Expected EPS × Trailing P/E*

Itzhak Ben-David† and Alex Chinco‡

July 3, 2024
[Click here for the latest version]

Abstract

When an analyst includes a price target in their earnings report, they are required to explain exactly how they calculated this one-year-ahead forecast. We read through these explanations to understand how analysts price their own subjective cash-flow expectations. Contrary to what textbooks assume, most do not apply present-value reasoning. Instead, they typically multiply a company's expected earnings per share (EPS) times its trailing price-to-earnings ratio (P/E) over the past year or two. We outline a simple model where this mostly backward-looking approach is correct on average because prices themselves are mostly backward-looking. Our empirical analysis shows that trailing P/E ratios explain 91% of the observed variation in analysts' price targets. Trailing P/E ratios also predict how realized prices respond to earnings news.

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

---

*We would like to thank Jeff Meli, Marco Sammon, and Jeff Wurgler for helpful comments.
†The Ohio State University and NBER. ben-david.1@osu.edu
‡Baruch College, City University of New York. alexchinco@gmail.com

Introduction

"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)" Researchers are willing to entertain the idea that some investors might have biased subjective beliefs or hold non-standard preferences. But no one ever questions whether real-world investors equate price with expected discounted payoff.

Suppose you buy a share of stock today. If you sell the share next year after collecting the dividend, your total future payout will be $\text{Price}_{t+1} + \text{Dividend}_{t+1}$. All standard asset-pricing models say that the current price will be given by

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

where $\mathbb{E}_t[\cdot]$ reflects investors' subjective beliefs about an asset's payoff next year and r_{t+1} is the discount rate they apply to these payoffs.

Given that textbook models relate a stock's current price to its expected future payouts, researchers would love to have data on the subjective beliefs of a broad swath of the investing public. Unfortunately, most market participants do not publicly announce their views prior to trading. Sell-side analysts are an exception to this rule. As such, the numerical values in their reports have played a critical role in asset-pricing research (Kothari, So, and Verdi, 2016).

To illustrate, Figure 1 shows a December 2019 report about Home Depot written by Chris Horvers, a senior analyst at JP Morgan. Chris Horvers started his report by recommending that investors "Overweight" Home Depot in their portfolios—i.e., by telling people to buy more shares of Home Depot. Then, he set a price target of $\mathbb{E}[\text{Price}] = \$241$ for December 2020—i.e., one year into the future. Chris Horvers also predicted that Home Depot's earnings per share (EPS) would be $\mathbb{E}[\text{EPS}] = \$11.50$ over the subsequent twelve months (FY2021).

Right now, researchers interpret these sorts of numerical values through the lens of forward-looking present-value relationships. If Chris Horvers had taken a present-value approach to setting his \$241 price target for Home Depot,

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.*

it would make sense to estimate the discount rate implied by this value given Home Depot's current price (Gebhardt, Lee, and Swaminathan, 2001). It would also be reasonable to aggregate sell-side analysts' earnings forecasts and plug them into a forward-looking Campbell and Shiller (1988) approximation. There is a world in which these exercises are exactly the right thing to do.

But we do not live in that world. In Section 1, we examine a sample of 513 sell-side analyst reports 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, $\text{PriceTarget}_t = \mathbb{E}_t[\text{Price}_{t+1}]$, equal to their near-term earnings forecast times a trailing price-to-earnings ratio (P/E).

How can we be so sure? Because sell-side analysts explicitly say so in their reports. These documents are more than just dry colorless lists of numbers. Analysts carefully explain how they convert their EPS forecasts for the next couple of years into a price target. They are legally required to do this. FINRA Rule 2241 states that "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 accompanied Chris Horvers' December 2020 price target for Home Depot. He could have said that he chose a value of $241 after carefully examining Home Depot's expected payouts from 2021 onward. He could have gone deep into the weeds, outlining precisely how

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.*

he discounted these expected future payoffs back to December 2020. He was more than capable of doing this sort of analysis. His report about Home Depot was 16 pages long and included numerous detailed calculations.

But Chris Horvers chose not to do this. Instead, his report 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"

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

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 use multiples analysis in a way that is consistent with present-value reasoning. The use of multiples on its own does not constitute a departure from textbook logic. For example, the classic Gordon model comes from iterating the pricing rule in Equation (1) forward, assuming constant discount and dividend-growth rates, and simplifying

$$\text{Price}_t = \mathbb{E}_t\left[\frac{\text{Price}_{t+1} + \text{Dividend}_{t+1}}{1 + r_{t+1}}\right] \tag{3a}$$

Iterate forward

$$= \mathbb{E}_t\left[\sum_{h=1}^{\infty} \frac{\text{Dividend}_{t+h}}{\prod_{i=1}^{h}(1 + r_{t+i})}\right] \tag{3b}$$

Assume: $(1+r)^h = \prod_{i=1}^{h}(1+r_{t+i})$

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

$$= \sum_{h=1}^{\infty} \frac{\text{Dividend}_t \cdot (1+g)^h}{(1+r)^h} \tag{3c}$$

…and simplify

$$= \mathbb{E}_t[\text{Dividend}_{t+1}] \times \left(\frac{1}{r-g}\right) \tag{3d}$$

Investors in this classic model price stocks with a forward-looking price-to-dividend (P/D) ratio. They scale up a company's dividend payout next year by a factor of $(\frac{1}{r-g}) = \left(\sum_{h=1}^{\infty} \frac{\mathbb{E}_t[\text{Dividend}_{t+h}]}{(1+r)^h}\right) \Big/ \mathbb{E}_t[\text{Dividend}_{t+1}]$ so that the price reflects the present discounted value of the company's entire future dividend stream.

The calculation in Equation (2) is not a special case of the Gordon model. Chris Horvers has clearly seen this framework. Why else would he have used Home Depot's expected 2021 earnings, $\mathbb{E}_t[\text{EPS}_{t+2}]$, to set his price target for December 2020, $\mathbb{E}_t[\text{Price}_{t+1}]$? But he still chose to use a *trailing* P/E. There is nothing forward-looking about the "Valuation Matrix" he provides in Figure 2. Chris Horvers arrived at his \$241 price target by asking: "Given how Home Depot has traded in the past, what would the company's price be if it were to announce earnings of \$11.50 per share?"

Frankly, it is surprising that so many papers have been written about biases in analyst earnings forecasts. At least analysts are trying to get 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 failed to apply the "one simple concept" at the root of "all asset-pricing theory". He did not even attempt to approximate $\sum_{h=1}^{\infty} \frac{\mathbb{E}_t[\text{Dividend}_{t+h}]}{(1+r)^h}$. Figure 1(a) even shows that Chris Horvers gave Home Depot an "overweight" rating, meaning that he chose to use a trailing P/E while simultaneously arguing that the company has been undervalued in the past.

We are not trying to pick on Chris Horvers or say that he is a bad analyst. The exact opposite is true. His 2019 report about Home Depot is emblematic of the kind of report that other analysts strive to write. Our point is that researchers need to take Chris Horvers seriously when he says he is not setting price equal to expected discounted payoff. Most analysts do not use present-value reasoning to price their own subjective cash-flow expectations. Researchers should write down models that reflect this reality.

But what exactly would such a model look like? The answer is not obvious. Behavioral finance typically studies models that deviate from a well-known rational benchmark in precisely one way (Rabin, 2013). Theory often becomes entirely intractable when more than one change is in play, and using a trailing P/E is much more than a one-step deviation from the textbook approach. If we take Equation (2) seriously, then there is nothing that pins down the price level. It could be that the resulting price dynamics are entirely incomprehensible.

In Section 2, we allay this concern by writing down a simple asset-pricing model where it makes sense to set PriceTarget = $\mathbb{E}[\text{EPS}] \times \text{TrailingPE}$. Investors in our model proportionally adjust their holdings over the next year based on the relative difference between analysts' price target and the current price. Likewise, realized price growth responds proportionally to these demand shocks. In this setup, observed prices will be mostly backward-looking. The only forward-looking information is the short-term earnings forecast, $\mathbb{E}[\text{EPS}]$. Because prices themselves are mostly backward-looking, we show that it is possible for price targets based on $\mathbb{E}[\text{EPS}] \times \text{TrailingPE}$ to be correct on average.

"In general, the claim that an instrument operates through a single known channel is called an *exclusion restriction*. (Angrist and Pischke, 2009)" Textbook models say that, if a piece of news affects a company's share price, it must do so by changing investors' beliefs about the expected discounted payoff to each shareholder. All price effects must operate through this one specific channel.

Our model points to a very different exclusion restriction. If a piece of news affects a company's expected return, then it must do so by changing investors' short-term earnings forecast, $\mathbb{E}[\text{EPS}]$. It does not matter if the signal also affects a company's long-term earnings growth g or the discount rate r investors should

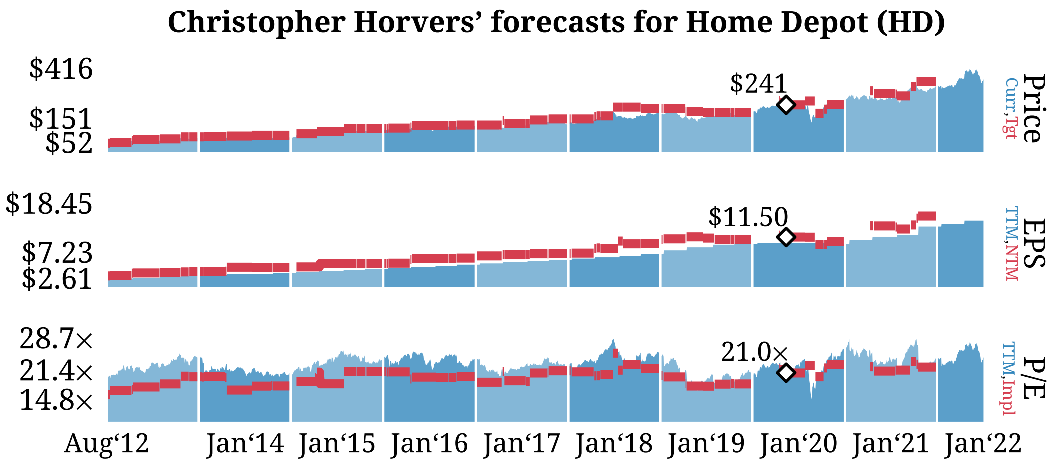

Figure 3. *y-axis shows min, median, and max. (Top) Blue ribbon is Home Depot's closing price on day t from CRSP, $Price_t$. Red line is Chris Horvers' price target, $PriceTarget_t = \mathbb{E}_t[Price_{\tau+1}]$, in IBES. (Middle) Blue is HD's trailing twelve-month (TTM) EPS on day t from IBES, EPS_t. Red is Chris Horvers' EPS forecast for the year following his target date, $\mathbb{E}_t[EPS_{\tau+2}]$. (Bottom) Blue is HD's TTM P/E ratio, $TrailingPE_t = Price_t \,/\, EPS_t$. Red is the P/E implied by Chris Horvers' forecasts, $ImpliedPE_t = PriceTarget_t \,/\, \mathbb{E}_t[EPS_{\tau+2}]$. White diamonds represent "primary source" values taken from Chris Horvers' December 12th 2019 report.*

apply in a Gordon model. In our model, analysts' price targets are based on trailing P/E ratios, which are already set in stone.

In Section 3, we show that trailing P/E ratios explain both price targets and realized returns. We start by studying analysts' price targets using data from IBES. Analysts set price targets for a fiscal-year end date roughly 12 months away, which we denote with $(\tau + 1)$. Analysts typically make their first forecast more than 12 months in advance and then revise this value as new information emerges. When the company's fiscal year-end is roughly 6 months away, analysts switch to targeting the following fiscal year-end (18 months away).

For each analyst tracking a given firm, we construct a panel of their most recent price target, $\text{PriceTarget}_t = \mathbb{E}_t[\text{Price}_{\tau+1}]$, on trading days from 18 months to 6 months prior to $(\tau+1)$. The red lines in the top panel of Figure 3 show Chris Horvers' price targets for Home Depot. The white diamond indicates the $241 price target Chris Horvers gave in his October 2019 report.

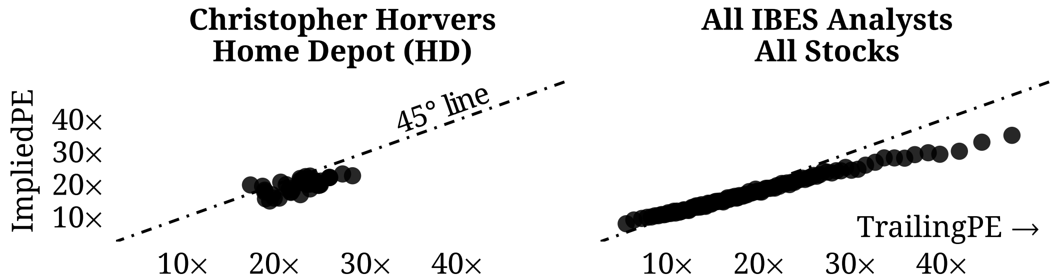

Figure 4. *(Left) Each dot denotes a day on which Chris Horvers updated his price target for Home Depot. x-axis is Home Depot's trailing twelve-month P/E,* $\mathit{TrailingPE}_t = \mathit{Price}_t \,/\, \mathit{EPS}_t$*. y-axis is the P/E ratio implied by Chris Horvers' forecast values,* $\mathit{ImpliedPE}_t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

The red lines in the middle panel of Figure 3 show Chris Horvers' earnings forecasts, and the white diamond indicates the \$11.50 EPS value from Chris Horvers' October 2019 report. Around half of all analysts in IBES post a two-year-ahead EPS forecast, $\mathbb{E}_t[\mathrm{EPS}_{\tau+2}]$, like the one Chris Horvers used. The remaining analysts only report $\mathbb{E}_t[\mathrm{EPS}_{\tau+1}]$ in IBES. We use the two-year-ahead EPS forecast whenever it is available in IBES; otherwise, we use the one-year-ahead value. We write this variable as $\mathbb{E}_t[\mathrm{EPS}]$ without a target-date subscript.

Chris Horvers explicitly told us that he was using a trailing P/E to price Home Depot's expected earnings, and the bottom panel of Figure 3 confirms that this was exactly what he did. The red line shows the P/E ratio implied by Chris Horvers' target price and subjective earnings expectation, $\mathrm{ImpliedPE}_t \overset{\text{def}}{=} \mathrm{PriceTarget}_t \,/\, \mathbb{E}_t[\mathrm{EPS}_{\tau+2}]$. These implied values closely track the Home Depot's trailing twelve-month P/E ratio, $\mathrm{TrailingPE}_t = \mathrm{Price}_t \,/\, \mathrm{EPS}_t$.

If Chris Horvers always set his price target for Home Depot based solely on the company's trailing twelve-month (TTM) P/E ratio, then each time he posted a new price target we would find that $\mathrm{ImpliedPE}_t = \mathrm{TrailingPE}_t$ exactly. The left panel of Figure 4 shows that this is a good first approximation to reality. Chris Horvers' implied P/E ratios move almost one-for-one with Home Depot's TTM P/E ratio. The right panel of Figure 4 shows that a tight linear relationship exists for other analysts covering other firms.

In the real world, analysts deploy variations on PriceTarget = $\mathbb{E}[\text{EPS}] \times$ TrailingPE. For example, the left panel of Figure 4 is a point cloud rather than a perfectly straight line because Chris Horvers used a trailing three-year average P/E rather than a trailing twelve-month value. Practical considerations also explain why the binned scatterplot on the right-hand side is slightly flatter than predicted by the theory. Analysts tend to consider outside factors when a company's trailing P/E is unusually high or low.

These wrinkles only serve to underscore the fact that analysts are pricing assets based on some sort of trailing P/E ratio. The fit in Figure 4 may not be perfect, but we know of no other asset-pricing theory that predicts investors' subjective price expectations nearly as well as PriceTarget = $\mathbb{E}[\text{EPS}] \times$ TrailingPE. Our simple model leaves out a lot of important details. Nevertheless, because it matches what analysts say they are doing, our model still explains over 90% of the observed variation in analysts' price targets.

$\mathbb{E}[\text{EPS}] \times$ TrailingPE also explains realized future returns. Our empirical tests focus on the way that prices respond to earnings surprises. If we hold constant the size of the earnings surprise, there is only one variable left in our simple model: TrailingPE. So the price of a TrailingPE = 20× stock should react twice as much to the same earnings surprise as the price of a 10× stock.

To test this prediction, we group all the firm-quarter observations that realized the same-sized earnings surprise into bins. Within each bin, we then regress each firm's realized price change following the earnings surprise on its trailing P/E ratio. Finally, we check if the estimated slope coefficients from these separate regressions increase linearly with the size of the earnings surprise.

Large earnings surprises often signal persist changes. If investors apply present-value logic, these persistent future effects should cause them to choose a different multiple. In a textbook world, we would not find a neat linear relationship in the data. A stock with a TrailingPE = 20× might have double the price reaction of a 10× stock for small earnings surprises. But, for large surprises, the effect of persistent future changes should take over.

This is not what we find. Our empirical results show a straight line exactly as predicted by our simple $\mathbb{E}[\text{EPS}] \times$ TrailingPE model.

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, Nissim, and Thomas, 2002; 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; De la O and Myers, 2021). While we agree that it is important to understand the implications of biased analyst forecasts, it is also important to correctly model how these biased forecasts 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

In May 2002, the SEC passed NASD Rule 2711 stating that: "If a research report contains a price target, the [analyst] must disclose in the research report the valuation methods used to determine the price target." In 2015, this rule was superseded by FINRA Rule 2241, which also requires a "price target [to be] accompanied by a clear explanation." In this section, we study the clear explanations that analysts provide to see how they price their own subjective earnings expectations. Subsection 1.1 describes the kind of explanations textbook models would predict. Subsection 1.2 then documents how analysts actually explain themselves. Finally, in Subsection 1.3, we provide corroborating evidence.

1.1 What Textbook Models Predict

Imagine that you are an alien on your way to visit the Earth. You have an interest in how humans organize their societies and want to better understand this thing they call a "financial market". Not knowing where else to start, you begin by reading through the academic literature. When you arrive on our planet, you talk to a human named "Chris" who calls himself a "sell-side analyst". Based on what you have read, how would you expect Chris to describe his approach to making one-year-ahead price forecasts?

Prediction #1. *Analysts should use a forward-looking pricing rule.*

To start with, textbook models predict that analysts will talk about forward-looking pricing rules. In all standard asset-pricing models, a company's current share price will reflect its expected discounted payoff to shareholders next year as shown in Equation (1) and reproduced below

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

Dividend_{t+1} is the company's dividend next year, Price_{t+1} is the sale price at the end of next year, and r_{t+1} is the discount rate applied to these future payoffs.

As discussed in the introduction, the assumption that prices are forward-looking is tantamount to an exclusion restriction (Manski, 1995). Textbook models say that, if a piece of information affects a company's share price, then it must do so by changing investors' beliefs about the discounted value of the firm's payout to shareholders next year. According to these models, analysts should describe all price effects as operating through this one channel.

Prediction #2. *Analysts should apply the same pricing rule at all points in time.*

Notice that there are two prices in Equation (1): the company's current share price, Price_t, and its price level next year, Price_{t+1}. Textbook models predict that an analyst will describe both these prices as following from the same underlying logic. For example, researchers think it is completely natural to replace the Price_{t+1} on the right-hand side of Equation (1) with $\mathbb{E}_{t+1}\big[\frac{\text{Dividend}_{t+2}+\text{Price}_{t+2}}{1+r_{t+2}}\big]$ and then swap out the Price_{t+2} in the resulting expression with $\mathbb{E}_{t+2}\big[\frac{\text{Dividend}_{t+3}+\text{Price}_{t+3}}{1+r_{t+3}}\big]$ and so on... This process of forward iteration yields an expression for a company's current share price in terms of its expected future dividend stream

$$\text{Price}_t = \mathbb{E}_t\left[\sum_{h=1}^{\infty} \frac{\text{Dividend}_{t+h}}{\prod_{i=1}^{h}(1+r_{t+i})}\right] \tag{4}$$

The dividend discount model comes from applying this procedure in a world with a constant discount rate, $(1+r)^h = \prod_{i=1}^{h}(1+r_{t+i})$. The whole derivation assumes that Equation (1) holds in all time periods.

Prediction #3. *Analysts should use multiples analysis to scale up a company's expected payout next year to capture the discounted value of all future payouts.*

There is no inherent conflict between multiples analysis and textbook models. As highlighted in the introduction, in a world with constant annual dividend growth, $(1+g)^h = \frac{\mathbb{E}_t[\text{Dividend}_{t+h}]}{\text{Dividend}_t}$, the dividend discount model predicts that the present discounted value of a company's entire future dividend stream will be proportional to its expected dividend next year, $\frac{1}{r-g} = \Big(\sum_{h=1}^{\infty}\frac{\mathbb{E}_t[\text{Dividend}_{t+h}]}{(1+r)^h}\Big)\Big/\,\mathbb{E}_t[\text{Dividend}_{t+1}]$. For these reasons, an analyst living in a textbook world might well describe using a forward-looking multiple.

In principle, the analyst could even describe using an extremely complicated forward-looking multiple. For example, Campbell and Shiller (1988) replaces $\left(\frac{1}{r-g}\right)$ with an approximate multiple which allows r and g to vary over time

$$\text{Price}_t \approx \mathbb{E}_t[\text{Dividend}_{t+1}] \times \left(\frac{1}{e^{\sum_{h=0}^{\infty} \rho^h \cdot \mathbb{E}_t[r_{(t+h)+1}] - \sum_{h=0}^{\infty} \rho^h \cdot \mathbb{E}_t[\Delta \log \text{Dividend}_{(t+h)+1}]}} \right) \tag{5}$$

where $\rho = \frac{1}{1+e^{\mathbb{E}[\log \text{DivYield}]}}$. This is just another way of scaling up a company's expected next-twelve-month dividend to reflect its entire future dividend stream.

However, to be consistent with textbook present-value logic, an analyst must use a forward-looking multiple. If an analyst were to explain that his price target was based on a trailing price-to-dividend ratio (P/D), then his approach would not be consistent with textbook present-value reasoning. No amount of algebraic complication could change this.

Prediction #4. *Analysts' should report subjective beliefs that respect ex ante accounting identities.*

Textbook models do not require analysts to have objectively correct expectations about a company's future payoffs. Analysts could describe using biased subjective beliefs to set price targets in a way that is consistent with textbook models. For example, Campbell (2017) explains that Equation (5) "holds ex post as an accounting identity. It should therefore hold ex ante, not only for rational expectations but also for irrational expectations that respect identities."

But those last three words are important. Yes, a company's current share price must satisfy the ex post identity

$$\text{Price}_t = (1 + \text{Return}_t) \cdot \text{Price}_{t-1} - \text{Dividend}_t \tag{6}$$

But analysts do not have to set price targets with this identity in mind

$$\text{PriceTarget}_t \neq (1 + \mathbb{E}_t[\text{Return}_{t+1}]) \cdot \text{Price}_t - \mathbb{E}_t[\text{Dividend}_{t+1}] \tag{7}$$

Price targets do not have to respect identities. Analysts do not have to create price forecasts by plugging their subjective beliefs into Equation (5).

Ex ante calculations take place in the analyst's head. Practical considerations do not have to play any role. There are hundreds of thousands of kids who expect to play in the NBA one day even though there are only 450 NBA roster spots at any one point in time. No one finds this the least bit troubling. Like a kid with unrealistic dreams, an analyst can expect a PriceTarget_t that is inconsistent with the reality implied by $\mathbb{E}_t[\text{Return}_{t+1}]$, $\mathbb{E}_t[\text{Dividend}_{t+1}]$, and Price_t.

Prediction #5. *Analysts should view earnings as the capacity to pay dividends.*

Textbook models assume that "price equals expected discounted payoff". Hence, according to these models, analysts should only care about a company's earnings insofar as they translate into future dividend payouts to shareholders. When a company's earnings go up, analysts should raise their price target to reflect an increase in the company's capacity to pay dividends going forward. Judging by textbook models, the mapping from expected earnings to expected dividend payouts should be a key part of analysts' write up.

Prediction #6. *Analysts should use a discount rate equal to the expected return.*

Finally, most standard asset-pricing models are actually models of expected returns. Nevertheless, researchers still refer to them as "asset-pricing models" because they assume that investors use the implied expected return as the discount rate when computing a company's share price. Textbook models predict that analysts will describe the discount rates they use in these terms.

For example, think about the Capital Asset-Pricing Model (CAPM; Treynor, 1961; Sharpe, 1964; Lintner, 1965). This model says that the difference between a company's expected return next year and the riskfree rate should be proportional to the covariance between its future returns and the market

$$\mathbb{E}_t[\text{Return}_{t+1}] - r_f = \lambda \times \left(\frac{\text{Cov}_t[\text{Return}_{t+1}, \text{Market}_{t+1}]}{\text{Var}_t[\text{Return}_{t+1}]} \right) \tag{8}$$

Researchers call it the CAPM and not the CERM because they assume that investors determine a company's share price by replacing the discount rate, r, with $\mathbb{E}_t[\text{Return}_{t+1}] = r_f + \lambda \cdot \left(\frac{\text{Cov}_t[\text{Return}_{t+1}, \text{Market}_{t+1}]}{\text{Var}_t[\text{Return}_{t+1}]} \right)$ in Equation (1).

1.2 What Analysts Actually Say

We now look at how sell-side analysts say that they create their price targets in earnings reports. None of these six textbook predictions is true in our data. Real-world analysts do not price their subjective earnings expectations in a way that is consistent with the forward-looking present-value logic at the heart of standard asset-pricing models.

Data description. To generate the results in this section, we read through the text of 513 sell-side analyst reports. After reading each report, we asked: "How did the author of this report convert their earnings forecast, $\mathbb{E}_t[\text{EPS}]$, into a price target, PriceTarget_t?" Every report in our sample has an entire section where the analyst directly answers this question. The passage is clearly labeled. We have already seen an example of this in Figure 2, and we will be providing many more examples throughout this section. Analysts explicitly state how they price their own subjective cash-flow expectations.

We downloaded our 513 analyst reports from Investext in two separate waves. The first wave contains reports written about the 30 largest publicly traded companies at year-end in 2004, 2011, and 2019. We list these 47 companies in Table 1. For each company in a given year, we include one report written by each brokerage in Table 2. This gives us a total of 339 analyst reports in our first sample: 91 in 2004, 93 in 2011, and 155 in 2019. Our goal in this first wave of downloads was to construct a representative snapshot of typical sell-side reports written about large publicly traded companies.

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. But you might still hold out hope that even if the typical analyst does not apply forward-looking present-value logic when writing the average report, surely the best analysts do this when writing reports that really matter.

To check whether this is the case, we also downloaded an additional 174 coverage-initiation reports written by 28 analysts on *Institutional Investor* magazine's All-American research team. These 174 reports come from the best

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.*

of the best (Stickel, 1992). Analysts on *Institutional Investor* magazine's All-American team are at the top of their field. "All-Star analysts earn 61% higher compensation than their unrated peers. (Groysberg, Healy, and Maber, 2011)"

Institutional Investor publishes their rankings in October. We read through these issues and cataloged which analysts made the All-American research team each year. The magazine ranks analysts by GICS sector. So, for each sector, we focused on the 10 analysts with the most years on the All-American team. The 174 documents in our second wave of downloads were written by the subset of these All-American analysts who could be matched to both Investext and IBES.

A "coverage-initiation report" is the first report an analyst writes about a particular company for a given brokerage. This means that, for the reports in our second wave of downloads, either the firm is new to public equity markets or the analyst is new to his/her current brokerage. Analysts put a disproportionate amount of effort into writing coverage-initiation reports (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.

If anything, analysts should be more likely to use forward-looking information in coverage-initiation reports simply because there is often not much trailing information to go on. 53 of our 174 coverage-initiation reports (30.5%

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.*

of sample) in our sample involve companies that went public within the previous three years. These firms have little historical data. Moreover, many firms start out with negative earnings. Both these considerations make it difficult for analysts to apply the formula $\mathbb{E}[\text{EPS}] \times \text{TrailingPE}$.

We only look at reports written by analysts that can be matched to IBES. 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 the entire time-series of reports from Investext for a subset of analyst-firm pairs. For example, Figure 5 shows the first page and methods section from an October 2019 report written by Andrea Teixeira about Coca-Cola (KO). The red lines in the top two panels of

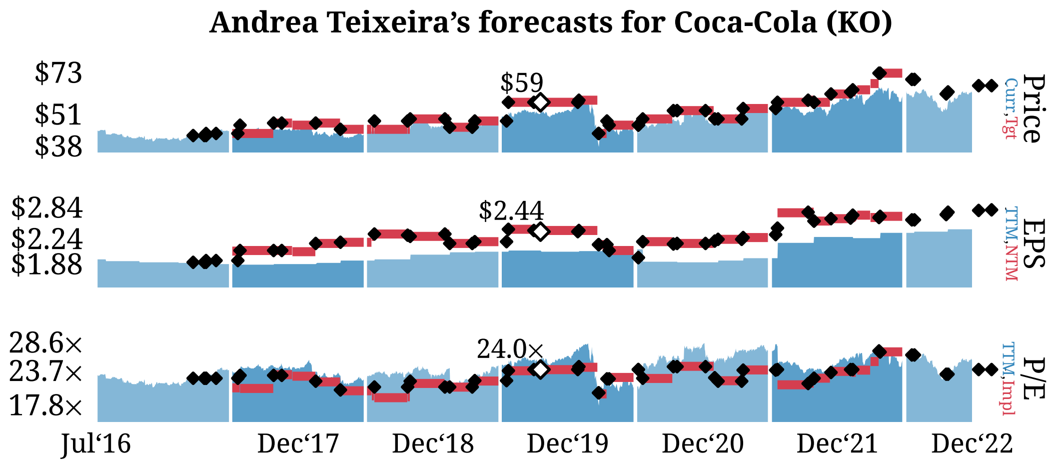

Figure 6. *y-axis shows min, median, and max. (Top) Blue ribbon is Coca-Cola's (KO)'s closing price on day t from CRSP, Price_t. Red line is Andrea Teixeira's price target, $\text{PriceTarget}_t = \mathbb{E}_t[\text{Price}_{\tau+1}]$, for target date $(\tau + 1)$ as reported in IBES. (Middle) Blue is KO's trailing twelve-month (TTM) earnings per share (EPS) on day t, EPS_t, as reported in IBES. Red is Andrea Teixeira's EPS forecast for the year following her target date, $\mathbb{E}_t[\text{EPS}_{\tau+2}]$. (Bottom) Blue is KO's TTM price-to-earnings (P/E) ratio, $\text{TrailingPE}_t = \text{Price}_t \,/\, \text{EPS}_t$. Red is the P/E implied by Andrea Teixeira's forecasts, $\text{ImpliedPE}_t = \text{PriceTarget}_t \,/\, \mathbb{E}_t[\text{EPS}_{\tau+2}]$. White diamonds represent "primary source" values taken from Andrea Teixeira's October 18th 2019 report about Coca-Cola in our document database. Black diamonds depict values from other reports written by Andrea Teixeira about Coca-Cola not in our sample.*

Figure 6 show all price targets and EPS forecasts that Andrea Teixeira made for Coke as reported in IBES. The white diamonds in these two panels depict the $59 price target and $2.44 EPS forecast found in her October 2019 report. The black diamonds represent analogous PriceTarget_t and $\mathbb{E}_t[\text{EPS}_{\tau+2}]$ values found in other Andrea Teixeira reports about Coke downloaded from Investext.

As you can see from Figure 6, the values found in these additional PDFs perfectly match up with the numbers in IBES. We also note that the implied P/E values that we calculated based on IBES data (red lines; bottom panel) match up with the P/E ratios that Andrea Teixeira says she used in her 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.

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%.

Price targets reflect trailing P/Es. The first thing that jumps out at you when reading through the reports is that multiples analysis is incredibly common. 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 $\mathbb{E}[\text{EPS}] \times \text{TrailingPE}$. It is common to see analysts use them in situations where a company's EPS has been negative in recent years or is so small that it would imply an enormous P/E ratio. In this section, we treat these calculations as separate valuation techniques to be as conservative as possible.

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%.

Again, there is nothing inherently wrong with multiples analysis. The issue is the way that analysts calculate their multiple. Sell-side analysts typically chose their multiple based on where the firm and others like it have been trading at in recent years. This general approach directly contradicts prediction #1.

Table 5 shows that analysts set a price target by looking at the firm's own trailing multiples in 63.5% of our sample (326 of 513 reports). They chose a multiple to reflect 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%). These firms often lack the requisite historical data that an analyst needs to compute a trailing average. However, as you can see in Figure 7, many coverage-initiation reports still use the formula $\text{PriceTarget} = \mathbb{E}[\text{EPS}] \times \text{TrailingPE}$ when setting a price target. This figure shows the first page and methods section from a RBC Capital Markets (2015) report about Chico's FAS written by Brian Tunick. There is nothing forward-looking about his choice of a 20× P/E ratio.

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

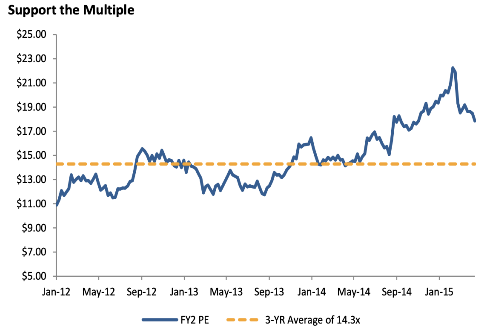

Source: Company reports and RBC Capital Markets estimates

(b) Methods section

Figure 7. *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.*

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%*.*

Also note how, at the top of the first page, Brian Tunick states that the CHS's "mid-to high-teens total returns" will come from "15% EPS CAGR from 2015–2017E and a ~2% dividend yield." If CHS's dividend yield had been 0% rather than 2% in the past, then Brian Tunick's expected returns would have been entirely determined by his forecast of short-term EPS growth.

There is every reason to believe that analysts are capable of calculating a forward-looking multiple. Table 6 shows that analysts performed "sum of the parts (SOTP)" analysis in 38.8% of our sample (199 out of 513 reports). This procedure is far more involved than simply calculating a model-implied value of $\left(\frac{1}{r-g}\right)$. For an example of what SOTP analysis looks like, see Figure 8 which 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, Figure 9 shows an April 2010 earnings report about Avis Budget in which the analyst uses multiple different kinds of earnings multiples. This piece of research was a coverage-initiation report written by an All-American analyst, Himanshu Patel. While this is a thorough report by a high-quality analyst, it does not use present-value reasoning to set a price target. Himanshu Patel uses a combination of backward-looking multiples.

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 8. *Earning report about Amazon, which was published on October 24th 2019 by Wolfe Research. The lead analyst on this report was Chris Bottiglieri. He computed a different multiple to value each of Amazon's four lines of business.*

J.P.Morgan

North America Equity Research
12 April 2010

Avis Budget and Dollar Thrifty

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

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

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

CAR, CAR US
Overweight
$14.91

Price Target: $20.00

DTG, DTG US
Neutral
$36.21

Price Target: $39.00

(a) Top of first page

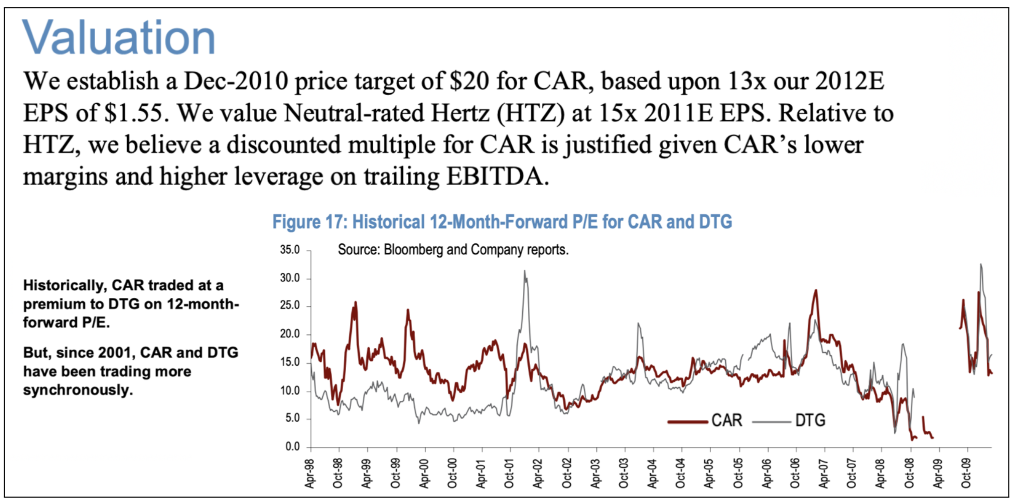

Valuation

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

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

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

(b) Methods section

Figure 9. *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.*

There is little evidence of present-value reasoning. In Table 7, we see that analysts mention a discounted cash-flow (DCF) or dividend discount model in 30.2% of reports (155 of 513). However, this table also reveals that analysts rarely used 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. 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).

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

What's more, this 30.2% 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 90% 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. For example, Figure 10 shows the entirety of the methods section from one of the 19 discount-model-only reports from Credit Suisse.

Prior to examining the data, we anticipated that analysts would be more likely to use a DCF model in our sample of 174 coverage-initiation reports. After all, many of these reports are for newly public firms with little historical data. However, it turns out that discount models are even less common in this subset of our data. The "All Am" column in Table 7 shows that only one in five coverage-initiation reports that an All-American analyst writes (34 of 174; 19.5%) will make use of a discount model in any capacity.

The All-American analysts who are responsible for these reports often talk about DCF models as a second-best option. For example, Figure 11 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

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%.

"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. "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 use a discount model, they often implemented the model in a way that is inconsistent with textbook present-value reasoning. For instance, Table 7 indicates that in 24.6% of our sample (126 reports), the analyst's price target is based on both a discount model and a trailing multiple. In these reports, the analyst typically just averages together the price targets implied by each method. This evidence contradicts prediction #3.

For example, Figure 12 shows a December 2019 report about Citigroup in which the lead analyst, Mike Mayo, describes his price target as a "simple average of six valuation techniques (PE, price-to-book, dividend discount model,

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 11. *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.*

PE/G ratio analysis and sum of the parts for both PE and PB)." Mike Mayo thought about $\left(\frac{1}{r-g}\right)$ as just another trailing multiple (Mukhlynina and Nyborg, 2020).

Analysts do not take a dynamically consistent approach. Textbooks predict that analysts will set price targets by asking: "How is the market going to price a company's expected earnings next year?" Instead, we find that analysts set price targets using a trailing P/E. This is like asking: "How would the company's expected earnings be priced under current market conditions?" This approach is not dynamically consistent (prediction #2).

The price targets that analysts report also do not respect ex ante accounting identities (prediction #4). For example, notice the fundamental tension between Andrea Teixeira's trading recommendation in Figure 13 and her use of a backward-looking multiple. She gave Pepsi an "Overweight" rating in her October 2019 report, 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.

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

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 13(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 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.

J.P.Morgan

PepsiCo

Resilient Growth Continues to Drive PEP Higher; Reiterate OW

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

North America Equity Research
03 October 2019

Overweight

PEP, PEP US

Price (03 Oct 19): $138.23

▲ **Price Target (Dec-20): $154.00**
Prior (Dec-20): $148.00

(a) Top of first page

Valuation

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

(b) Methods section

Key Metrics (FYE Dec)

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

(c) Table of key metrics

Figure 13. *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.*

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.

Analysts care about earnings for earnings' sake. Prediction #5 states that, according to textbook models, sell-side analysts should care about a company's earnings because these cash flows allow a firm to pay dividends to each shareholder. 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 12 is one of the few reports that specifically talks about using a dividend discount model.

On top of that, when an analyst does mention a company's dividend yield, this information typically only plays a role in computing the firm's expected return over the next year. Analysts do not use it when setting their price target. They "track capital gains and dividends as separate and largely independent variables. (Hartzmark and Solomon, 2019)"

Figure 14 provides an illustrative example. This figure 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

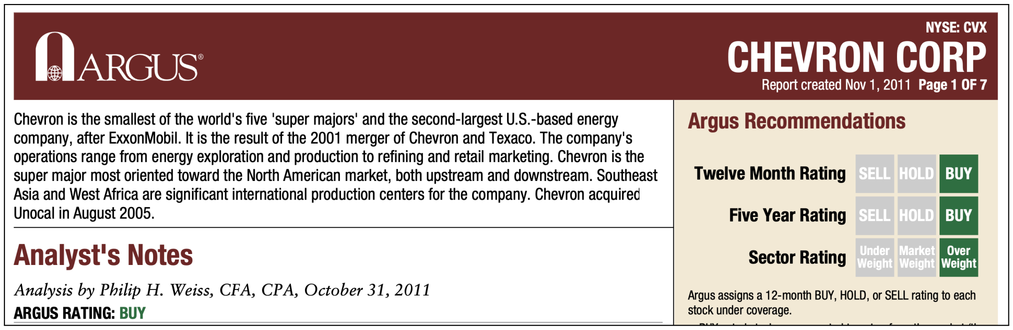

NYSE: CVX

ARGUS

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

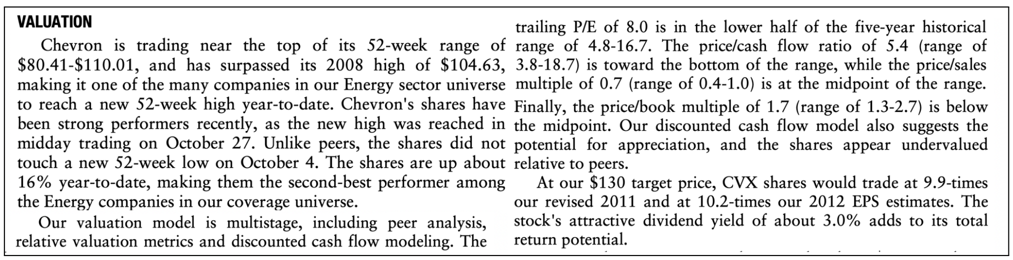

VALUATION

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

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, $\frac{\$3.12}{\$105.05} \approx 3\%$,

$$\underbrace{\mathbb{E}_t[\text{Return}_{t+1}]}_{27\%} = \underbrace{\left(\frac{\text{PriceTarget}_t - \text{Price}_t}{\text{Price}_t}\right)}_{(\$130.00-\$105.05)/\$105.05 \approx 24\%} + \underbrace{\left(\frac{\text{Dividend}_t}{\text{Price}_t}\right)}_{\$3.12/\$105.05 \approx 3\%} \tag{9}$$

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 (9). 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, $\text{Dividend}_t \neq \mathbb{E}_t[\text{Dividend}_{t+1}]$.

It might at first seem like Mr Weiss was following textbook logic, but the nature of these two differences shows that he was not. "Asset-pricing theory all stems from one simple concept: price equals expected discounted payoff. (Cochrane, 2009)" But Mr Weiss chose to calculate an expected return based on Chevron's past pricing.

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.

Analysts focus on expected earnings not discount rates. 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 (prediction #6).

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 14(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 prices. It is clear from reading analysts' reports that they are not.

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.

Every profession does some things on autopilot. 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) that researchers obsess over. 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. They are not even trying to calculate the present-value formula at the heart of every standard asset-pricing model. This seems like the bigger issue.

1.3 General Discussion

Our main results come from reading a sample of earnings reports to see how analysts describe their own pricing rule. We conclude this section by answering common questions that researchers often have about this empirical strategy.

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"?

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. Again, think back to Figure 10, which shows the entirety of the methods section from one of the 28 DCF-only reports in our sample. "Discounted Cash Flow (DCF) Valuation. (Credit Suisse, 2004)"

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 low value of 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.

Is our data sample representative? Our analysis in this section is based on reading a sample of 513 earnings reports. We read through each of these reports ourselves. This takes time and puts a practical limit on the size of our data set. Given the number of observations, you might be concerned about the representativeness of our sample. In particular, Table 7 indicates that analysts used a DCF or dividend discount model in less than 40% of our sample (155 of

513 reports). If we examined a more complete data set, would we still find that most analysts do not apply present-value logic?

Yes. 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 non-representative sample. Most reports do not even mention a discount model.

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.

Do other market participants think like analysts? We find that most analysts do not set price equal to expected discounted payoff. However, in principle, other kinds of investors could still be following the textbook approach.

We respond to this concern in two ways. First, suppose for the sake of argument that sell-side analysts were the only ones who did not apply present-value reasoning. In that case, we would not expect a company's trailing P/E to have a sizable effect on its price. Nevertheless, our results would still have important implications for researchers. Sell-side analysts do not determine a company's share price by plugging their subjective cash-flow expectations into some version of the Campbell and Shiller (1988) approximation in Equation (5). We learn nothing about how investors price assets by plugging analysts' subjective cash-flow expectations into this formula.

Second, there are good reasons to believe that sell-side analysts are not the only ones using trailing P/Es. 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 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.

In addition, we also document that, when regulatory filings include a price calculation, market participants are much more likely to use multiples analysis than a discount model to create this number. This is another more direct piece of evidence suggesting that other market participants set prices in a similar way to sell-side analysts.

We look at seven different kinds of regulatory filings submitted to the SEC 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; 40-Act funds use this form to report holdings, performance, assets under management, etc on a quarterly basis.

After downloading these forms, we restrict our sample to filings that include a discussion of firm pricing. 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.

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.*

When we look at the remaining 70k observations, 34.0% mention a discount model of some sort while 90.6% talk about multiples analysis.

Some niche industries do rely on DCF modeling. Sell-side analysts do regularly use present-value logic to set price targets for a small subset of industries. Analysts use DCF models to value shipping companies, which are frequently 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). These sorts of firms are not the ones that most researchers have in mind when they write down their asset-pricing models.

What's more, when we look at sell-side reports in these niche industries, it is obvious that DCF models play a much more prominent role in analysts' thinking. For example, Figure 15 shows a coverage-initiation report written by Michael Webber about GasLog Ltd in January 2014. This report is completely different from sell-side research about a company like Apple.

January 13, 2014 **Outperform / V**

Equity Research

Sector: Marine GP
Overweight

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.**

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

(a) Top of first page

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

(b) Methods section

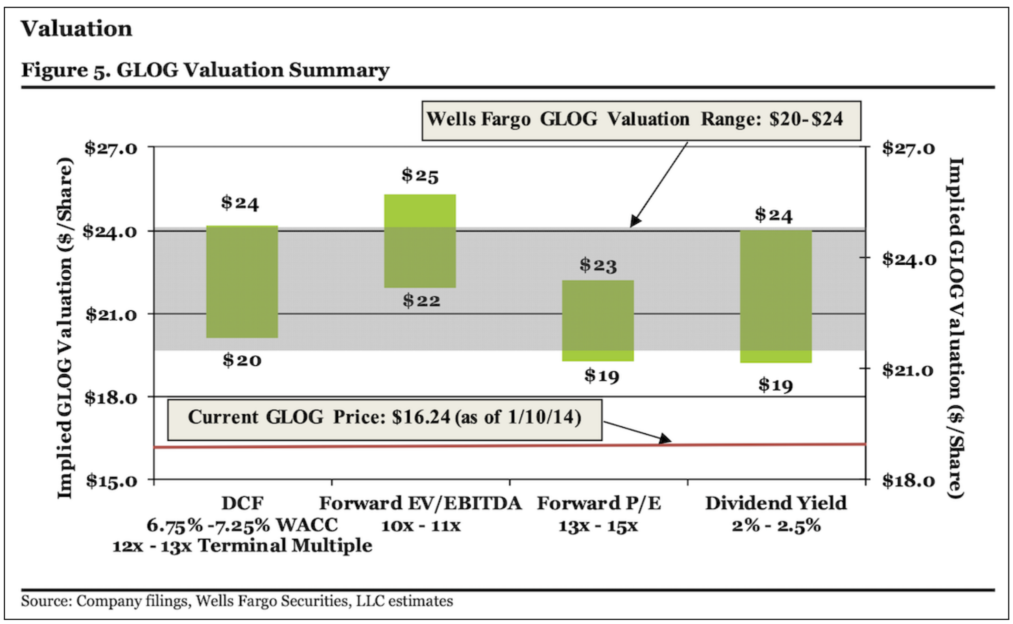

(c) Valuation summary

Figure 15. *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.*

Michael Webber's report clearly states that he used a DCF model to set his price target for GasLog Ltd. He gives the precise numerical inputs needed to do the calculation. Our method of analysis is capable of finding evidence of present-value reasoning when such evidence exists. This is just not what most analysts are doing when they write sell-side reports about a prototypical large publicly traded company.

2 A Simple Model

Right now, the first step when writing down any asset-pricing model is to make an assumption about future payoff distributions. The second step is to specify investors' preferences over these future payoffs. After hammering in these guideposts, a researcher then asks: Which discount rate will investors use to price each asset according to the model? Unfortunately, we have just seen that real-world analysts do not set price equal to expected discounted payoff.

In this section, we show that it is possible to write down a simple model that is consistent with how analysts describe their own approach. Analysts focus on earnings not payoffs. So, 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. Because market participants do not apply present-value reasoning, the company's share price will not equal expected discounted payoff in our model. Nevertheless, we are still able to provide conditions under which analysts' price targets are correct on average. This agreement occurs because prices are almost entirely backward-looking. Finally, in Subsection 2.4, we highlight how our model yields a sharp testable exclusion restriction. For a piece of news to affect a company's expected return, the news must change short-term earnings expectations.

2.1 Earnings Growth

In textbook models, investors are supposed to care about a firm's future payouts to each shareholder. They care about earnings only insofar as these earnings translate into future payouts to each shareholder. However, as noted in the previous section, many sell-side analysts appear to focus on earnings for earnings' sake. So we start our model by describing how earnings evolve.

We model the pricing of a single company. Let EPS_t denote this company's earnings per share over the previous twelve months. We assume that these earnings change over time according to the law of motion below

$$\left(\frac{\mathrm{EPS}_{t+1} - \mathrm{EPS}_t}{\mathrm{EPS}_t} \right) = X_t + \epsilon_{t+1} \tag{10}$$

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

Think about a firm that had earnings of $\mathrm{EPS}_t = \$1.00$/sh over the past year. Over the next twelve months, its earnings are expected to grow by $X_t = 5\%$ on average. But investors would not be surprised to see the firm's earnings growth rate by $\sigma = 2\%$pt higher or lower. Given these assumptions, investors expect the company to generate earnings of $\mathbb{E}_t[\mathrm{EPS}_{t+1}] = \$1.05 \pm \$0.02$ for each shareholder over the next year and $\mathbb{E}_t[\mathrm{EPS}_{t+2}] = \$1.11 \pm \$0.04$ 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. If we were writing that sort of model, we would also need to give a law of motion for how X_t will change.

But we are not writing down that sort of model. In our model, all that matters is analysts' near-term earnings forecast, $\mathbb{E}_t[\mathrm{EPS}_{t+2}]$. We will give conditions under which analysts' price target will match the firm's realized price next year. However, there will be nothing tying this realized price level to the present discounted value of the firm's expected future payoffs.

2.2 Investor Demand

Textbook models assume that investors choose their portfolio allocation to maximize the expected utility from its future payoff stream. However, in the previous section, we saw that sell-side analysts made explicit trading recommendations based on their price target. 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)"

This is how investors trade in our model. Specifically, at each time t, investors tell their broker how many shares of the company's stock they want to hold next year, Demand_{t+1}. They choose this quantity by comparing the company's current share price, Price_t, to analysts' one-year-ahead price target

$$\text{PriceTarget}_t = \mathbb{E}_t[\text{EPS}_{t+2}] \times \text{TrailingPE}_t \tag{11}$$

where $\text{TrailingPE}_t = \text{Price}_t \,/\, \text{EPS}_t$ is the firm's trailing-twelve-month P/E ratio.

When analysts' price target is higher than the current price, investors tell their broker to buy shares over the next year. When the price target is lower, they reduce their position. Moreover, they do so in proportion to the percent difference between analysts' price target and the current share price

$$\left(\frac{\text{Demand}_{t+1} - \text{Demand}_t}{\text{Demand}_t}\right) = \mu \cdot \left(\frac{\text{PriceTarget}_t - \text{Price}_t}{\text{Price}_t}\right) \tag{12}$$

where the demand multiplier, $\mu > 0$, is a positive constant.

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

Proposition 2.2 below shows that, even though investors in our model make their portfolio decisions by comparing analysts' price target to the company's current price, it is "as if" they were making the decision based solely on the company's expected near-term earnings growth.

Proposition 2.2 ("As If" Demand Schedule)**.** *If analysts use trailing P/E ratios to set price targets (Equation 11) and investors proportionally adjust their holdings next year (Equation 12), it is "as if" investors adjust their demand based on changes in the firm's expected near-term earnings growth*

$$\left(\frac{Demand_{t+1} - Demand_t}{Demand_t}\right) = \mu \cdot \left(\frac{\mathbb{E}_t[EPS_{t+2}] - EPS_t}{EPS_t}\right) \tag{13}$$

Analysts need not be able to describe every detail of an asset-pricing model. But analysts do need to be able to describe key details of a model if the model requires them to know those details. It is not possible to tell an "as if" story for every part of a model. There are some things that an analyst must be aware of. If analysts calculate price targets using a trailing P/E ratio, their approach cannot be seen as present-value logic in disguise.

2.3 Asset Prices

The first page of Cochrane (2009) explains how, in its current form, "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." Since this is how we have been taught to think, many researchers simply assume that investors must also be thinking carefully about how to discount a company's expected future payoffs to each shareholder. But it is not so.

The previous section documented that analysts largely ignore discount rates. Instead, they focus on getting their near-term earnings forecast right. 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 $\mathbb{E}_t[\mathrm{EPS}_{t+2}]$ rather than EPS_t today?"

Obviously, in a world where sell-side analysts use trailing P/E ratios, there is no reason to expect price targets to line up with the present discounted value of a firm's expected dividend stream. But, if realized prices also do not reflect present-value logic, 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.

Suppose that there exists a strictly positive constant, $\nu > 0$, such that

$$\left(\frac{\text{Price}_{t+1} - \text{Price}_t}{\text{Price}_t}\right) = \nu \cdot \left(\frac{\text{Demand}_{t+1} - \text{Demand}_t}{\text{Demand}_t}\right) + \varepsilon_{t+1} \tag{14}$$

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 even 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* (12) *and that realized price growth is governed by the law of motion in Equation* (14)*. If* $\nu = 1/\mu$*, then*

$$\hat{\mathbb{E}}_t[\mathit{Price}_{t+1}] = \mathbb{E}_t[\mathit{EPS}_{t+2}] \times \mathit{TrailingPE}_t \tag{15}$$

where $\hat{\mathbb{E}}_t[\mathit{Price}_{t+1}]$ *denotes the company's average price next year as estimated by an econometrician using market data.*

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 (11). 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 (12). 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 (14) 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.

Proposition 2.3 shows that analysts' price targets can be correct on average even if the firm's average price level next year does not equal expected discounted payoff. There is nothing pinning down the company's price level in our model. However, the absence of this boundary condition does imply a lack of economic structure. It is noteworthy that the law of motion for prices (Equation 14) involves the same sort of proportional thinking found in investors' demand rule (Equation 12). 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 (12) and (14) 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 said, our model focuses on demand multipliers and price elasticities for a very different reason. The demand-system framework focuses on μ and ν because they play a pivotal role in models where investors solve a forward-looking portfolio problem and fully 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.

Finally, it is important to emphasize that the expectation operator on the left-hand side of Equation (15) is different from the one on the right-hand side. On the left, $\hat{\mathbb{E}}_t[\text{Price}_{t+1}]$ represents a researcher's best guess about a company's share price next year at time $(t+1)$ given the observed market data today at time t. On the right, $\mathbb{E}_t[\text{EPS}_{t+2}]$ denotes analysts' subjective beliefs about the firm's expected EPS in the subsequent year. Proposition 2.3 maps a belief that analysts hold into an estimate that researchers can observe.

2.4 Exclusion Restriction

We now develop the key testable implication that emerges from taking seriously the economic logic behind this mapping. Think about a standard asset-pricing test, which looks at how returns react to new information. Let News_t denote a piece of information revealed about a firm at time t. In textbook asset-pricing models, if News_t predicts the company's future returns, then it must be correlated with investors' subjective beliefs about the expected discounted payoff. The entire effect must operate via this one specific channel.

"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 returns*

$$\widehat{\text{Corr}}\big(\,\hat{\mathbb{E}}_t \Delta \log Price_{t+1}, News_t\,\big) = 0 \tag{16}$$

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, $\widehat{\text{Corr}}(\mathbb{E}_t[r_{t+h}], News_t) \neq 0$.

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 Price $= A + B \cdot$ Signal $- C \cdot$ 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. In a world where B was larger, it would be less valuable to buy the private signal since it is possible to learn much of the same information for free by studying prices.

By contrast, the interesting thing about our model is all the things that *should not* affect prices. When setting price targets, analysts mainly focus on getting a company's expected near-term earnings right. They choose a trailing P/E ratio almost as an afterthought. So, provided that the basic functional forms in Equations (12) and (14) are correct, there is not much more to be learned from μ and ν. We should not expect μ and ν to satisfy subtle constraints. Proposition 2.4 shows that, if a piece of news affects a company's share price, it must do so via an effect on expected short-term earnings. Everything must operate through this one narrow channel according to our model.

How could prices fail to reflect information about a company's earnings three years from now? 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.

It may seem obvious to us as researchers that prices *should* reflect expected discounted payoffs. But this does not imply that market participants think in these terms or that market prices obey this constraint. Our analysis of sell-side research in Section 1 suggests that market participants do not. The simple model in this section predicts that market prices will not either.

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.*

3 Explanatory Power

We wrap up our analysis by showing that $\mathbb{E}[\text{EPS}] \times \text{TrailingPE}$ has a substantial amount of explanatory power in real-world data. We describe our data in Subsection 3.1. Then, in Subsection 3.2 we demonstrate that $\mathbb{E}[\text{EPS}] \times \text{TrailingPE}$ explains over 90% of the observed variation in analysts' price targets. Finally, in Subsection 3.3 we demonstrate that this expression also helps explain realized market prices, not just the price that analysts expect to observe.

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/sh for Home Depot in December 2020 (target date). We distinguish between trading days t and target dates τ because an analyst can revise his/her forecast for the same target date on successive trading days. For each analyst a tracking a particular firm n, we record their most recent price target, $\text{PriceTarget}_{n,t}^{a} = \mathbb{E}_t^a[\text{Price}_{n,\tau+1}]$, from 18 months to 6 months prior to each target date $(\tau + 1)$.

We write the analyst's corresponding EPS forecast as $\mathbb{E}_t^a[\text{EPS}_n]$. We use the two-year-ahead EPS forecast when available in IBES, $\mathbb{E}_t^a[\text{EPS}_{n,\tau+2}]$, otherwise we use the one-year-ahead value, $\mathbb{E}_t^a[\text{EPS}_{n,\tau+1}]$. We restrict our sample to include observations with a positive EPS forecast, $\mathbb{E}_t^a[\text{EPS}_n] \geq \$0.01/\text{sh}$. We also require firms to have a price target greater than \$1/sh and less than \$10,000/sh.

The resulting panel data set is organized by firm × analyst × target date. We have already shown what this panel looks like for Chris Horvers' coverage of Home Depot in Figure 3 and Andrea Teixeira's coverage of Coca-Cola (KO) for JP Morgan in Figure 6. See Figures B1(a)-B1(n) in Appendix B for more examples.

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

$$\text{ImpliedPE}_{n,t}^{a} \overset{\text{def}}{=} \frac{\text{PriceTarget}_{n,t}^{a}}{\mathbb{E}_t^a[\text{EPS}_n]} \tag{17}$$

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 $\text{ImpliedPE}_{n,t}^{a} = \text{TrailingPE}_{n,t}$ exactly. Figure 17 shows the cross-sectional distribution of $\text{ImpliedPE}_{n,t}^{a}$ and $\text{TrailingPE}_{n,t}$.

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 require both these P/E ratios to be 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. However, we show in Appendix B Figures B2(a)-B2(e) that our findings extend outside this range.

3.2 Predicting Price Targets

In Section 1 our paper, we documented that sell-side analysts tend to set price targets by multiplying a near-term earnings forecast times a trailing P/E ratio. This first part of our analysis came from reading through a sample of 513 earnings reports. We now verify that this finding extends to the price targets set by other analysts covering other firms.

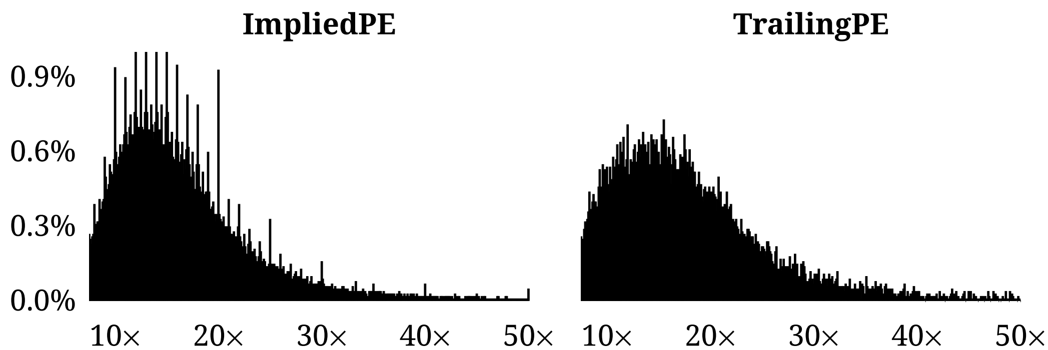

Figure 17. *Histograms showing the distribution of ImpliedPE$^a_{n,t}$ (left panel) and TrailingPE$_{n,t}$ (right panel) for all sell-side analyst reports in our sample with* $\mathbb{E}^a_t[EPS_n] \geq \$1.00$ *from 2003 to 2022. x-axis denotes the P/E ratio in increments of* $0.1\times$. *y-axis represents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

Table 10 shows the results of estimating the regression specification below

$$\log(\text{PriceTarget}^a_{n,t}) \sim \hat{\alpha} + \hat{\beta} \cdot \log(\mathbb{E}^a_t[\text{EPS}_n]) + \hat{\gamma} \cdot \log(\text{TrailingPE}_{n,t}) + \cdots \tag{18}$$

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 set price targets using the formula $\mathbb{E}[\text{EPS}] \times \text{TrailingPE}$, 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 trailing P/E formula generates an adjusted $R^2 = 91.0\%$. Most of the variation in analysts' price targets can be explained by $\mathbb{E}[\text{EPS}] \times \text{TrailingPE}$. Columns (2)-(4) in Table 10 verify that firm, analyst, and month fixed-effects do not change this main message.

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

Table 10. *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.*

Asset-pricing researchers are used to seeing R^2s in the low single digits (Campbell and Thompson, 2008; Welch and Goyal, 2008). This can make it difficult to really appreciate what $R^2 = 91.0\%$ really means. At the very least, we know of two asset-pricing researchers who ran into this problem when trying to make sense of our early results.

We have found that binned scatterplots do a much better job of conveying the tight fit between theory and data. 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 $\mathbb{E}[\text{EPS}] \times \text{TrailingPE}$, 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.

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*

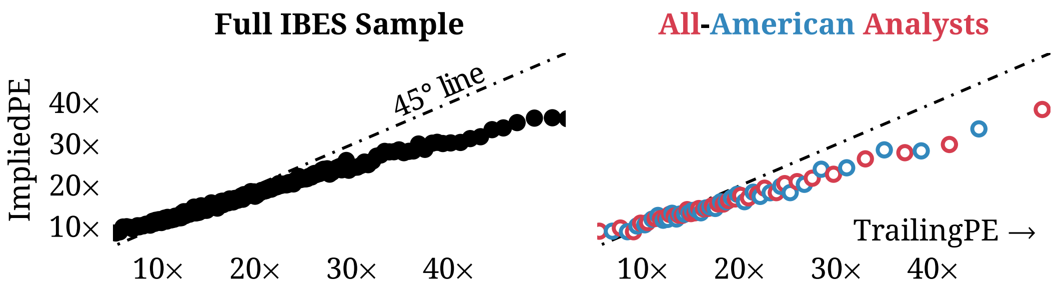

Figure 18. *(Left) Binned scatterplot using data from the full sample of IBES reports. x-axis shows the firm's trailing twelve-month P/E, $TrailingPE_{n,t} = Price_{n,t} \,/\, EPS_{n,t}$. y-axis shows the P/E ratio implied by the analyst's price target and EPS forecast, $ImpliedPE^a_{n,t} \overset{\text{def}}{=} PriceTarget^a_{n,t} \,/\, \mathbb{E}^a_t[EPS_n]$. (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.*

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. This time around, the specification is given by

$$\text{ImpliedPE}^a_{n,t} \sim \hat{\eta} + \hat{\theta} \cdot \text{TrailingPE}_{n,t} + \cdots \tag{19}$$

If sell-side analysts were exclusively using $\mathbb{E}[\text{EPS}] \times \text{TrailingPE}$ 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 smart people. When they think other information might be relevant, they will use more than just $\mathbb{E}[\text{EPS}] \times \text{TrailingPE}$ when setting a price target. This

Dep variable:	$\text{ImpliedPE}^a_{n,t}$			
	(1)	(2)	(3)	(4)
$\text{TrailingPE}_{n,t}$	0.58***	0.43***	0.58***	0.52***
	(0.01)	(0.01)	(0.01)	(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* (19)*. 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.*

simple formula is the standard starting point for their analysis. Our line is flatter than one because analysts are more likely to deviate toward the mean 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 set a target price using a P/E ratio of 20×. The same thing is true when a firm has a trailing P/E of 20.1×.

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.

Figure 19 shows what happens when we regress an analyst's implied P/E on the company's realized P/E in each of the last 20 quarters

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

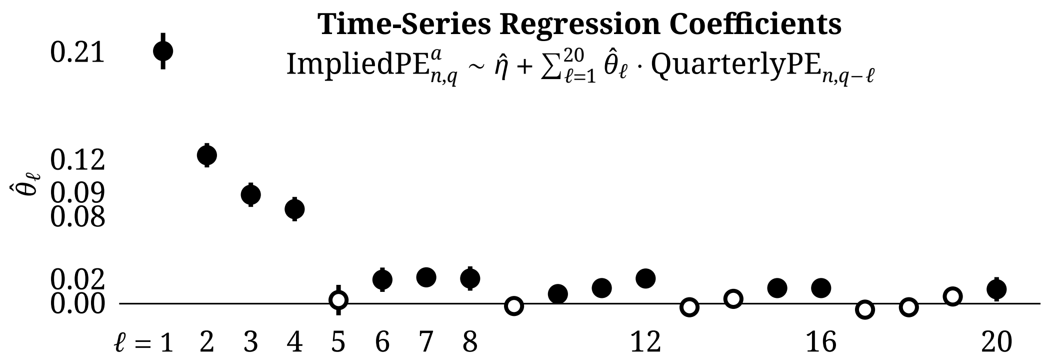

Figure 19. *Each dot denotes one of the* 20 *estimated slope coefficients,* $\{\hat{\theta}_\ell\}_{\ell=1}^{20}$*, from the regression specification in Equation* (20)*. ImpliedPE*$^a_{n,q}$*: P/E ratio implied by an analyst's price target and the relevant annual EPS forecast. QuarterlyPE*$_{n,q}$*: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

Let $\text{eps}_{n,q}$ 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. 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 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.

Also 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 Predicting Realized Returns

In principle, even if sell-side analysts are using PriceTarget $= \mathbb{E}[\text{EPS}] \times$ TrailingPE, other market participants might be applying textbook present-value logic. In such a scenario, analysts' price targets would reflect trailing P/E ratios, but realized future returns would not. In this section, we provide evidence that trailing P/E ratios predict realized future returns as well.

To provide the cleanest possible empirical setting, we study firms' returns following an earnings surprise. Suppose there are a large number of investors who forecast future price levels based on $\mathbb{E}[\text{EPS}] \times \text{TrailingPE}$. If we hold constant the size of the earnings surprise, then there is only one free parameter in this forecasting rule: TrailingPE. Hence, if stocks A and B both realized the same earnings surprise in a given quarter, $\text{eps}_A - \mathbb{E}_{t_{\text{Ancmt}}-1}[\text{eps}_A] = \text{eps}_B - \mathbb{E}_{t_{\text{Ancmt}}-1}[\text{eps}_B] = s$, then our model says that each stock's return following the earnings announcement will be proportional to $s \times \text{TrailingPE}$. If stock A has $\text{TrailingPE}_A = 20\times$ while stock B has $\text{TrailingPE}_B = 10\times$, then stock A's price reaction should be double that of stock B's.

To test this prediction, we first group all the firm-quarter observations that realized the same-sized earnings surprise into bins. Let $N_s \stackrel{\text{def}}{=} \{ (n,q) : \text{eps}_{n,q} - \mathbb{E}_{t_{\text{Ancmt}}-1}[\text{eps}_{n,q}] = s \}$ denote the set of all observations that realized an earnings surprise of s dollar per share. Then, within each bin, we regress the realized price change over the quarter following the earnings surprise on a firm's trailing P/E ratio at the time of the announcement

$$\Delta\text{Price}_{n,q+1} \sim \hat{\kappa}_s + \hat{\lambda}_s \cdot \text{TrailingPE}_{n,q} \qquad \begin{array}{c}\text{using data on firm-qtr}\\ \text{obs where } (n,q) \in N_s\end{array} \tag{21}$$

Finally, we run a second-stage regression to see if the estimated slope coefficients, $\{\hat{\lambda}_s\}$, increase linearly in the size of the underlying earnings surprise, s.

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_{Market}, by running a separate time-series

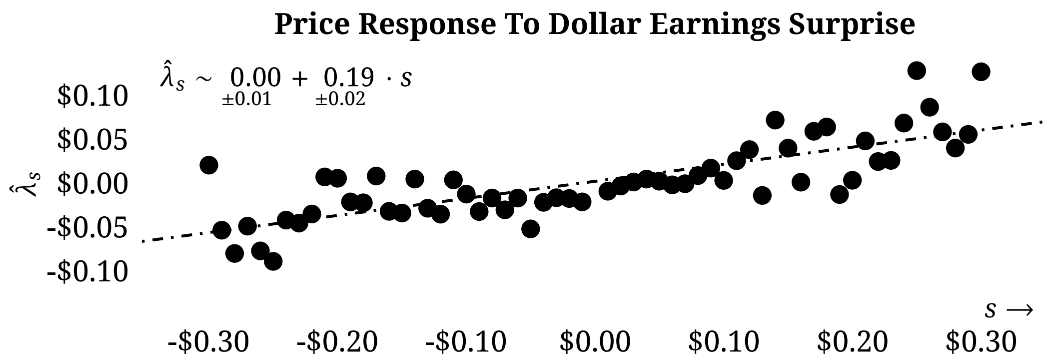

Figure 20. *Each dot denotes an estimated slope coefficient, $\hat{\lambda}_s$, from one of* 60 *separate regressions like the one shown in Equation* (21)*. The y-axis shows the first-stage slope coefficient, $\hat{\lambda}_s$, estimated using data on all firm-quarter observations that had the same dollar earnings surprise, s. The x-axis shows the corresponding value of s in* \$0.01*/sh bins. The highest bin is centered at* \$0.30*/sh while the lowest bin is centered at* $-$0.30*/sh. This gives us* 60 *data points since we omit the* $s = \$0.00$*/sh bin containing observations with no earnings surprise. The dashed line is the best-fit OLS equation.*

regression. Finally, you check whether there is a linear relationship between each portfolio's excess returns and its estimated β_{Market}.

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 $s = \$0.00$/sh bin containing firm-quarter observations with no earnings surprise, $\text{eps}_{n,q} = \mathbb{E}_{t_{\text{Ancmt}}-1}[\text{eps}_{n,q}]$.

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 company's future earnings. Hence, in a world where investors were applying present-value logic, we should see them apply a different multiple following large earnings surprises. This is clear evidence of the exclusion restriction we outline in Proposition 2.4. The only way that earnings surprises affect subsequent returns is by changing investors' beliefs about short-term earnings.

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

Table 12. *Each column reports the results of a separate second-stage regression. The left-hand-side variable in each regression is an estimated slope coefficient, $\hat{\lambda}_s$, for each earnings surprise bin like in Equation (21). Column (1) reports results using 60 separate \$0.01/sh bins centered at $\{-\$0.30/sh, \ldots, -\$0.01/sh, \$0.01/sh, \ldots, \$0.30/sh\}$. These results match the dashed best-fit line in Figure 20. Column (2) shows results where we group observations into 30 separate \$0.02/sh bins centered at $\{-\$0.30/sh, \ldots, -\$0.02/sh, \$0.02/sh, \ldots, \$0.30/sh\}$. Column (3) shows a similar analysis using 12 bins that are \$0.05/sh wide, $\{-\$0.30/sh, \ldots, -\$0.05/sh, \$0.05/sh, \ldots, \$0.30/sh\}$. All three columns omit the bin centered at $s = \$0.00/sh$—i.e., observations without an earnings surprise.*

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 a $s = \$0.10/\text{sh}$ earnings surprise, the price of stock A with $\text{TrailingPE}_A = 20\times$ will increase by $\$0.19/\text{sh} = 0.19 \cdot \{\$0.10/\text{sh} \cdot 20 - \$0.10/\text{sh} \cdot 10\}$ more than the price of stock B with $\text{TrailingPE}_B = 10\times$. If both firms had realized a $s = \$0.20/\text{sh}$ earnings surprise, then stock A's price would go up by \$0.38/sh more than stock B's over the subsequent quarter. A $s = \$0.30/\text{sh}$ earnings surprise would lead to a \$0.57/sh difference between the two stocks' price growth.

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/sh wide and \$0.05/sh wide. We get quantitatively similar results no matter how finely we divide our portfolios. The intercept is always a precisely estimated zero. This straight line exists because investors are using a trailing P/E ratio to update a firm's price following earnings surprises.

Conclusion

How do investors choose the rate to discount a company's future cash flows when pricing its equity shares? In his 2010 presidential address, Cochrane (2011) calls this "the central organizing question of current asset-pricing research." Discount rates may be important to researchers. But, to many investors, they are a sideshow. In this paper, we show that sell-side analysts do not take a forward-looking present-value approach to setting price targets. Instead, they multiply a firm's near-term EPS forecast times a trailing P/E ratio.

Perhaps this is a bad thing. Maybe investors should be thinking more carefully about discount rates. No matter. This is not what investors are doing. Going forward, when a researcher wants to predict how investors will forecast future price levels, they should model investors who multiply an earnings forecast times a defensible recent multiple. That should be the starting point of the model…not because it is optimal but because it is what investors actually do. This is the central premise of our paper, and we provide a tractable theoretical model to help researchers get started down this path.

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.

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, 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.

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

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

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

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

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

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.

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

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

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

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

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

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

Groysberg, B., P. Healy, and D. Maber (2011). What drives sell-side analyst compensation at high-status investment banks? *Journal of Accounting Research 49*(4), 969–1000.

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.

Lintner, J. (1965). The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets. *Review of Economics and Statistics 47*(1), 13–37.

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.

Manski, C. (1995). *Identification problems in the social sciences*. Harvard University Press.

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.

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

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

Sharpe, W. (1964). Capital asset prices: A theory of market equilibrium under conditions of risk. *Journal of Finance 19*(3), 425–442.

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.

Treynor, J. (1961). Market value, time, and risk. Working paper.

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* (11)

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

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

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

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

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

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

Proof. **(Proposition 2.3)** *Suppose that year-over-year price growth is governed by the law of motion in Equation* (14)*. 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 can 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 beliefs about the firm's earnings and 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

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

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

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

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 $\hat{\mathbb{E}}_t \Delta \log Price_{t+1}$. The right-hand side is a function of the analysts' expectations about short-term earnings as defined in Equation (10), which we have reproduced below

$$\left(\frac{EPS_{t+1} - EPS_t}{EPS_t} \right) = X_t + \epsilon_{t+1} \tag{10}$$

In this equation, $X_t \approx \mathbb{E}_t[\Delta \log EPS_{t+1}]$ 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, if a signal is uncorrelated with X_t, then it cannot predict $\hat{\mathbb{E}}_t \Delta \log Price_{t+1}$. □

B Additional Results

Katy Huberty's forecasts for Apple (AAPL)

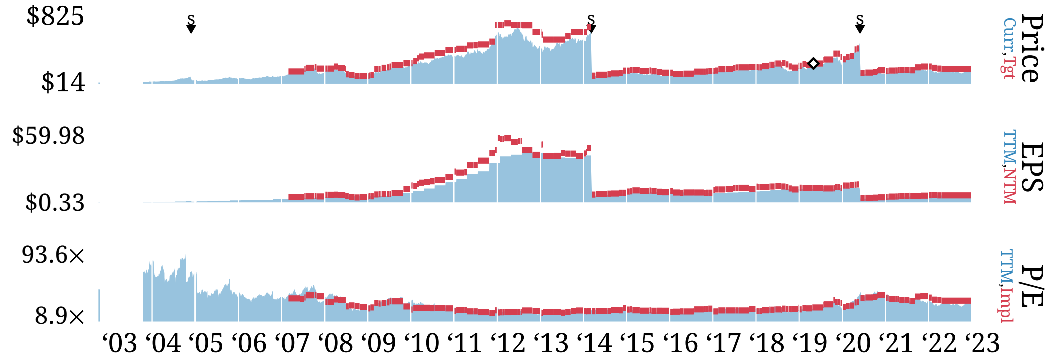

Figure B1(a). *y-axis shows min, median, and max. (Top) Blue ribbon is Apple's closing price on day t from CRSP, $Price_t$. Red line is Katy Huberty's price target, $PriceTarget_t = \mathbb{E}_t[Price_{\tau+1}]$, in IBES. (Middle) Blue is AAPL's trailing twelve-month (TTM) EPS on day t from IBES, EPS_t. 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.*

Jamie Baker's forecasts for Alaska Airlines (ALK)

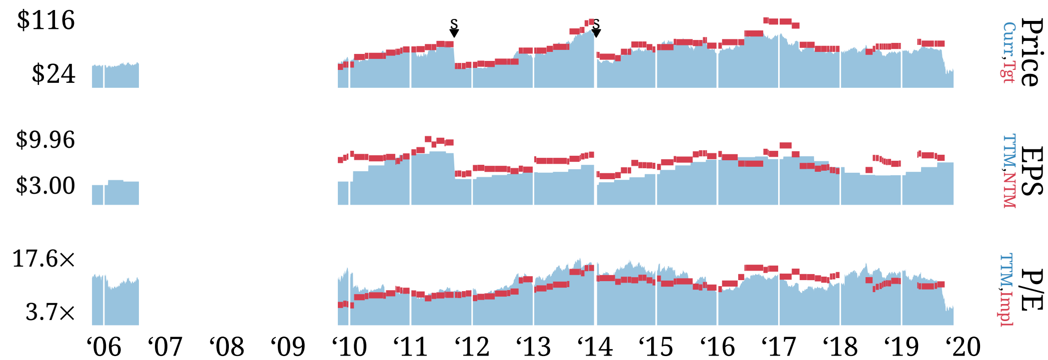

Figure B1(b). *y-axis shows min, median, and max. (Top) Blue ribbon is Alaska Airlines' closing price on day t from CRSP, $Price_t$. Red line is Jamie Baker's price target, $PriceTarget_t = \mathbb{E}_t[Price_{\tau+1}]$, in IBES. (Middle) Blue is ALK's trailing twelve-month (TTM) EPS on day t from IBES,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

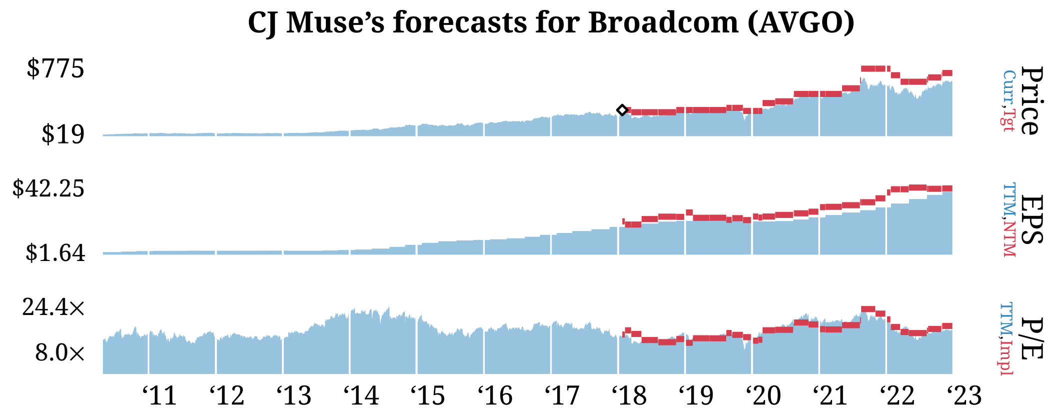

Figure B1(c). *y-axis shows min, median, and max. (Top) Blue ribbon is Broadcom's closing price on day t from CRSP, $Price_t$. Red line is CJ Muse's price target, $PriceTarget_t = \mathbb{E}_t[Price_{\tau+1}]$, in IBES. (Middle) Blue is Broadcom's trailing twelve-month (TTM) EPS on day t from IBES, EPS_t. 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.*

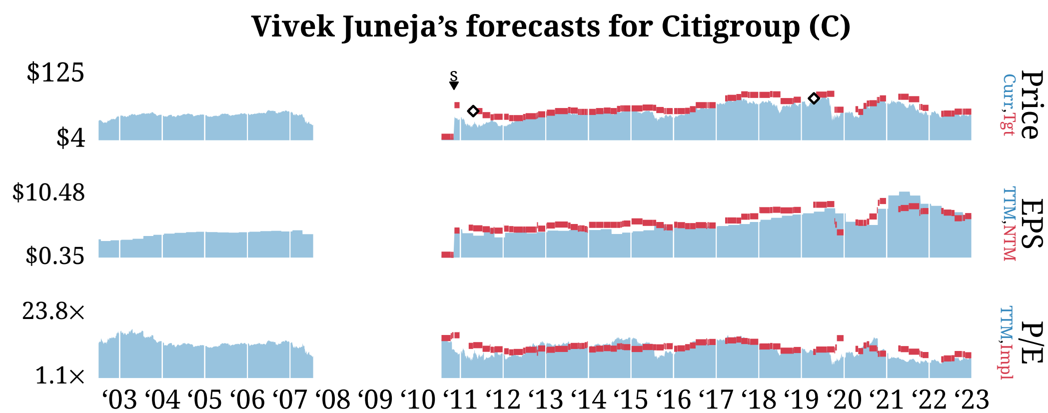

Figure B1(d). *y-axis shows min, median, and max. (Top) Blue ribbon is Citigroup's closing price on day t from CRSP, $Price_t$. Red line is Vivek Juneja's price target, $PriceTarget_t = \mathbb{E}_t[Price_{\tau+1}]$, in IBES. (Middle) Blue is Citi's trailing twelve-month (TTM) EPS on day t from IBES,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

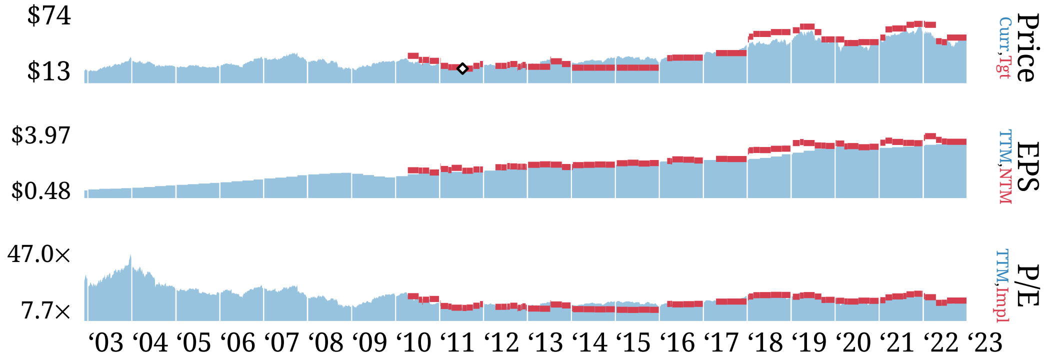

Figure B1(e). *y-axis shows min, median, and max. (Top) Blue ribbon is Cisco System's closing price on day t from CRSP,* $Price_t$. *Red line is Rod Hall's price target,* $PriceTarget_t = \mathbb{E}_t[Price_{\tau+1}]$, *in IBES. (Middle) Blue is Cisco's trailing twelve-month (TTM) EPS on day t from IBES,* 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)

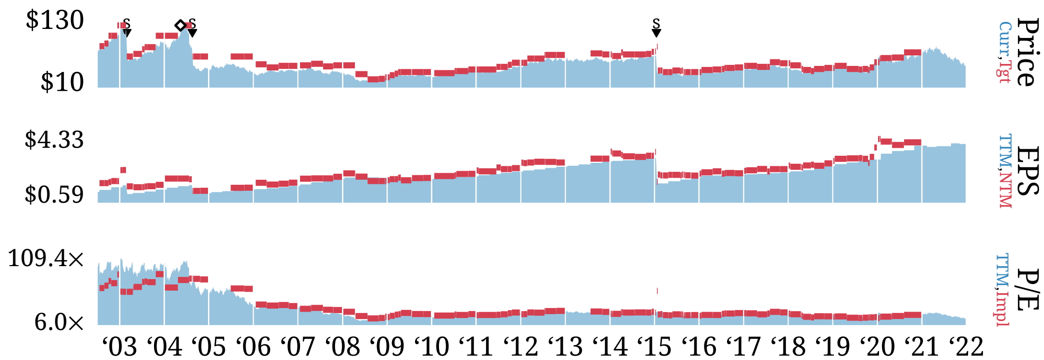

Figure B1(f). *y-axis shows min, median, and max. (Top) Blue ribbon is eBay's closing price on day t from CRSP,* $Price_t$. *Red line is Heath Terry's price target,* $PriceTarget_t = \mathbb{E}_t[Price_{\tau+1}]$, *in IBES. (Middle) Blue is eBay's trailing twelve-month (TTM) EPS on day t from IBES,*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

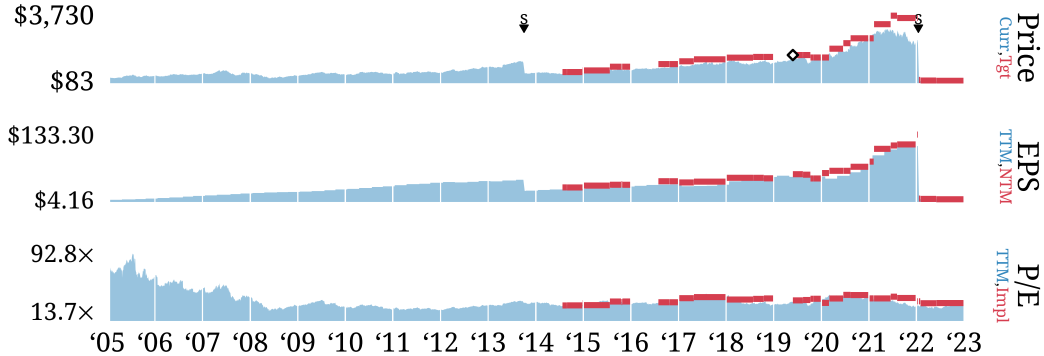

Figure B1(g). *y-axis shows min, median, and max. (Top) Blue ribbon is Google's closing price on day t from CRSP, $Price_t$. Red line is Christophe Cherblanc's price target, $PriceTarget_t = \mathbb{E}_t[Price_{\tau+1}]$, in IBES. (Middle) Blue is Google's trailing twelve-month (TTM) EPS on day t from 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)

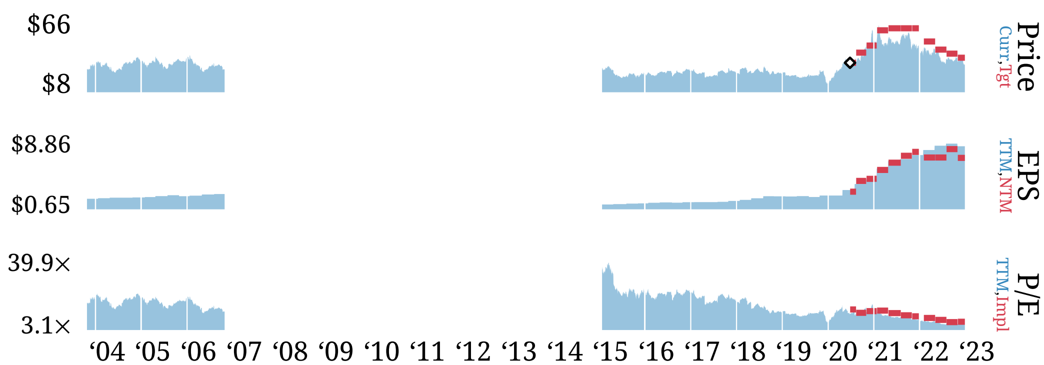

Figure B1(h). *y-axis shows min, median, and max. (Top) Blue ribbon is MarineMax's closing price on day t from CRSP, $Price_t$. Red line is Greg Badishkanian's price target, $PriceTarget_t = \mathbb{E}_t[Price_{\tau+1}]$, in IBES. (Middle) Blue is HZO's trailing twelve-month (TTM) EPS on day t from IBES,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

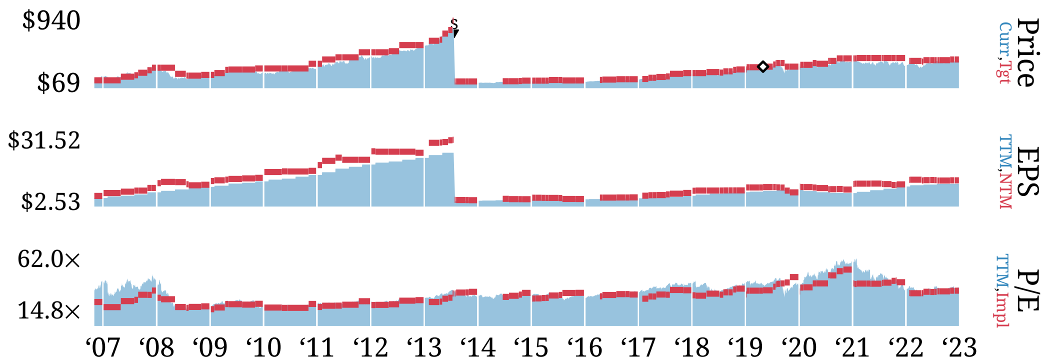

Figure B1(i). *y-axis shows min, median, and max. (Top) Blue ribbon is Mastercard's closing price on day t from CRSP, $Price_t$. Red line is Moshe Orenbuch's price target, $PriceTarget_t = \mathbb{E}_t[Price_{\tau+1}]$, in IBES. (Middle) Blue is MA's trailing twelve-month (TTM) EPS on day t from IBES, 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)

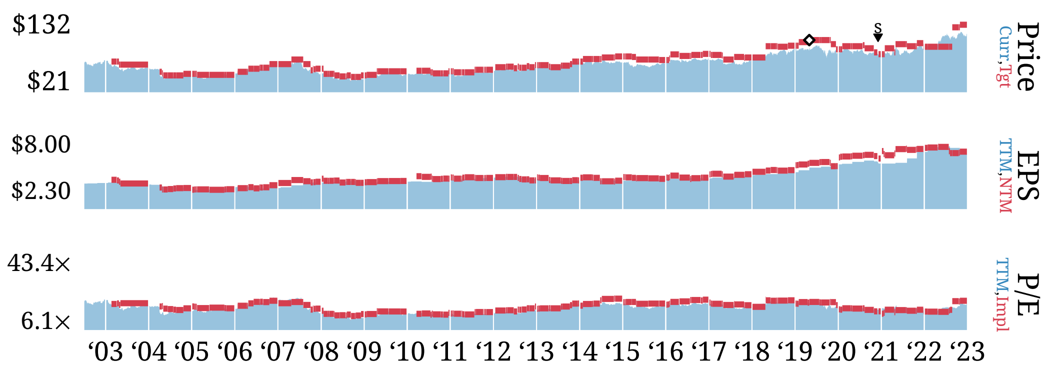

Figure B1(j). *y-axis shows min, median, and max. (Top) Blue ribbon is Merck's closing price on day t from CRSP, $Price_t$. Red line is Tim Anderson's price target, $PriceTarget_t = \mathbb{E}_t[Price_{\tau+1}]$, in IBES. (Middle) Blue is MRK's trailing twelve-month (TTM) EPS on day t from IBES, EPS_t.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

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

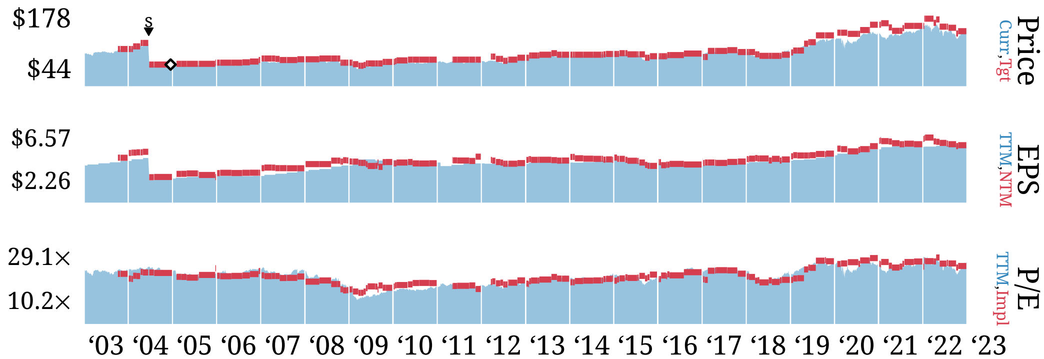

Figure B1(k). *y-axis shows min, median, and max. (Top) Blue ribbon is Procter & Gamble's closing price on day t from CRSP, $Price_t$. Red line is Lauren Lieberman's price target, $PriceTarget_t = \mathbb{E}_t[Price_{\tau+1}]$, in IBES. (Middle) Blue is PG's trailing twelve-month (TTM) EPS on day t from 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.*

Brian Tunick's forecasts for Ross Stores (ROST)

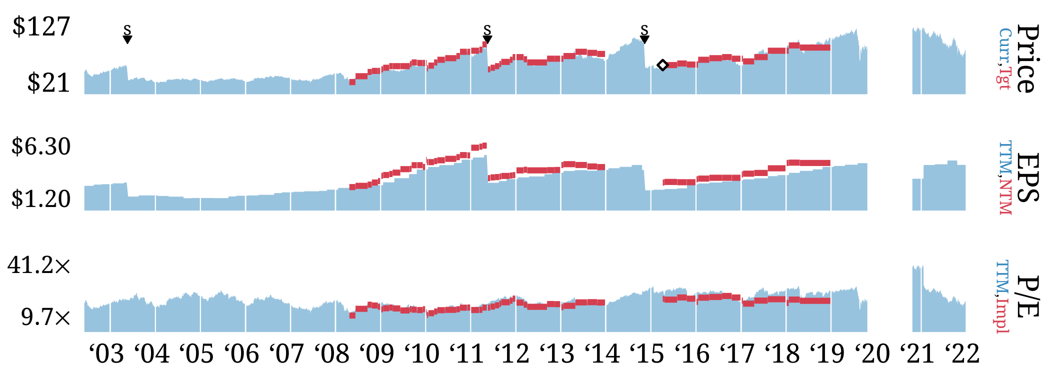

Figure B1(l). *y-axis shows min, median, and max. (Top) Blue ribbon is Ross' closing price on day t from CRSP, $Price_t$. Red line is Brian Tunick's price target, $PriceTarget_t = \mathbb{E}_t[Price_{\tau+1}]$, in IBES. (Middle) Blue is Ross' trailing twelve-month (TTM) EPS on day t from IBES, EPS_t.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

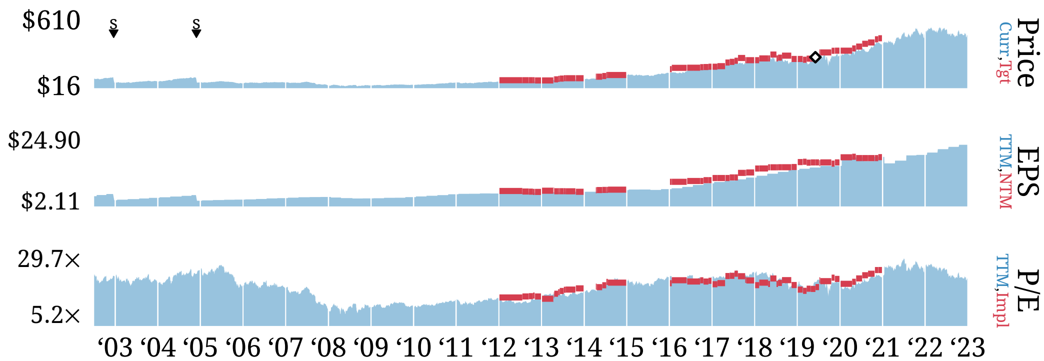

Figure B1(m). *y-axis shows min, median, and max. (Top) Blue ribbon is United Healthcare's closing price on day t from CRSP, $Price_t$. Red line is Justin Lake's price target, $PriceTarget_t = \mathbb{E}_t[Price_{\tau+1}]$, in IBES. (Middle) Blue is UNH's trailing twelve-month (TTM) EPS on day t from 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)

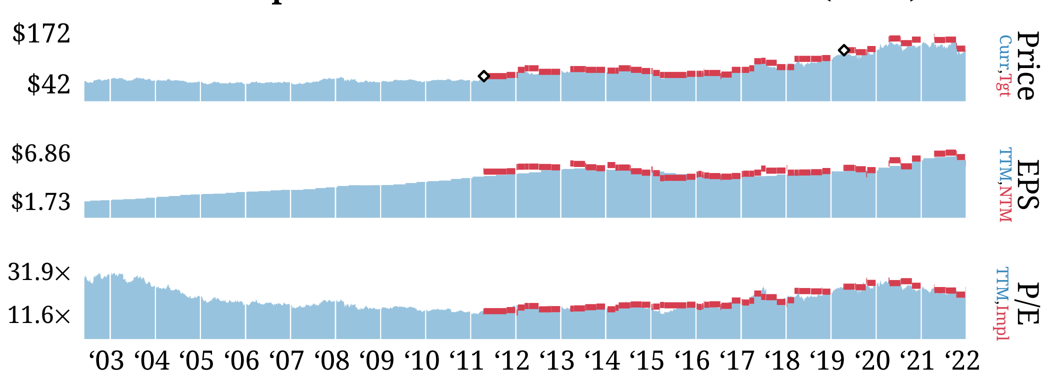

Figure B1(n). *y-axis shows min, median, and max. (Top) Blue ribbon is Walmart's closing price on day t from CRSP, $Price_t$. Red line is Chris Horvers' price target, $PriceTarget_t = \mathbb{E}_t[Price_{\tau+1}]$, in IBES. (Middle) Blue is WMT's trailing twelve-month (TTM) EPS on day t from IBES, EPS_t.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

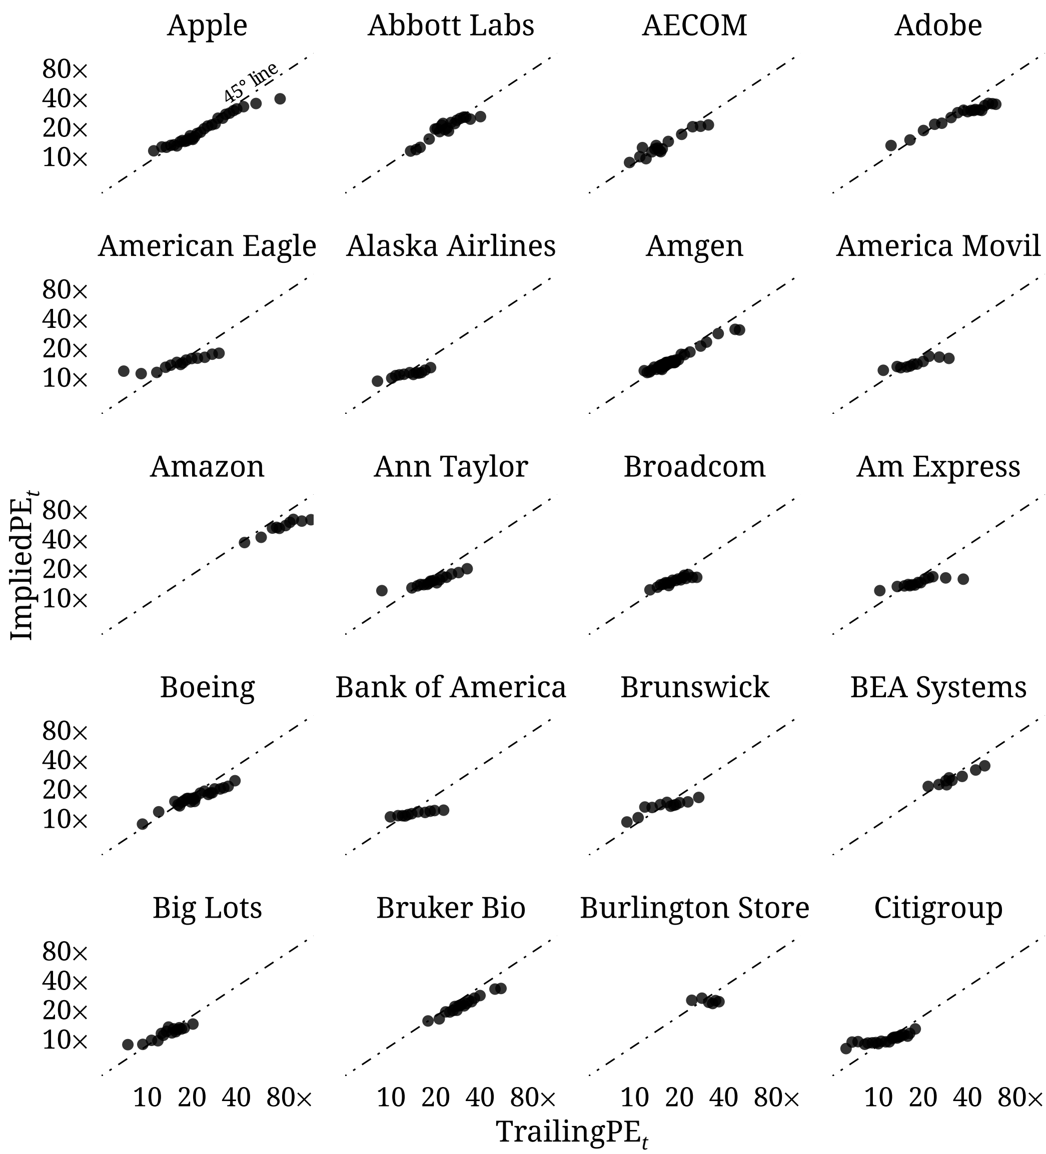

Figure B2(a). *Each panel shows a binned scatterplots using data from the full sample of IBES reports for a single firm. x-axis shows the firm's trailing twelve-month P/E,* $TrailingPE_{n,t} = Price_{n,t} \,/\, EPS_{n,t}$*. y-axis shows the P/E ratio implied by the analyst's price target and EPS forecast,*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

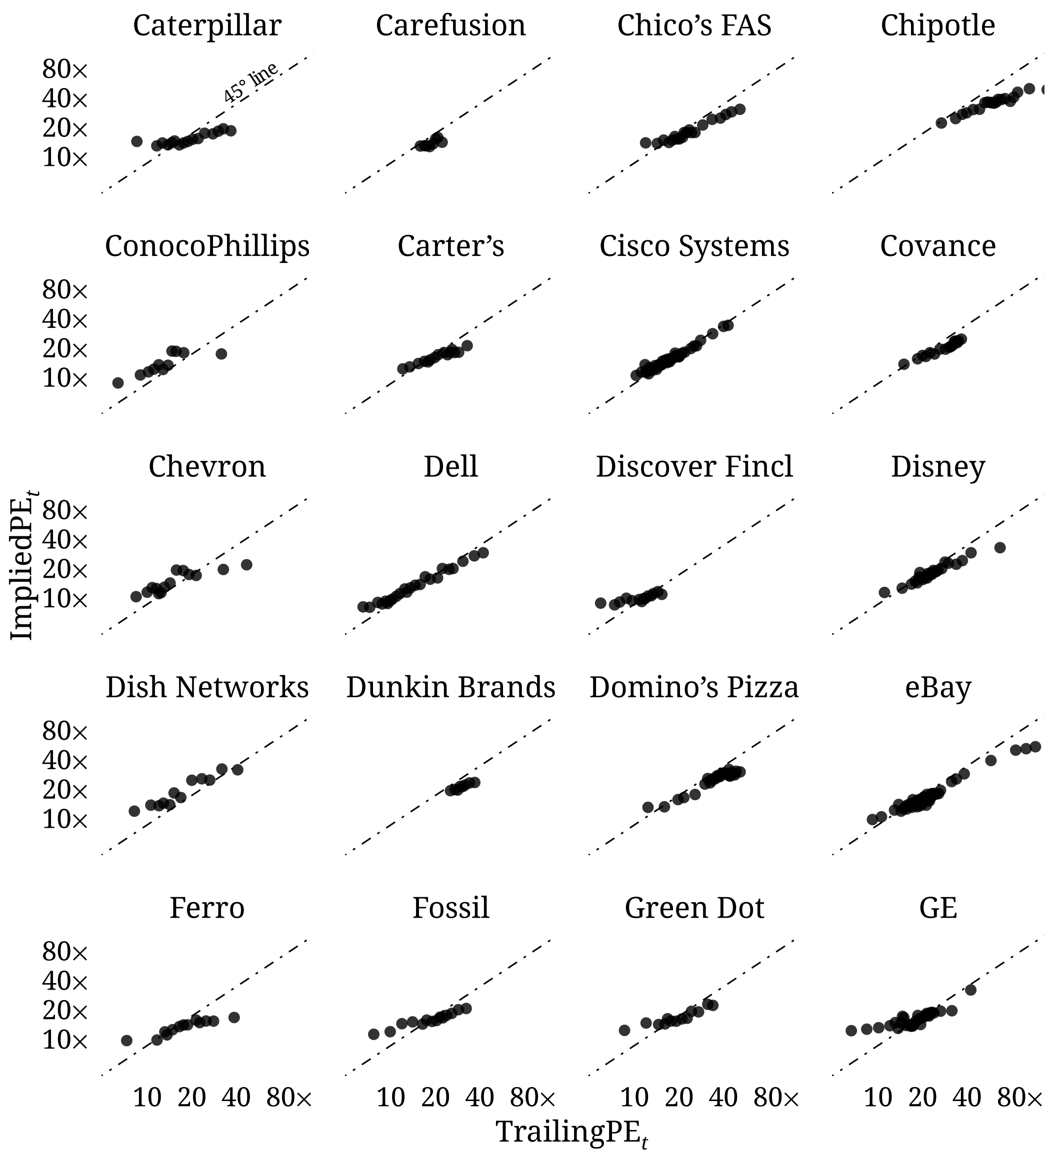

Figure B2(b). *Each panel shows a binned scatterplots using data from the full sample of IBES reports for a single firm. x-axis shows the firm's trailing twelve-month P/E,* $TrailingPE_{n,t} = Price_{n,t} \,/\, EPS_{n,t}$. *y-axis shows the P/E ratio implied by the analyst's price target and EPS forecast,*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

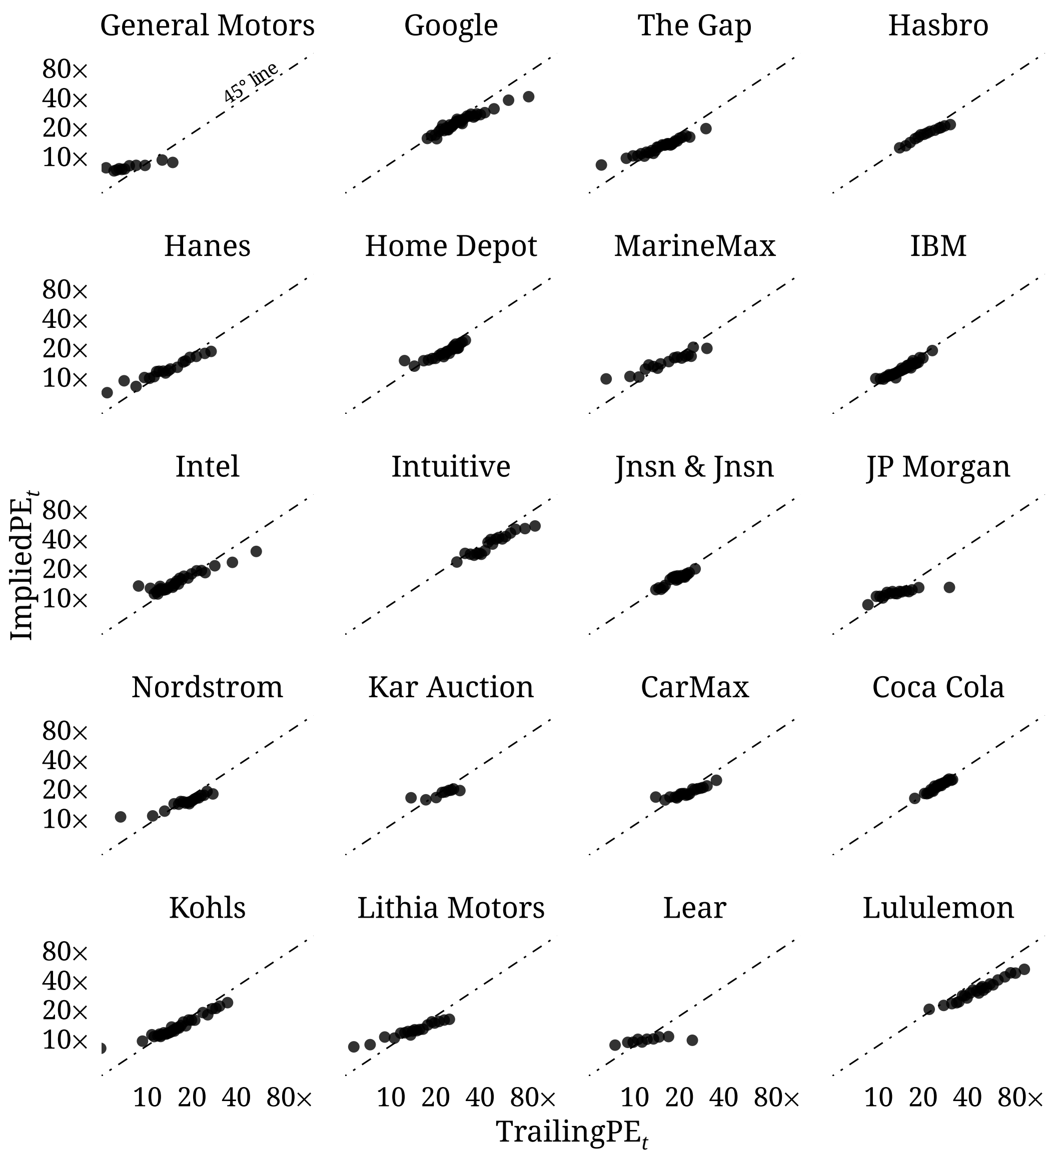

Figure B2(c). *Each panel shows a binned scatterplots using data from the full sample of IBES reports for a single firm. x-axis shows the firm's trailing twelve-month P/E, $TrailingPE_{n,t} = Price_{n,t} / EPS_{n,t}$. y-axis shows the P/E ratio implied by the analyst's price target and EPS forecast,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

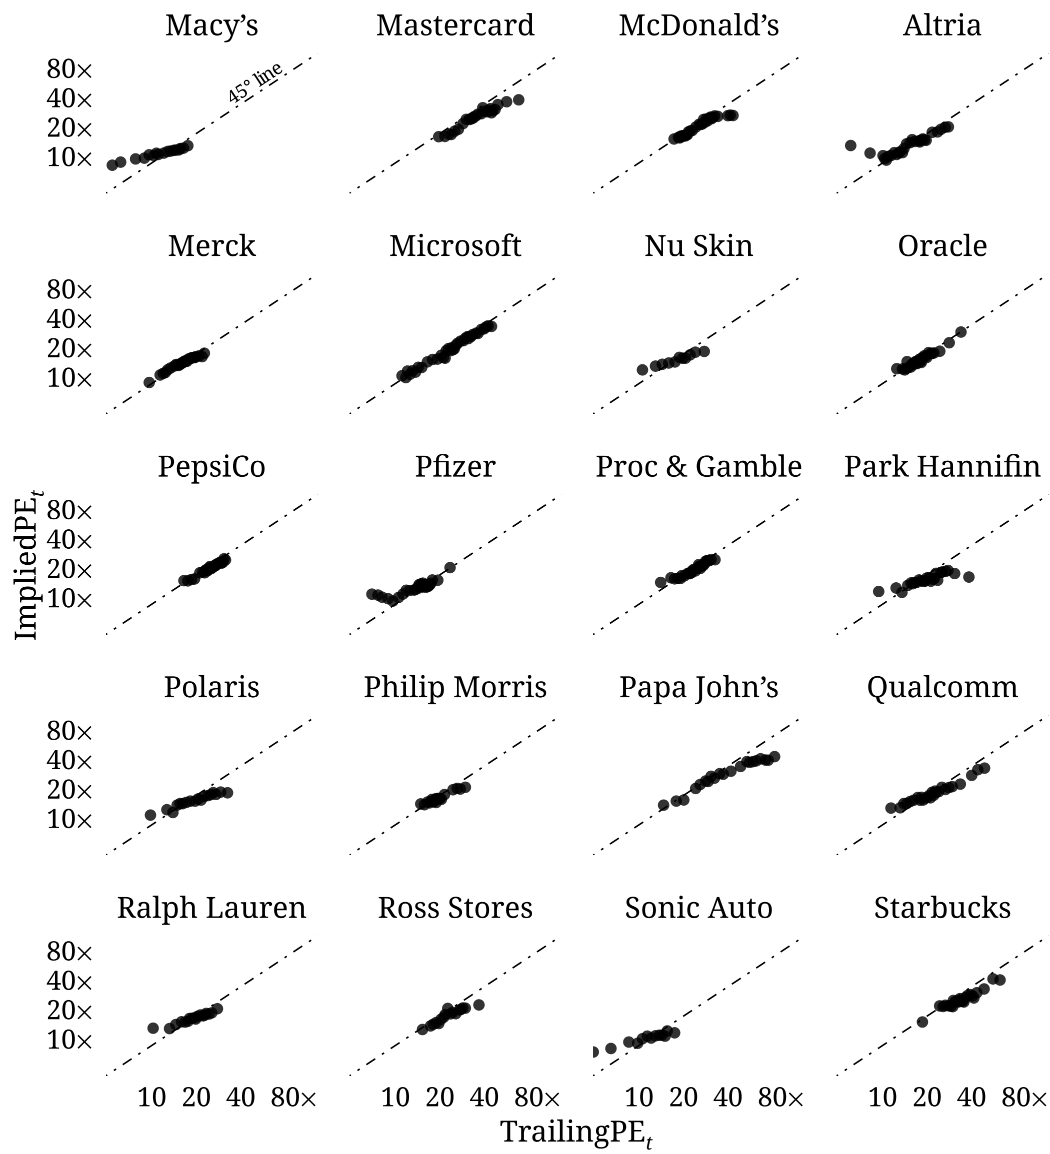

Figure B2(d). *Each panel shows a binned scatterplots using data from the full sample of IBES reports for a single firm. x-axis shows the firm's trailing twelve-month P/E,* $\text{TrailingPE}_{n,t} = \text{Price}_{n,t} \,/\, \text{EPS}_{n,t}$*. y-axis shows the P/E ratio implied by the analyst's price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

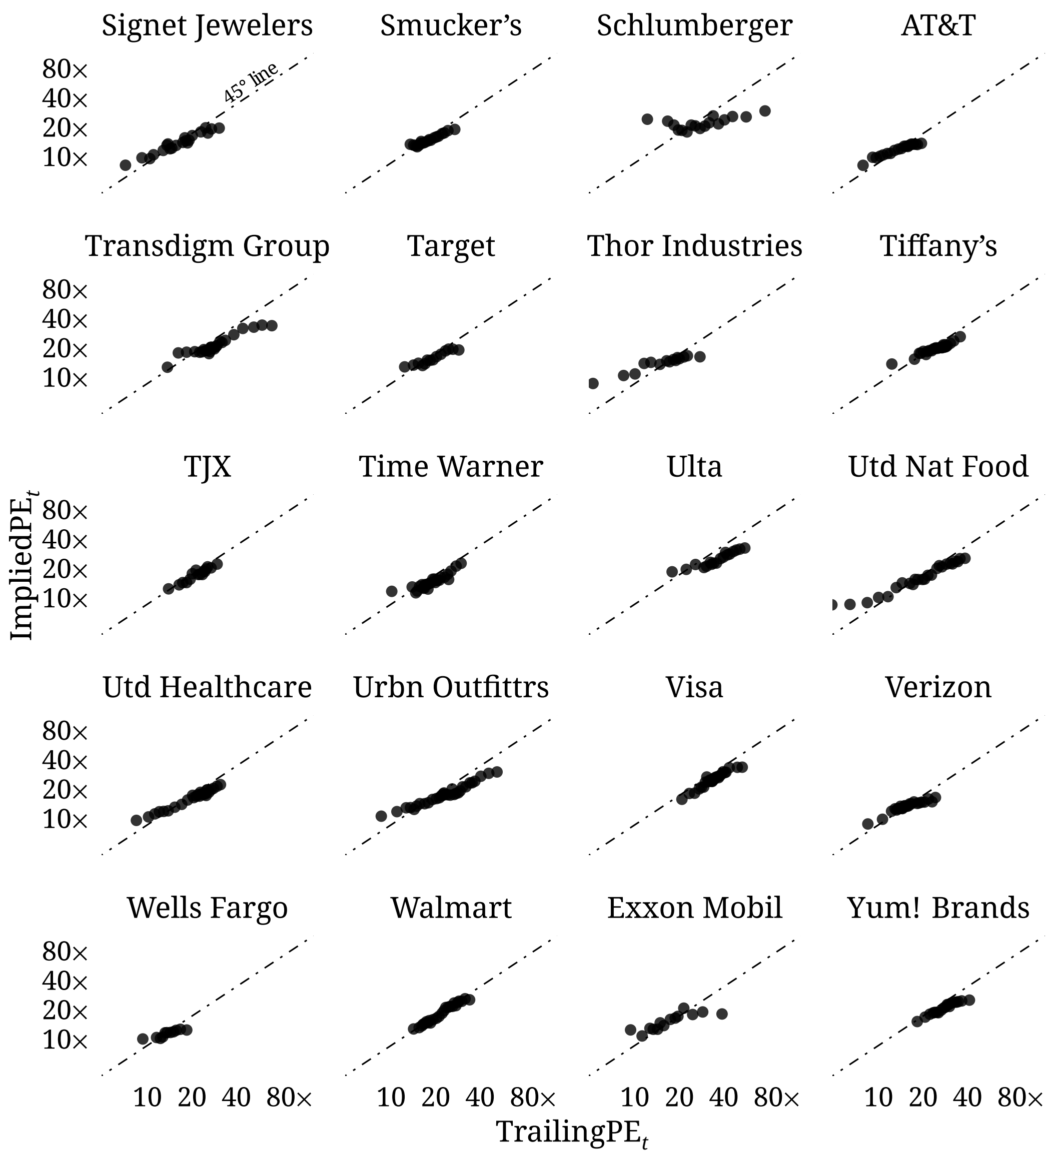

Figure B2(e). *Each panel shows a binned scatterplots using data from the full sample of IBES reports for a single firm. x-axis shows the firm's trailing twelve-month P/E, $TrailingPE_{n,t} = Price_{n,t} \,/\, EPS_{n,t}$. y-axis shows the P/E ratio implied by the analyst's price target and EPS forecast,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