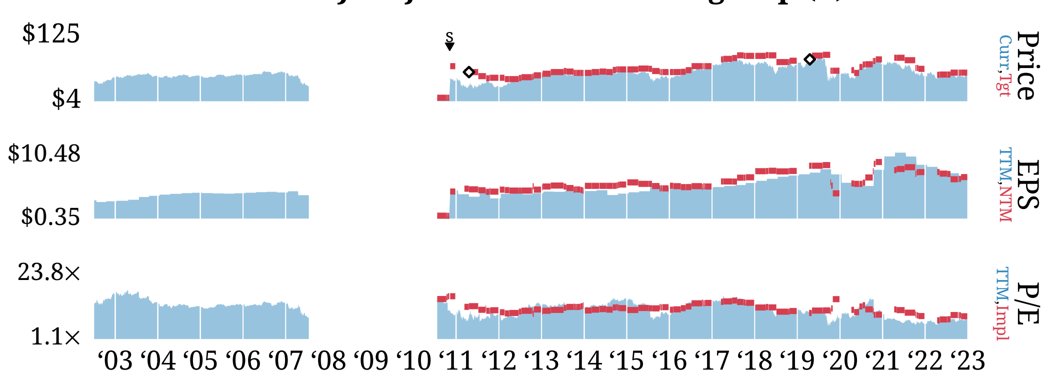

Figure B1(d). *y-axis shows min, median, and max. (Top) Blue ribbon is Citigroup's closing price on day t from CRSP, $Price_t$. Red line is Vivek Juneja's price target, $PriceTarget_t = \mathbb{E}_t[Price_{\tau+1}]$, in IBES. (Middle) Blue is Citi's trailing twelve-month (TTM) EPS on day t from IBES,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

Rod Hall's forecasts for Cisco Systems (CSCO)

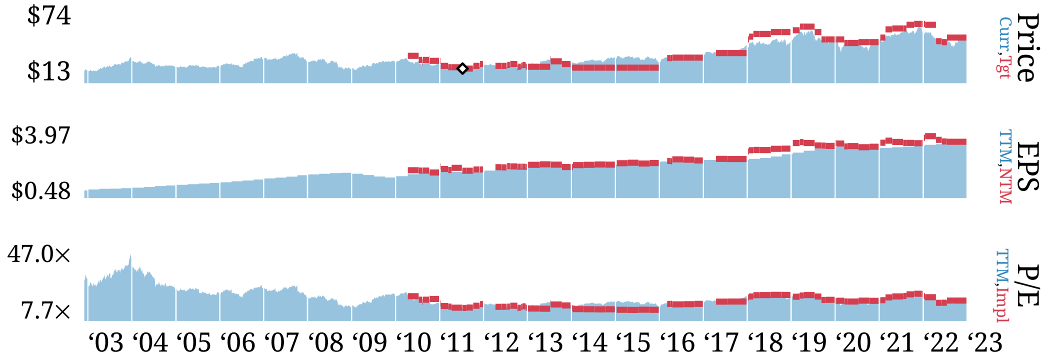

Figure B1(e). *y-axis shows min, median, and max. (Top) Blue ribbon is Cisco System's closing price on day t from CRSP, $Price_t$. Red line is Rod Hall's price target, $PriceTarget_t = \mathbb{E}_t[Price_{\tau+1}]$, in IBES. (Middle) Blue is Cisco's trailing twelve-month (TTM) EPS on day t from IBES,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

Heath Terry's forecasts for eBay (EBAY)

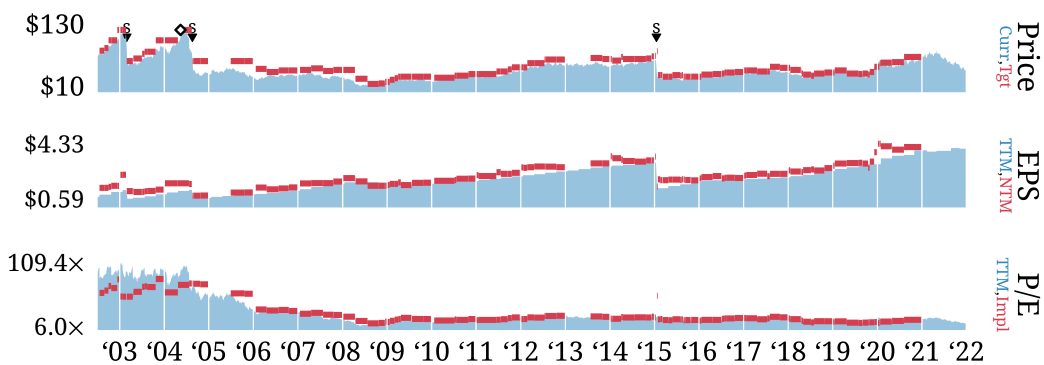

Figure B1(f). *y-axis shows min, median, and max. (Top) Blue ribbon is eBay's closing price on day t from CRSP, $Price_t$. Red line is Heath Terry's price target, $PriceTarget_t = \mathbb{E}_t[Price_{\tau+1}]$, in IBES. (Middle) Blue is eBay's trailing twelve-month (TTM) EPS on day t from IBES, EPS_t.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

Christophe Cherblanc's forecasts for Google (GOOGL)

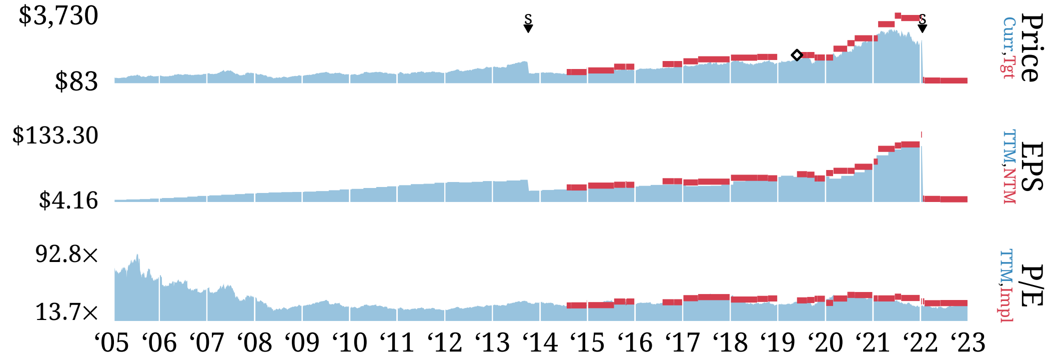

Figure B1(g). *y-axis shows min, median, and max. (Top) Blue ribbon is Google's closing price on day t from CRSP, $Price_t$. Red line is Christophe Cherblanc's price target, $PriceTarget_t = \mathbb{E}_t[Price_{\tau+1}]$, in IBES. (Middle) Blue is Google's trailing twelve-month (TTM) EPS on day t from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

Greg Badishkanian's forecasts for MarineMax (HZO)

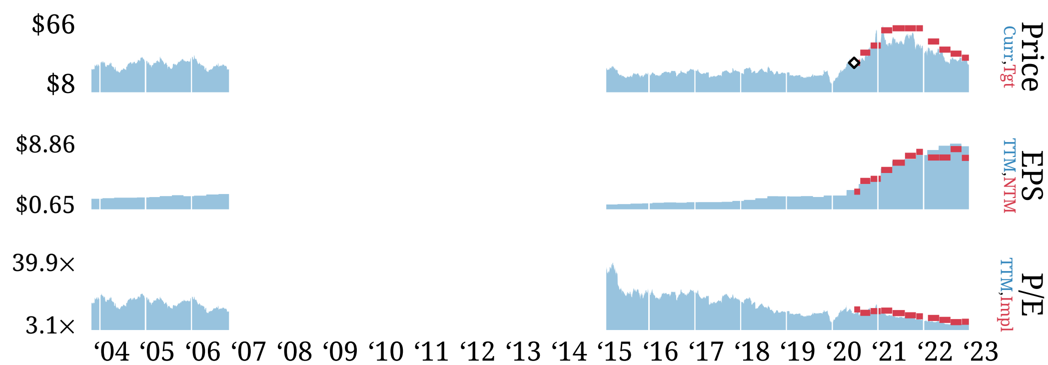

Figure B1(h). *y-axis shows min, median, and max. (Top) Blue ribbon is MarineMax's closing price on day t from CRSP, $Price_t$. Red line is Greg Badishkanian's price target, $PriceTarget_t = \mathbb{E}_t[Price_{\tau+1}]$, in IBES. (Middle) Blue is HZO's trailing twelve-month (TTM) EPS on day t from IBES,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

Moshe Orenbuch's forecasts for Mastercard (MA)

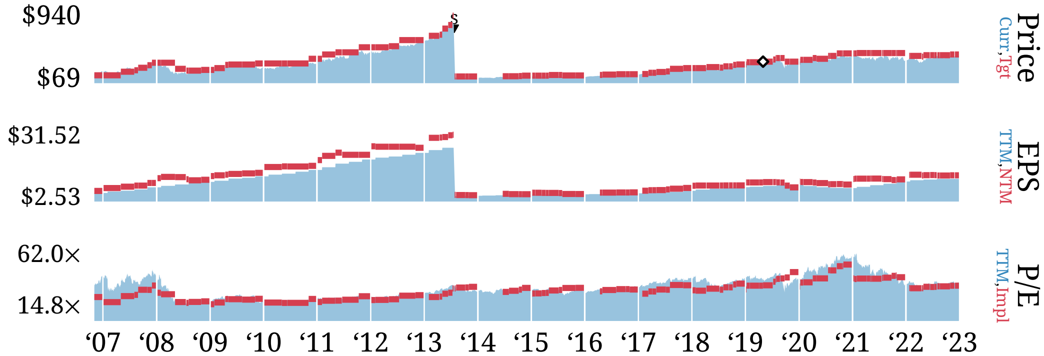

Figure B1(i). *y-axis shows min, median, and max. (Top) Blue ribbon is Mastercard's closing price on day t from CRSP, $Price_t$. Red line is Moshe Orenbuch's price target, $PriceTarget_t = \mathbb{E}_t[Price_{\tau+1}]$, in IBES. (Middle) Blue is MA's trailing twelve-month (TTM) EPS on day t from IBES,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

Tim Anderson's forecasts for Merck (MRK)

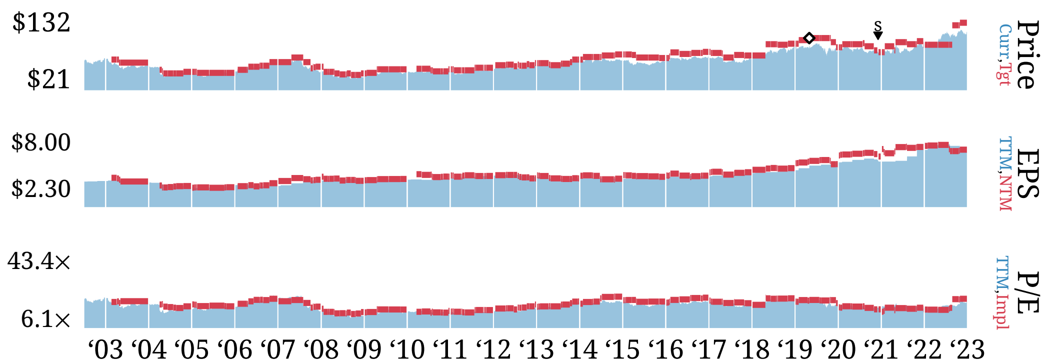

Figure B1(j). *y-axis shows min, median, and max. (Top) Blue ribbon is Merck's closing price on day t from CRSP, $Price_t$. Red line is Tim Anderson's price target, $PriceTarget_t = \mathbb{E}_t[Price_{\tau+1}]$, in IBES. (Middle) Blue is MRK's trailing twelve-month (TTM) EPS on day t from IBES, EPS_t.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

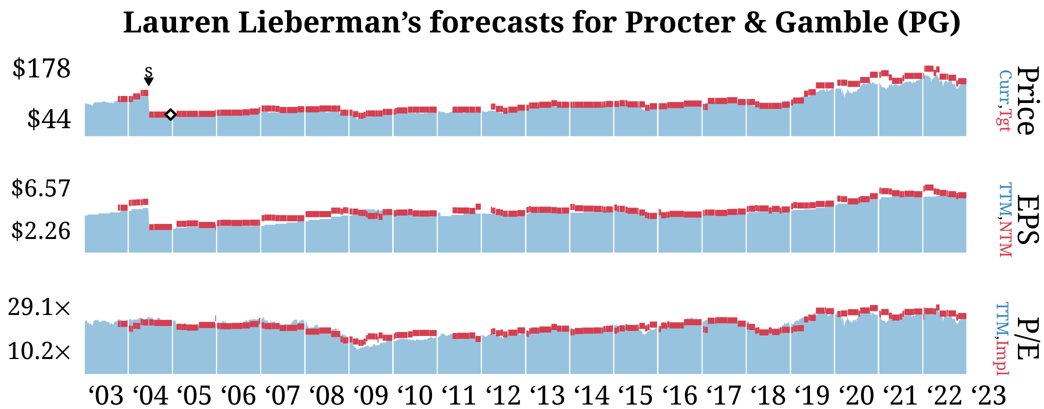

Figure B1(k). *y-axis shows min, median, and max. (Top) Blue ribbon is Procter & Gamble's closing price on day t from CRSP, $Price_t$. Red line is Lauren Lieberman's price target, $PriceTarget_t = \mathbb{E}_t[Price_{\tau+1}]$, in IBES. (Middle) Blue is PG's trailing twelve-month (TTM) EPS on day t from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

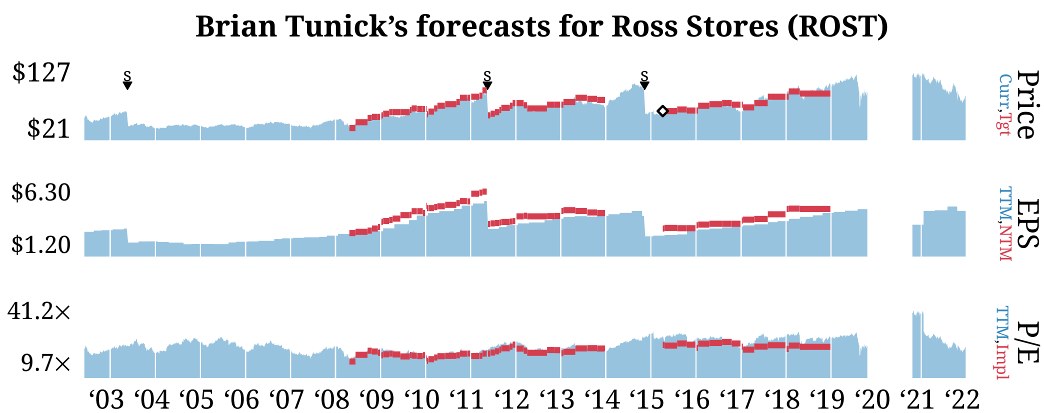

Figure B1(l). *y-axis shows min, median, and max. (Top) Blue ribbon is Ross' closing price on day t from CRSP, $Price_t$. Red line is Brian Tunick's price target, $PriceTarget_t = \mathbb{E}_t[Price_{\tau+1}]$, in IBES. (Middle) Blue is Ross' trailing twelve-month (TTM) EPS on day t from IBES, EPS_t.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

Justin Lake's forecasts for United Healthcare (UNH)

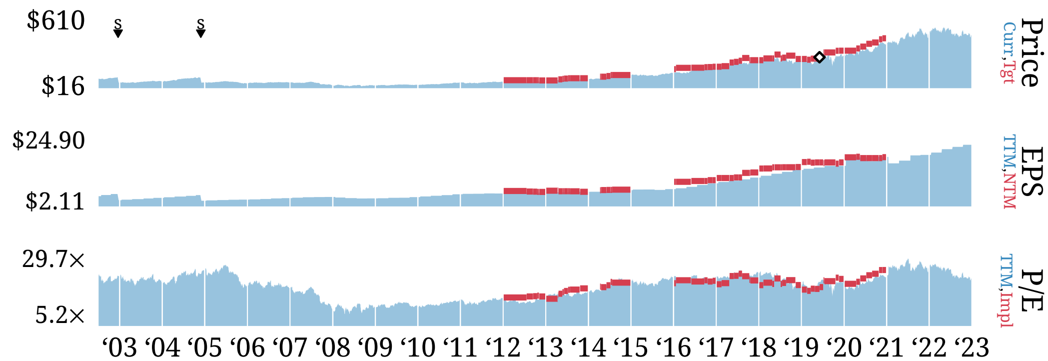

Figure B1(m). *y-axis shows min, median, and max. (Top) Blue ribbon is United Healthcare's closing price on day t from CRSP, $Price_t$. Red line is Justin Lake's price target, $PriceTarget_t = \mathbb{E}_t[Price_{\tau+1}]$, in IBES. (Middle) Blue is UNH's trailing twelve-month (TTM) EPS on day t from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

Christopher Horvers' forecasts for Walmart (WMT)

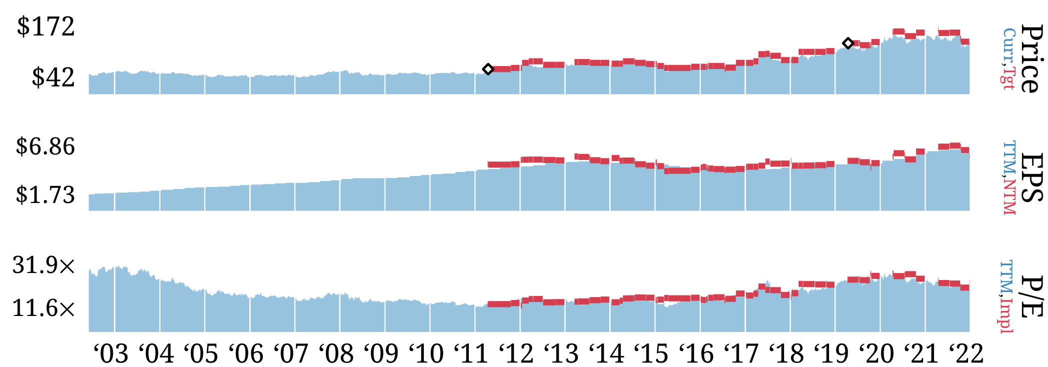

Figure B1(n). *y-axis shows min, median, and max. (Top) Blue ribbon is Walmart's closing price on day t from CRSP, $Price_t$. Red line is Chris Horvers' price target, $PriceTarget_t = \mathbb{E}_t[Price_{\tau+1}]$, in IBES. (Middle) Blue is WMT's trailing twelve-month (TTM) EPS on day t from IBES, EPS_t.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

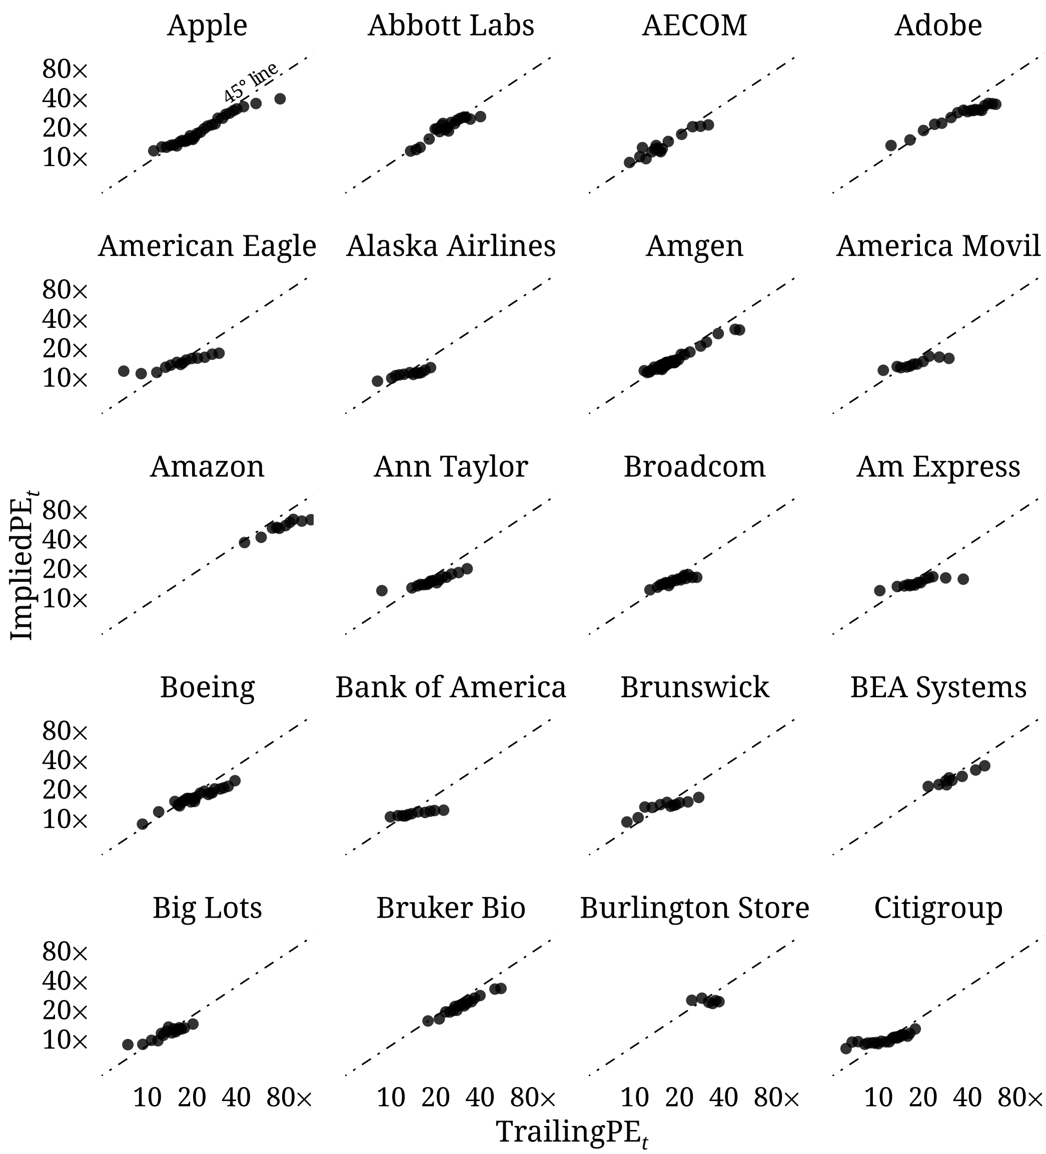

Figure B2(a). *Each panel shows a binned scatterplots using data from the full sample of IBES reports for a single firm. x-axis shows the firm's trailing twelve-month P/E, $TrailingPE_{n,t} = Price_{n,t} / EPS_{n,t}$. y-axis shows the P/E ratio implied by the analyst's price target and EPS forecast,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

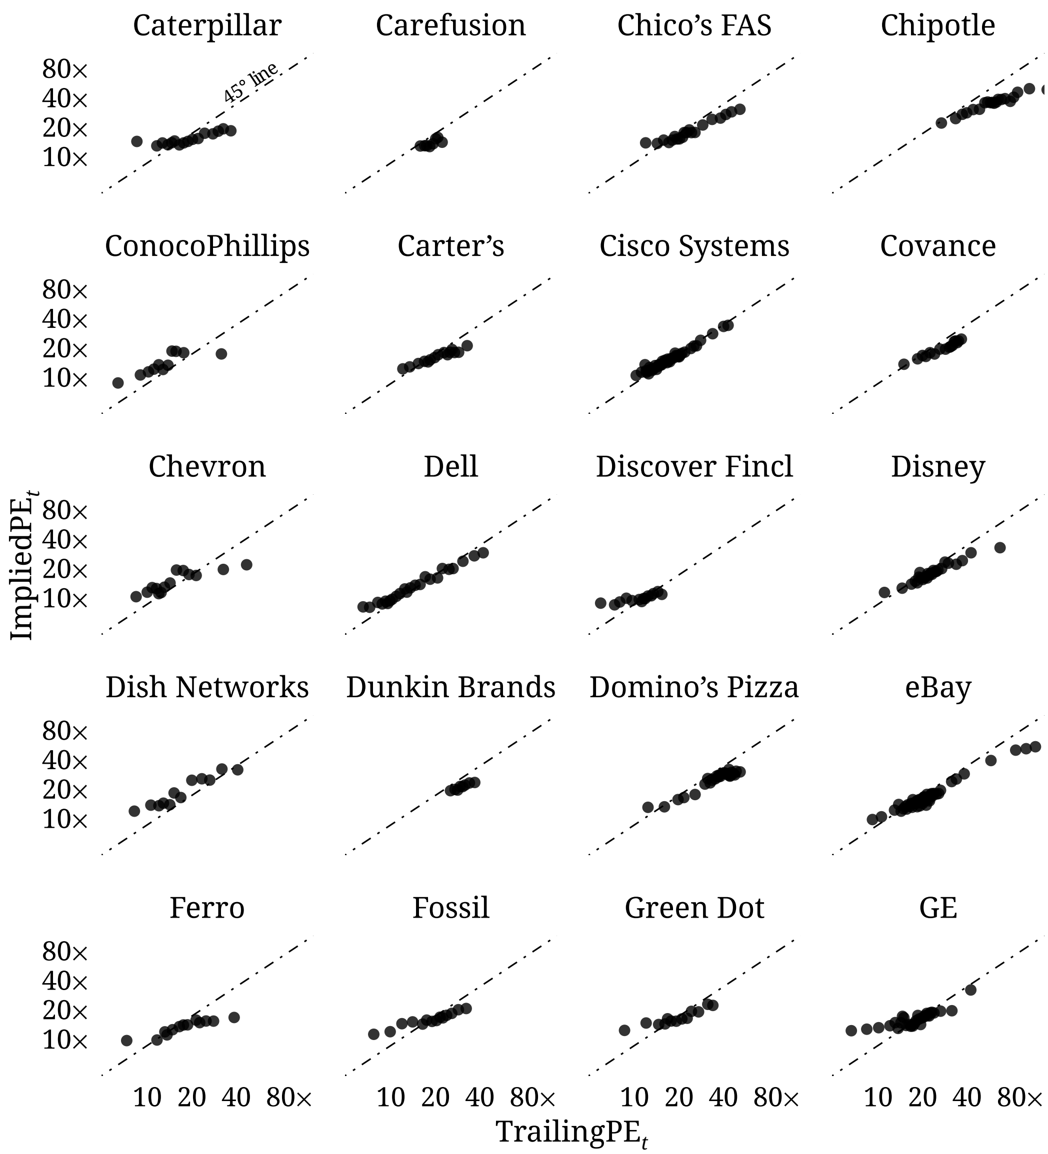

Figure B2(b). *Each panel shows a binned scatterplots using data from the full sample of IBES reports for a single firm. x-axis shows the firm's trailing twelve-month P/E,* $\text{TrailingPE}_{n,t} = \text{Price}_{n,t} \,/\, \text{EPS}_{n,t}$. *y-axis shows the P/E ratio implied by the analyst's price 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.*

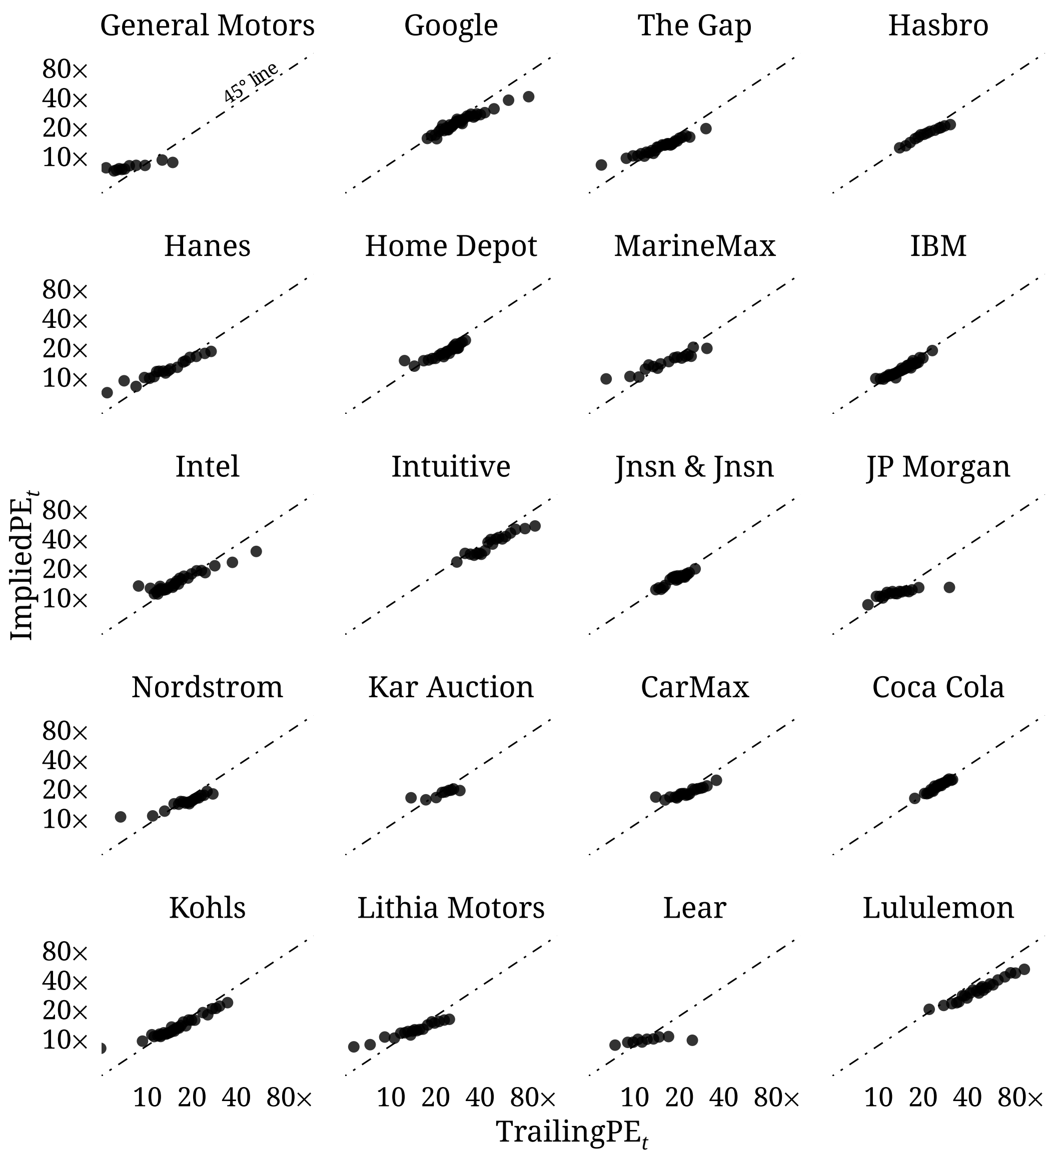

Figure B2(c). *Each panel shows a binned scatterplots using data from the full sample of IBES reports for a single firm. x-axis shows the firm's trailing twelve-month P/E,* $TrailingPE_{n,t} = Price_{n,t} / EPS_{n,t}$*. y-axis shows the P/E ratio implied by the analyst's price target and EPS forecast,*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

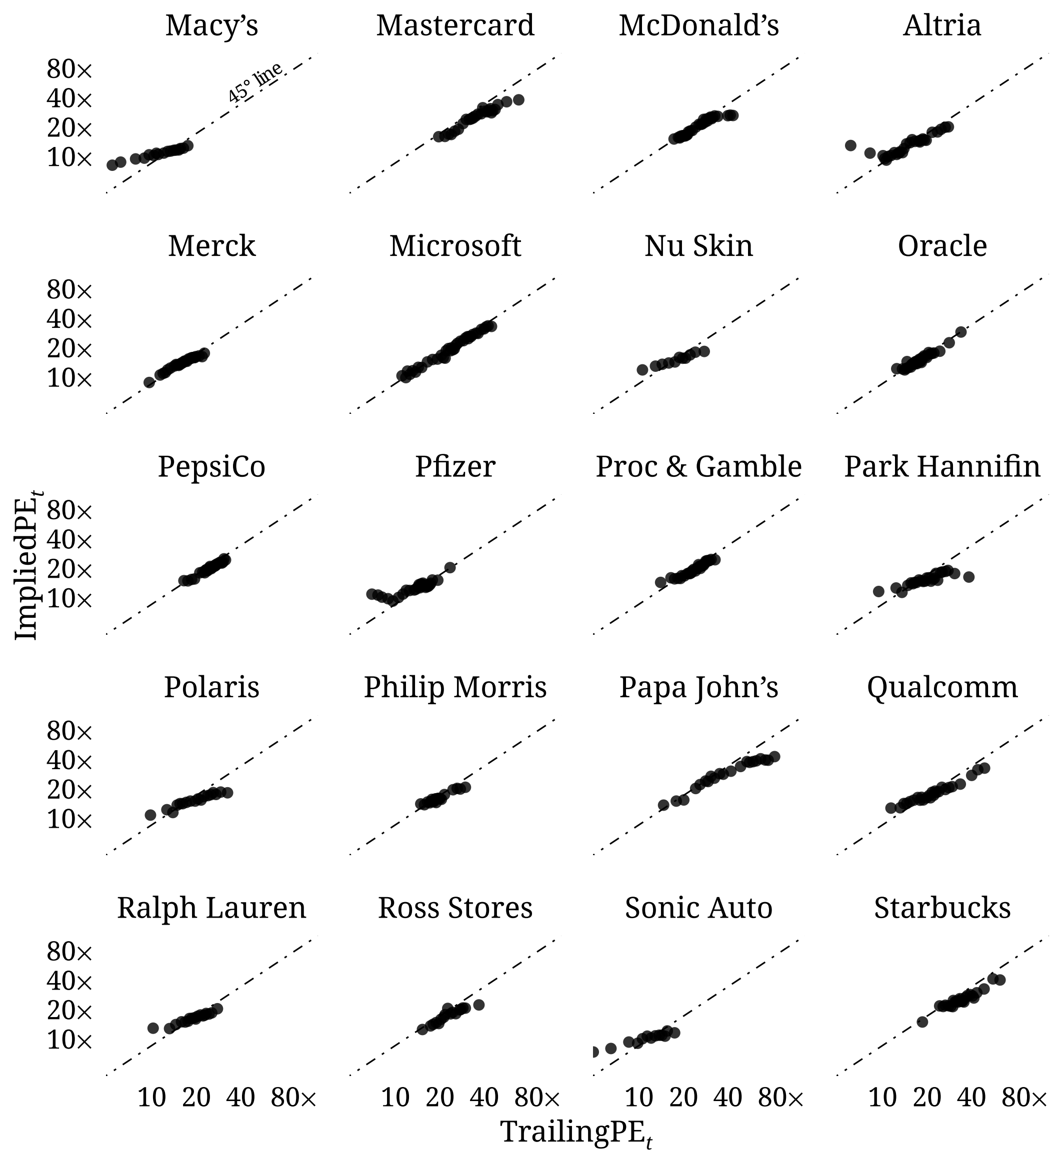

Figure B2(d). *Each panel shows a binned scatterplots using data from the full sample of IBES reports for a single firm. x-axis shows the firm's trailing twelve-month P/E,* $TrailingPE_{n,t} = Price_{n,t} / EPS_{n,t}$. *y-axis shows the P/E ratio implied by the analyst's price target and EPS forecast,*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

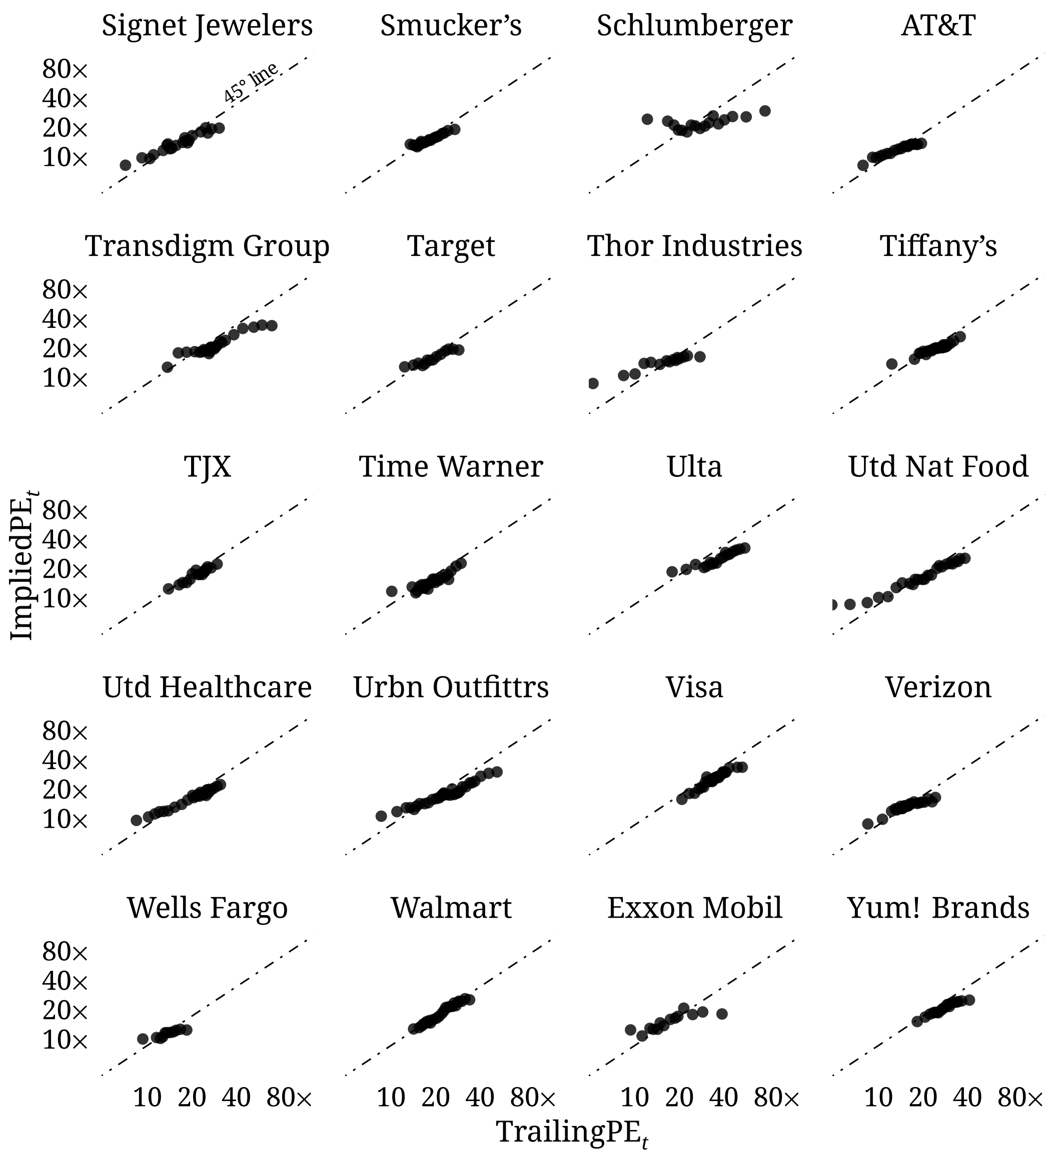

Figure B2(e). *Each panel shows a binned scatterplots using data from the full sample of IBES reports for a single firm. x-axis shows the firm's trailing twelve-month P/E, $TrailingPE_{n,t}$ = $Price_{n,t}$ / $EPS_{n,t}$. y-axis shows the P/E ratio implied by the analyst's price target and EPS forecast,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