THREE ESSAYS ON FINANCIAL MARKETS

THREE ESSAYS ON FINANCIAL MARKETS

By Qiao Wang,

A Thesis Submitted to the School of Graduate Studies in the Partial Fulfillment of the Requirements for the Degree Doctor of Philosophy

McMaster University

Doctor of Philosophy (2024)

Hamilton, Ontario (School of Business)

TITLE: THREE ESSAYS ON FINANCIAL MARKETS

AUTHOR: Qiao Wang (McMaster University)

SUPERVISOR: Dr. Ronald Balvers (McMaster University)

COMMITTEE: Dr. Narat Charupat (McMaster University)

COMMITTEE: Dr. Richard Deaves (McMaster University)

EXTERNAL: Dr. Marie-Hélène Gagnon, (Université Laval)

NUMBER OF PAGES: xii, 159

Abstract

This thesis comprises three distinct but interconnected studies in the field of financial markets, each exploring different facets of financial markets: asset pricing, sustainable investment and behavioral finance.

The first paper derives stock returns for firms producing non-renewable commodities employing the investment-based asset pricing approach. By identifying the appropriate time-varying discount rate the investment-based approach allows an alternative test of the Hotelling Valuation Principle. The empirical results support the principle and enable predicting returns from sorting firms into quintiles by expected return, producing a 16-20 percent realized difference between top and bottom quintile. The return differences cannot be explained by standard risk factors or a commodity-specific factor, suggesting that an important risk factor is still missing from standard models. The approach permits cost-of-capital estimation that circumvents identifying systematic risk factors.

The second paper examines whether the carbon pricing risk factor is priced in the cross-section of commodity futures. Analyzing unexpected pricing shocks in carbon emission allowances, it is shown that carbon pricing risk carries a significant positive risk premium in commodity markets. The study reveals that commodity sensitivities to carbon pricing risk vary, influenced by commodity-specific characteristics such as basis and hedging pressure. Additionally, a portfolio of commodity futures constructed based on carbon pricing beta offers superior out-of-sample

hedging performance for climate change risk compared to hedge portfolios constructed from equities or ETFs.

The third paper investigates the accuracy of target price forecasts made by sell-side analysts, employing machine learning approaches to predict the forecasts' accuracy. Using a dataset of target price forecasts for U.S. listed companies from 1999 to 2021, ensemble methods incorporating market-level, firm-level, and analyst-level information are used to predict target price accuracy in terms of errors and achievement. A long-short portfolio constructed based on these predictions significantly outperforms the benchmark in terms of cumulative return and Sharpe ratio.

Acknowledgements

I would like to express my deepest gratitude to my supervisory professor, Dr. Ronald Balvers, for his invaluable guidance, insightful feedback, and unwavering support throughout my research journey. His expertise and encouragement have been instrumental in the completion of this thesis. I am also grateful to Dr. Narat Charupat and Dr. Richard Deaves for their constructive comments.

I thank my family for their unconditional love, patience, and unwavering support. Your belief in me has been a constant source of strength and encouragement. To my beloved children, your presence in my life brings me immense joy and inspiration. I am profoundly thankful for the support and love of my husband, Yugang, who has been my rock throughout this journey. Your understanding, sacrifices, and encouragement have made this achievement possible. Finally, I would like to thank my friends and colleagues, Laite, Rongzhao and Michael for their encouragement and support. This accomplishment would not have been possible without the collective support of everyone in my life.

Thank you all from the bottom of my heart.

Contents

List of Figures

List of Tables

Declaration of Authorship

I, Qiao Wang, declare that this thesis, entitled, "THREE ESSAYS ON FINANCIAL MARKETS", and the work presented in it, are my own or to which I have made significant contribution. I confirm that the thesis comprises the following chapters:

- Determinants and Predictability of Commodity Producer Returns
- Carbon pricing and the Commodity Risk Premium
- Are analysts' Forecasts Reliable? A Machine Learning-Based Analysis of the Target Price Accuracy

The first chapter is prepared under the close supervision of my Ph.D. supervisor, Dr. Ronald Balvers, the second chapter is entirely on my own original work, and the third chapter is a collaboration with Rongzhao Ou (McMaster University). This thesis is entirely my own or me and my coauthors' original work unless otherwise indicated. Any use of the work of other authors is acknowledged at their point of use.

Introduction

The financial markets are dynamic and multifaceted, continuously evolving in response to global economic shifts, technological advancements, and emerging risks. This thesis delves into three areas of finance: asset pricing for commodity producers, carbon pricing risk in commodity markets, and the accuracy of the analysts' forecasts. Each of these topics is critical for understanding the complex landscape of financial markets.

Commodity markets play a pivotal role in the global economy, influencing everything from raw material supply chains to investment strategies. Traditional asset pricing models fall short in capturing the unique risks and returns associated with commodity-producing firms. The first chapter in this thesis addresses this gap by adapting the investment-based asset pricing approach of Cochrane (1991) to determine the stock returns for producers of non-renewable resources. This approach allows prediction of returns without specifying systematic risk factors, leveraging well-known principles of optimal exhaustible resource extraction, the original Hotelling Rules (Hotelling (1931)). By developing a modified Hotelling Valuation Principle (HVP), we identify how the firms' expected stock returns relate to commodity prices, commodity reserves, production costs, and other factors. Unlike traditional evaluations of firm value which often overestimate reserves, our approach focuses on changes in levels, captured by stock market returns, thus avoiding biases related to valuation levels. The empirical results confirm the qualitative predictions of the modified HVP for the stock returns of commodity producers. Sorting firms into quintiles based on expected stock returns yields substantial annualized returns, suggesting that unidentified systematic factors are central for commodity-producing firms.

The second chapter examines carbon pricing risk in commodity markets. Using carbon emission allowance prices as a proxy, we investigate whether carbon pricing risk carries a positive risk premium in commodity markets. Our findings confirm that carbon pricing risk is priced, with a significant positive risk premium. Furthermore, this study demonstrates that commodity-specific characteristics, such as basis and hedging pressure, influence sensitivities to carbon pricing risk. By constructing a climate change hedge portfolio based on carbon pricing beta, we highlight the superior hedging performance of commodities compared to the benchmark portfolio of equities. This research extends the scope of climate finance, offering new insights into managing climate change risks in commodity markets.

In the third chapter, we focus on the accuracy of target price forecasts made by sell-side analysts. Utilizing machine learning approaches, we analyze a dataset of target price forecasts for U.S. listed companies from 1999 to 2021. Our analysis reveals substantial forecast errors and the significant impact of market conditions on target price achievement. By employing ensemble methods of machine learning, we predict target price accuracy and identify key drivers of the target price accuracy such as implied returns, market sentiment, and volatility. The predictive models developed in this study outperform benchmarks, demonstrating the potential of machine learning in enhancing investment strategies and decision-making processes.

Chapter 1

Determinants and Predictability of Commodity Producer Returns

1.1 Introduction

We adapt the investment-based asset pricing approach of Cochrane 1991 to determine the stock returns for producers of non-renewable commodities. The approach allows prediction of returns and calculation of the costs of capital for these firms without the need to specify systematic risk factors[1]. The focus on nonrenewable commodity producers facilitates the investment-based approach because the investment returns of these firms are driven by well-known principles of optimal exhaustible resource extraction. In-ground commodity reserves may be viewed as inputs to produce commodities, allowing direct calculation of investment returns.

[1]Specifying the correct factors affecting average returns is problematic. Even the mainstay models of Hou et al. 2015 and Fama and French 2015 struggle to explain the return variation within portfolios of test assets and leave numerous anomalies unexplained (Linnainmaa and Roberts 2018; Jacobs and Müller 2020.

The investment-based approach then identifies how the firms' expected stock returns relate to investment returns and thus commodity reserves and other inputs. [23] The literature on optimal exhaustible resource extraction provides guidance for applying the investment-based approach. By the original Hotelling Rule (Hotelling 1931) firms adjust resource extraction until the expected increase in the commodity spot price equals the risk-free rate. Building on previous modifications of the Hotelling analysis we determine the commodity-producing firm's output more generally where the marginal net profit from extraction of the commodity equals the firm's required return. Identification of the required return on equity follows from the determinants of the firm's net profit margin.

Miller and Upton 1985 recast the Hotelling Rule into what they call the Hotelling Valuation Principle (HVP): given intertemporal profit maximization, the market equity of the firm as a fraction of its total reserves must equal the current commodity price net of the marginal extraction cost. Computing the present value of the firm, in principle, entails evaluating the future revenues from all reserves at the prices and profit margins prevalent at all future times when the reserves are

[2] The investment-based approach has featured, apart from tangible capital, different inputs as determinants of stock returns. For instance, Lin 2012 finds that R&D investment predicts stock returns; Da et al. 2017 show that electricity usage predicts stock returns; and Belo et al. 2023 demonstrate that labor hiring predicts stock returns. Our approach adds commodity reserves to the input list.

[3] Two previous papers have related stock returns of commodity producers to production variables. Yang 2013 presents a production-based asset pricing model. Commodity producers are viewed as regular firms who manufacture non-storable consumption goods. Yang's approach does not treat commodity reserves as inputs and cannot treat commodities as nonrenewable. The approach also does not relate stock returns to investment returns and, accordingly, requires (exogenous) specification of a stochastic discount factor. Chen (2016) investigates the links between the stock returns of (an index of) commodity-influenced producers and commodity price increases. But examines the effect of exogenous stock returns on commodity prices, instead of the effect of exogenous commodity prices on stock returns. Chen does not consider additional determinants and works outside the production-based approach by taking the expected stock return as exogenous and constant.

sold. By the HVP, however, the only relevant price (and marginal extraction cost) to consider for valuing the entire reserve stock is today's price.

The HVP takes commodities prices as given and checks if firm valuation is consistent with the HVP. Subsequent research, however, concludes that it typically overestimates the value of proven and probable reserves by as much as fifty percent. Our approach looks at the HVP from a different angle. Instead of evaluating the valuation levels of the commodity-producing firms we examine the change in levels, captured by stock market returns. The difference is that omitted variables related to, for instance, market power, taxation, and real options may impact the level of firm value, but do not interfere with the intertemporal equalization of discounted profit margins if these variables are stable over time. From this perspective, the HVP will fail in level terms, but remain useful in difference terms, which we target by focusing on stock returns.

As a secondary motivation, consider that the commodityproducing firms represent a market segment that has interesting potential as a component in investment portfolios. The returns of these firms are, on the one hand, strongly tied to the underlying commodities.[4] On the other hand, the underlying commodity prices are weakly or even negatively correlated with stock market returns in general.[5] The question is if the commodity-producing firms act more like the commodities or more like the general stock market, and how the connections change with aggregate fluctuations. These issues are of practical importance because they determine the effectiveness of these firms as part of diversified investment portfolios or their

[4]For instance, Tufano 1998, Baur 2014, Zhang 2015, and Dar et al. 2019.

[5]See Pindyck and Rotemberg 1990,Zapata et al. 2012, and Daskalaki et al. 2014. Commodity prices also forecast general stock returns. See Huang and Kilic 2019

value as hedging instruments that may be more liquid than real commodities.

To preview the results, we find that the qualitative predictions of the modified Hotelling model directed to the explanation of stock returns are confirmed entirely for industrial metals and for energy resources, and partially for precious metals. The empirical results pertain to explaining the stock returns of the firms, but also to predicting their returns. We find that sorting the exhaustible commodity-producing firms into quintiles based on their expected stock returns, using only past information, and holding the top quintile while shorting the bottom quintile, yields an annualized return of around 16% for sorting based on prior pooled estimation and above 20% for sorting based on prior firm fixed effects estimation. These returns are hardly diminished by adjusting for standard risk factors, or for a commodity-specific risk factor. The fact that risk-adjusted returns are so large suggests that other, unidentified, systematic factors are central for commodityproducing firms. Unless the unidentified factors that are priced for the commodity-producing firms are non-systematic (i.e., only relevant for pricing commodities), unidentified systematic factors must be central for financial assets in general.[6]

1.2 Theoretical development

1.2.1 Background

Our theoretical development builds on two separate literatures: the Hotelling approach for optimal exhaustible resource extraction and investment-based asset

[6] Recent literature has identified factors which explain co-movement in commodities prices. Bakshi et al. 2019, Boons and Prado 2019, and Szymanowska et al. 2014 find common factors related to momentum, basis, and basismomentum. These, however, do not explain average stock returns.

pricing. The latter allows the appropriate discount rate (cost of capital) to replace the risk-free rate usually employed in the Hotelling framework.

By the original Hotelling Rule (Hotelling 1931) firms adjust resource extraction until the expected increase in the commodity spot price equals the risk-free rate. This implies that commodity prices should increase monotonically over time. Various modifications of the Hotelling analysis, however, determine the commodityproducing firm's output more generally with different implications.

Whereas Hotelling did not model extraction costs,Miller and Upton 1985 discuss the extraction cost as depending positively on current output and negatively on the level of reserves.[7] The motivation of the latter is the Ricardian principle of mining the cheaper resources first. Hence, as reserves decrease, the marginal extraction costs increase. In deriving the HVP they set the firm's discount rate equal to the real interest rate, tacitly assuming the firm does not face systematic risk.

Slade and Thille 1997 extend the theoretical contribution of Gaudet and Khadr 1991 to move beyond viewing the discount rate as risk free and employ an arbitrage argument to establish the appropriate discount rate as the CAPM-based required return. Although this contribution allows firm-specific discount rates that account for systematic risk, the discount rates are determined exogenously and are constant over time.

[7] Slade 1982 had previously formally incorporated the dependence of costs on the reserve level, as well as modeling technological progress in extraction. This modified the Hotelling Rule, allowing commodities prices in equilibrium to follow a U-shaped pattern increase over time.

Further refinements of the Hoteling approach include technological progress as considered by Lin and Wagner 2007, building on Slade 1982. Ellis and Halvorsen 2002 incorporate the impact of market power in commodities markets. Slade 1984 considers government regulation and taxation affecting the value and decisions of mining firms. Falls and Wilson n.d. argue that the reserve values are lower than predicted by the HVP due to the existence of real options. Cairns and Van Quyen 1998tackle the additional investment decisions of mining firms for exploring additional reserves.

Anderson et al. 2018 for the oil industry distinguish production from current wells and investment in drilling additional wells. They view the maximum output from a well as constrained by a fraction of total reserves in the well. Thus, increased output may arise from the intensive margin (more production from current wells, unless the constraint binds) and the extensive margin (drilling more wells). The Hotelling Rule is then amended to entail that the marginal return to drilling must rise at the discount rate.

We present an equilibrium model for exhaustible resource companies (mining firms) to derive the firm-specific elements determining their stock returns. The model contributes to the Hotelling (1931) setting by allowing for endogenous time-varying discount rates using the production-based asset pricing approach. [8] Utilizing the investment-based version of Cochrane 1991, Restoy and Rockinger 1994, and Liu et al. 2009, we derive how stock returns for the commodity-producing companies are related to their investment returns determined in the Hotelling framework.

[8]Established by Cox et al. 1985, Brock 1982, Berk et al. 1999, Balvers et al. 1990, Cochrane 1991,Cochrane 1996, and Zhang 2005

The mechanism by which investment returns become related to stock returns is internal arbitrage by managers at the individual firm level. The firm takes both commodity prices and financial market prices as given. The expected stock return (cost of equity capital) viewed by management is revealed in the expected investment return which we can predict in advance for the individual (mining) companies. The link between stock returns and investment returns holds irrespective of what the risk factors are or how firm value loads on the risk factors. Accordingly, the resulting expected returns may be determined without knowing what the systematic risk factors are. Moreover, the investment and production decisions of firms that we focus on are guided by relatively transparent profitability considerations in place of the investments decisions of households guided by the more opaque utility considerations of the consumption-based asset pricing approach (See e.g., Lin and Zhang 2013, and Zhang 2017).

1.2.2 The model

Consider a firm producing a nonrenewable commodity. The firm is competitive and takes the market price at time t of its resource (commodity) q_t as given. As in the Hotelling model, we assume the reserve level is finite and known, so the extraction ("production") quantity at time t, y_t, is equal to the difference in the reserve level x_t between two consecutive time periods, i.e., $y_t = x_t - x_{t+1}$.

The cost-of-extraction function, $c(y_t, x_t)$ is assumed to be homogeneous of degree one in the production level and the reserve level, strictly convex increasing in production, $c_y(y_t, x_t) > 0$, $c_{yy}(y_t, x_t) > 0$ (single and double non-time subscripts indicate first and second partial derivatives, respectively), and decreasing in the reserve level, $c_x(y_t, x_t) < 0$. Extraction is costly, and marginal costs are increasing in the extraction amount. Extraction also is relatively easier and cheaper when there is a larger quantity of reserves. If both production and reserves increase by a particular percentage then the mining costs increase by the same percentage, which is implied by the homogeneity assumption.

The assumption that production costs be homogeneous of degree one is not commonly made in the literature on optimal resource extraction. However, it is necessary to apply the investment-based asset pricing approach. The assumption superimposes the reasonable requirement that the average production costs increase monotonically in the ratio of production to reserves, y_t/x_t. However, it rules out "well effects" as we discuss later. For later use we state here that, by the Euler homogeneous function theorem (see e.g., Varian 1992 pp. 481-482), $c(y_t, x_t) = yc_y(y_t, x_t) + xc_x(y_t, x_t)$, and that the first partial derivatives are homogeneous of degree zero: $c_y(y_t, x_t) = c_y(y_t/x_t, 1) > 0$ with, furthermore, $dc_y(y_t/x_t, 1)/d(y_t/x_t) > 0$, and $c_x(y_t, x_t) = c_x(y_t/x_t, 1) < 0$, while no restriction is imposed on the sign of $dc_x(y_t/x_t, 1)/d(y_t/x_t)$.

The commodity-producing company is assumed to issue riskless debt, b_t, which is renewed each period. All operating profits, net of the interest on the debt and the revenue from additional debt issuance, are disbursed to the shareholders as dividends, which accordingly equal:

$$d_t = q_t y_t - c(y_t, x_t) - r_{t-1} b_t + (b_{t+1} - b_t) \tag{1}$$

where r_{t-1} is the risk-free rate at time t (pre-determined at time $t-1$). The cum-dividend market value of the mining firm to its shareholders is given as

$$V(s_t) = E_t\left(\sum_{j=0}^{\infty} m_{t+j} d_{t+j}\right) \tag{2}$$

where (with some abuse of notation) the cumulative stochastic discount factor between time t and time $t+j$ is given by m_{t+j} which is determined at the aggregate level and, for expositional simplicity is considered exogenous here as in Berk et al. (1986), even though its value and determination do not affect the ultimate solution for expected returns; s_t indicates a set of state variables at time t. The commodity-producing firms are price takers in financial as well as commodity markets.

The firm's optimal extraction decision problem is expressed by the Bellman equation:

$$\begin{aligned} V(s_t) = \mathrm{Max}_{y_t, b_{t+1}} \{ & q_t y_t - c(y_t, x_t) + b_{t+1} - (1 + r_{t-1}) b_t \\ & + E_t [m_{t+1} V(s_{t+1})] \}. \end{aligned} \tag{3}$$

Here $s_t = \{x_t, b_t, M_t\}$, with the first two firm-specific state variables and M_t reflecting any number of macro state variables (including q_t and parameters affecting the distribution of the stochastic discount factor and future commodity prices). The firm-specific equations of motion are

$$x_{t+1} = x_t - y_t, q_{t+1} = h(q_t, \varepsilon_{t+1}). \tag{4}$$

The second equation indicates that the commodity prices follow a stochastic process exogenous to the firm and with ε_{t+1} a random variable with distribution parameters included in M_t. The implicit function theorem implies that the state variables s_t in the Bellman equation are the non-choice variables (pre-determined or exogenous) in the decision problem of Eqs. (3) and (4): the reserve level x_t, firm debt b_t, and the commodity price q_t. Any additional parameters determining the distribution of the stochastic discount factor m_{t+1} and future commodity prices ε_{t+1} also impact the state (together with q_t captured by M_t).

Along the lines of Restoy and Rockinger 1994 the Appendix uses the first-order conditions, properties of the homogeneous cost function, and other constraints of the model in Eq. (1) thru (4) to obtain

$$E_t\left[m_{t+1}\left(\frac{q_{t+1}-c_y\left(y_{t+1},x_{t+1}\right)-c_x\left(y_{t+1},x_{t+1}\right)}{q_t-c_y\left(y_t,x_t\right)}\right)\right]=1. \tag{5}$$

The term in parentheses may be interpreted (following Cochrane 1991) as the gross investment returns of the firm - the marginal return to leaving an extra unit of the commodity in the ground instead of mining it:

$$1+r_{t+1}^I=\frac{q_{t+1}-c_y\left(y_{t+1},x_{t+1}\right)-c_x\left(y_{t+1},x_{t+1}\right)}{q_t-c_y\left(y_t,x_t\right)}. \tag{6}$$

The interpretation is that the denominator represents the marginal cost of investing (leaving an extra unit in the ground), equal to the opportunity cost of forgoing the margin, $q_t-c_y(t)$. (Note that here and subsequently we use the single function argument " t " as short-hand notation for all function argument values at time t). The numerator represents the marginal benefit (discounted by

m_{t+1}) of investing (extracting in the next period), $q_{t+1} - c_y(t+1) - c_x(t+1)$: the revenue in the next period net of the next-period marginal extraction cost, $q_{t+1} - c_y(t+1)$, which is mitigated, compared to what it otherwise would have been in equilibrium, by $-c_x(t+1) > 0$ because keeping more available reserves lowers the cost of extraction.

Eq. (6) represents a modified Hotelling Rule: the commodity spot price q_t, when adjusted for marginal production costs $c_y(t)$ and the marginal cost impact of resource depletion $c_x(t)$, grows at the investment hurdle rate r_t^I. This would occur in general equilibrium if all firms were similar. The implication of the Hotelling Rule then is that spot prices rise over time (especially in Hotelling's original formulation when marginal production costs are ignored) which is easily refuted by the observation that spot prices of most commodities have not monotonically increased over time. [9] Our focus is on the other direction, in which we take as given a stochastic path for commodity prices and use that to explain stock returns.

Miller and Upton 1985 introduced the idea of applying the Hotelling Rule in reverse by using the optimal intertemporal production decisions to value a commodity-producing firm by what they call the Hotelling Valuation Principle (HVP). For our model, the Appendix derives a solution for the (ex-dividend) value of the firm:

$$p_t + b_{t+1} = \left[q_t - c_y\left(y_t/x_t\right)\right] x_{t+1}. \tag{7}$$

The value of the firm, the market value of equity p_t (the number of outstanding

[9] See for instance Livernois 2009 and Schwerhoff, Stuermer, et al. 2019. The empirical evidence further is mixed about the performance of the various augmented versions of the Hotelling Rule (Livernois 2009; Slade and Thille 2009

shares is normalized to one) and debt b_{t+1} together, in equilibrium equal the total end-of-period reserves x_{t+1} times the current marginal unit profit margin $q_t - c_y(y_t/x_t)$, which is the commodity spot price net of the marginal cost of producing (extracting) the commodity. Miller and Upton (1985) use Eq. (7) (though with constant marginal costs) to test the HVP. They neatly confirm the HVP by finding empirically that a linear regression of $(p_t + b_{t+1})/x_{t+1}$ on q_t generates a slope close to one.

Subsequent research, however, shows less support for the HVP. Adelman 1993 concludes that the predicted reserves are only about half of the measured reserves: linear regression of $(p_t + b_{t+1})/x_{t+1}$ on q_t generates a slope of barely above one-half.[10] The reason may be limitations of the model (discussed at the end of this section) but may also relate to the accounting method for reserves - how it incorporates probable and possible reserves and the potential for developing and discovering reserves, and whether it has a conservative bias. By focusing on differences in the market value across firms or over time, as measured by stock returns, we move away from assessing the level of the market value of the firm. If accounting biases are stable, focus on returns instead of prices will avoid the bias.

Eq. (7) holds based on a cost function like that used in Slade and Thille 1997, but with an additional restriction of linear homogeneity imposed to be able to apply the investmentbased asset pricing approach. Accordingly and notably, Eq. (7) is derived taking into consideration the relevant cost of capital of the firm. The contributions of Gaudet and Khadr 1991 and Slade and Thille 1997 allow for the important addition of risk and riskbased discounting in the Hotelling framework

[10] See also Adelman 1990, McDonald 1994,Cairns and Davis 2001, and Falls and Wilson n.d.

but they take the firm's discount rate as constant over time. Slade and Thille apply the CAPM to obtain empirical estimates of the firm's discount rate and then confirm that the spot price dynamics of the commodity are consistent with the Hotelling Rule. Our intent is the reverse. Considering the spot price dynamics of the commodity and firm production decisions we explain and predict the discount rate, i.e., the stock return.

The previous literature has not considered endogenous firmlevel returns in this context. Employing the definition of the gross market return on the firm's equity, $1 + r^S_{t+1} = (p_{t+1} + d_{t+1})/p_t$, we can derive directly from Eqs. (1) to (7), and given that $c(y_t, x_t) = c_y(y_t/x_t)\, y_t + c_x(y_t/x_t)\, x_t$ for a homogeneous cost function:

$$r^s_{t+1} - r_t = \frac{[q_{t+1} - (1 + r_t)\, q_t] + [(1 + r_t)\, c_y(y_t/x_t) - c_y(y_{t+1}/x_{t+1}) - c_x(y_{t+1}/x_{t+1})]}{q_t - c_y(y_t/x_t) - (b_{t+1}/x_{t+1})}. \tag{8}$$

The excess return expression is best understood with reference to the investment return. Given the assumption that the cost function is homogeneous of degree one, it is known from Hayashi 1982 that average and marginal returns are equal. Therefore, the (marginal) investment return is equal to the overall (average) return on assets.[11]

[11] The excess stock return expression in Eq. (8) applies to firms processing exhaustible inputs and may be compared to the excess returns for the "normal" firms covered by Cochrane 1991, Eq. (15)). These depend on investment-to-capital ratios $I_{t+1}/K_{t+1}, I_t/K_t$, and a stochastic marginal product of capital. To determine the dynamics of the capital stock and satisfy second-order conditions, the derivation requires capital adjustment costs that are convex and homogeneous of degree one in I_t and K_t. In our exhaustible-resource case, identifying the investment return does not require specifying the endogenous process of capital accumulation because the input is already in place and is reduced by the quantity of the commodity processed. It would be useful to benchmark the explanatory power of our exhaustible commodities model to that for normal firms. However, the two formulations have no determinants in common and in previous empirical work the investment-based approach has not examined returns of individual firms but rather has focused on portfolio returns.

By optimization (internal arbitrage) the firm will continue to invest until the investment return equals the cost of capital. It follows that observing the investment return is equivalent to observing the cost of capital as management perceives it to be. Thus, the factors that affect the investment return also affect our assessment of the asset return. In equilibrium, the factors determining the investment return exactly identify the cost of capital. As first pointed out by Cochrane 1991 Cochrane 1996, the investment returns are equal to the stock returns. This is not exactly the case here because the firm need not be fully equity financed. We have investment return equal to return on assets, $r_t^I = r_t^A$, and the excess return on assets is equal to the equity fraction of firm value times the excess stock return, as in Liu et al. (2009): $r_t^A - r_{t-1} = (1 - \lambda_{t-1}) \left(r_t^S - r_{t-1} \right)$, where $\lambda_{t-1} \equiv b_{t+1} / \left(p_t + b_{t+1} \right)$. [12] This follows directly from Eqs. (6) to (8). Hence, the excess stock return in Eq. (8) is simply the levered excess investment return. The assumption of a homogeneous cost function allows us, empirically, to replace the difficult-to-observe investment hurdle rate by a linear function of the observable stock return.[13]

The intuition for excess stock returns Eq. (8) is as follows. The first term in brackets (the basic Hotelling's Rule term) indicates that the impact on excess return is positive whenever $\left(q_{t+1} - q_t \right) / q_t > r_t$. All else equal, when commodity prices grow faster than the risk-free rate, then investment returns increase since the benefit of leaving the commodity buried until the next period rises. The second term in brackets indicates that strictly increasing marginal cost of extraction, $c_y(t+1) > c_y(t)$, negatively affects stock returns. It does, since, along the optimal

[12] Harris and Pringle 1985 and Cooper and Nyborg 2006 show that, if the firm continuously adjusts its capital structure to a fixed leverage ratio, this equation is the correct way to relate stock returns and equity returns, even if taxes are involved.

[13] This is analogous to a similar assumption that allows unobservable marginal Q to be replaced by observable average Q in Tobin's investment analysis Tobin 1969

production path, the future higher marginal production costs imply a lower investment return. In addition, the inexorable reduction over time of reserves implies a lower investment return as it raises the production costs, $-c_x(t+1) > 0$. The final term, the denominator, indicates the current benefit of producing instead of investing. If it is lower, the investment return is higher. Alternatively, this denominator term also equals, from Eq. (7), the (scaled) equity value, which is associated with more leverage and so higher returns if the term is lower.

The explanation for the stock returns associated with the production-based asset pricing approach differs from that of the consumption-based approach which requires discussion of risk premia and exposure to the risk factors. In principle the results should be consistent. The stock return differences across firms should be related to differences in loadings on the risk factors that we did not need to specify in our production-based approach. Higher stock returns in Eq. (8), because of lower anticipated production growth (lower marginal costs) or higher commodity price growth (increasing revenue), means higher investment returns so that current production is lower and more of the commodity is left in the ground. From the alternative consumption based perspective the increased reserves and uncertainty about their future value imply the firm is more exposed to just about any systematic risk, requiring higher stock returns. In view of the difficulty in the literature to identify common systematic risk factors, a major advantage of our approach is that we do not need to identify these factors.

Our model specification has some limitations. The model ignores taxes and regulation (see Slade 1984), monopoly power in commodity markets (see Ellis and Halvorsen 2002), and real options (see Falls and Wilson n.d.). More significantly investment in exploration is missing from the model (see Cairns and Van Quyen 1998). The model also does not allow for re-evaluation of economic accessibility of reserves when commodity prices or mining technology change. As a result, reserve quantities may only decrease. However, 56% of the firms in our sample report at least one increase in the quantity of reserves over the sample period. To address this issue, we may extend the model based on Pindyck 1978, which considers exploration, to include the time varying discount rate. It would generate the same determinants of equity returns as our featured model plus an extra factor which is the expected marginal benefit of exploration per unit of exploration cost. The excess return expression in Eq. (8) omits this variable.

A further limitation of the model is related to the assumption of linearly homogeneous costs which we impose to apply the investment-based asset pricing approach. The assumption implies that unit production costs are a function of the ratio of production to reserves only and hence cannot easily account for issues such as changes in well pressure: In practice, especially for crude oil extraction, independently of the level of reserves, a high extraction speed raises the unit production costs disproportionately (or is simply impossible) at the well level as emphasized by Ahmed et al. 2024. This would generate lower investment returns and cause marginal returns to deviate from average returns, which our approach does not permit.

Model limitations may be responsible for firm-level market valuation not conforming to the HVP in empirical tests. Our theoretical derivation of equilibrium firm-level stock returns in the Hotelling framework allows the different perspective of focusing on changes in market valuation. Accordingly, relatively stable deviations from valuation due to the model limitations will cancel; or if value biases are not stable but unpredictable, they show up as random return shocks. On the other hand, commodity-producing firms have in common the tradeoff between current and future liquidation of their assets as an essential driver of profitability. We believe that the associated incentives are reasonably well captured by our modified Hotelling framework and constitute key determinants of stock returns.

1.3 Implications and empirical specification

The results from the model may be summarized in the following proposition:
Proposition. Dynamic maximization of profitability by a commodity-producing firm of Eq. (2) subject to Eqs. (1) and (4), and given a cost function $c(y,x)$ that is homogeneous of degree 1 in reserves x and output y, implies that the firm's expected excess stock returns are given by:

$$E_t\left(r_{t+1}^S\right) - r_t = \frac{(q_{t+1}/q_t) - (1+r_t) + \left\{(1+r_t)\, c_y\left(y_t/x_t\right) - E_t\left[c_y\left(y_{t+1}/x_{t+1}\right) + c_x\left(y_{t+1}/x_{t+1}\right)\right]\right\}/q_t}{p_t/q_t x_{t+1}} \tag{9}$$

Here subscripts represent either time t or a derivative with respect to the indicated function argument. Further, r is the risk-free return; r^S is the stock return and p the market value of the firm's stock; q is the commodity price.

Given auxiliary assumptions:

$$c\left(y_t, x_t\right) = (c/2) y_t^2/x_t, \quad \mathrm{Var}_t\left(y_{t+1}/x_{t+1}\right) = V \quad for \quad all \quad t, \tag{10}$$

six variables affect returns, with directions as follows:

$$\begin{array}{lcccccc} \text{Variables} & E_t q_{t+1} & q_t & r_t & \frac{E_t y_{t+1}}{x_{t+1}} & \frac{y_t}{x_t} & \frac{q_t x_{t+1}}{p_t} \\ E_t\left(r^S_{t+1} - r_t\right) & + & - & - & - & + & + \end{array} \tag{11}$$

Proof. Taking expectations in Eq. (8), using Eq. (7) to replace the denominator by p_t/x_{t+1}, and then dividing numerator and denominator by q_t, yields Eq. (9). We may work with $E_t y_{t+1}/x_{t+1}$ (instead of $E_t\left[c_y\left(y_{t+1}/x_{t+1}\right)\right]$) if the marginal cost function is assumed to be linear (i.e., costs are quadratic). So, for empirical purposes it is convenient to assume $c\left(y_t, x_t\right) = (c/2)y_t^2/x_t$ as in Eq. (10). This specific cost function satisfies the conditions we discussed, including homogeneity. Applying Eqs. (10) to (9) gives

$$\begin{aligned} E_t r^S_{t+1} - r_t = \Bigg(& \frac{E_t q_{t+1} - q_t}{q_t} - r_t - \frac{c}{q_t}\left[\frac{E_t y_{t+1}}{X_{t+1}} - \frac{1}{2}\left(\frac{E_t y_{t+1}}{x_{t+1}}\right)^2 \right. \\ & \left. -\frac{1}{2}\operatorname{Var}_t\left(\frac{y_{t+1}}{x_{t+1}}\right) - \left(1+r_t\right)\frac{y_t}{x_t}\right]\Bigg)\frac{q_t x_{t+1}}{p_t} \end{aligned} \tag{12}$$

Apart from the production variance, which we assume constant, $\operatorname{Var}_t\left(y_{t+1}/x_{t+1}\right) = V$, in Eq. (10) six variables affect returns. [14] The impact directions are indicated in (11).

The first two reflect opposite directions of the revenue effects from producing now compared to the next period. The second two relate to opposite directions regarding the marginal costs of producing now compared to the next period. The fifth variable represents the Hotelling value of reserves relative to the market value

[14]Since we can only work with annual data on the production and reserve variables it is not practical to allow for time variation in the conditional production variance.

of equity. Since $q_t x_{t+1}/p_t = [1/(1-\lambda_t)]\{1/[1-(c_y(t)/q_t)]\}$ from Eq. (7), a higher value indicates a combination of higher financial leverage λ_t (raising stock returns for given investment returns) and lower current profit margin $c_y(t)/q_t$ (raising investment returns) both implying higher stock returns. From Eq. (8), the impact of the final determinant, the risk-free rate, is to raise the opportunity cost of investing. The impact on the excess stock return is proportionate to $-[q_t - c_y(y_t/x_t)]$ which must be negative. Only the sign of the link between $\frac{E_t y_{t+1}}{x_{t+1}}$ on $E_t\left(r^S_{t+1}\right) - r_t$ is not mathematically obvious. Given the cost function of Eq. (10) the negative sign requires $\frac{E_t y_{t+1}}{x_{t+1}} < 1$ which holds as production cannot exceed total reserves.

To predict returns, i.e., generate $E_t\left(r^S_{t+1}\right)$, by Eq. (12) both $\frac{q_{t+1}}{q_t}$ and $\frac{y_{t+1}}{x_{t+1}}$ need to be forecast. First, growth in commodity prices is predictable but the anticipated growth rate may change over time. We obtain the forecast as the mean growth rate based on the 36 previous monthly observations. Although we do not derive this formulation from first principles, it allows the forecast of the growth rate to vary over time: [15] $E_t q_{t+1}/q_t = (1/S)\sum_{s=1}^{S} q_{t+1-s}/q_{t-s}, \quad S = 36$. Second, to forecast y_{t+1}/x_{t+1} we run a linear regression, using only data up to time t to forecast

[15]In an earlier draft we approximated the expected future spot rate by the futures rate, assuming a constant bias in the futures rate as a forecast of the future spot rate, but there are problems with this proxy. First, the futures rate and bias are endogenous and both stock returns and the futures rate may be affected by extraneous variables. E.g., investment may lower both stock returns and the commodity futures rate (as in David 2019, e.g., for the oil market). Second, the risk premia giving rise to the bias are not constant, as shown, for instance, by Szymanowska et al. 2014. Akin to this, as emphasized in Koijen et al. 2018, the futures premium (our earlier proxy) is closely related to the carry (or basis), whereas the expected commodity price appreciation (the variable we are seeking to approximate) is the complementary component of the expected commodity return that may not be highly correlated with the carry. Our empirical results when the futures rate proxies for the expected future spot rate are in Web Appendix Wang and Balvers 2021. Compared to our main results they are very similar except that the coefficient on the expected commodity price increase is quantitatively smaller. Just as for risk premia of stocks, there is considerable discussion about which systematic risk factors explain commodity price risk premia. See Bakshi et al. 2019, Beckmann et al. 2014, Boons and Prado 2019,Daskalaki et al. 2014,Ratti and Vespignani 2016, Szymanowska et al. 2014, and (Yang 2013, p.165).

y_{t+1}/x_{t+1} for time $t+1$. Our production forecast $E_t\left(y_{t+1}/x_{t+1}\right)$ is then used to forecast the stock return r^S_{t+1}, using exclusively past information. The forecast variables in this regression are the same as those used in Eqs. (9) and (12):

$$E_t y_{t+1}/x_{t+1} = f \quad [E_t q_{t+1} \quad q_t, \quad r_t, \quad y_t/x_t, \quad q_t x_{t+1}/p_t]. \tag{14}$$

In the following we estimate Eq. (12) using Eqs. (13) and (14).

1.4 Data

1.4.1 Sample selection

We use all mining firms in the Compustat Industry Specific Annual database with available production and reserve data. Compustat contains operational data for North American companies in particular industries including airlines, gaming, mining, oil and gas, etc. The mining industries covered include gold, diversified metals and mining, precious metals and minerals, and oil and gas. The products include gold, silver, copper, nickel, zinc, coal, metallurgical coal, iron ore, oil, natural gas and natural gas liquid. These are all exhaustible resources subject to the theoretical forces of the Hotelling analysis.

We exclude firms whose main mining products are iron ore, coal and met coal since they are not highly liquid commodities in North America.[16] Our sample consists of the North American firms for which COMPUSTAT has annual data on

[16]The lack of liquidity means that quoted commodity prices may not reflect fundamentals well at each time. These same commodities are also excluded in other recent asset pricing studies involving commodities. E.g.,Bakshi et al. 2019, Bianchi et al. 2016,Boons and Prado 2019, Daskalaki et al. 2014,Gorton et al. 2013,Koijen et al. 2018, and Szymanowska et al. 2014

reserves for the fiscal year ending in calendar year $t-1$ and have production data on one or more of the following commodities: gold and silver (precious metals), copper, nickel and zinc (industrial metals), WTI crude oil, and natural gas (energy fuels) for calendar year t. Because the industry-specific data from Compustat are available from 1999 forward, our return sample starts in July 2001 (since we use data from the previous full calendar year 2000 for reserves and production to explain returns in 2001) and ends in December 2018, containing a total of 52,337 firm months that meet our data criteria. Reserves for each firm are measured as the sum of their proven reserves and probable reserves (excluding possible reserves).[17]

Stock returns and market equity at the monthly frequency are also from Compustat for both the commodity-producing firms listed on Canadian and United States exchanges. They are computed from end-of-month closing prices adjusted for dividends and stock splits. For dual-listed firms we use the listing in the country of origin. Monthly real risk-free returns are measured by the US 3-month T-Bill rate available from the Board of Governors of the Federal Reserve System minus the realized CPI-based (urban consumers, seasonally adjusted) inflation rate from the Bureau of Labor Statistics.

Table 1.1 provides an overview of the available data. Panel A lists the number of firms producing each of the commodities by country (U.S., Canada, and Other,

[17]For the mineral mining commodities (gold, silver, copper, zinc, and nickel in our case) Compustat provides "proven and probable reserves" whereas for oil and gas it provides "total proved reserves". Both are similar: SEC reporting guidelines define proven and probable reserves as deposits that may be economically and legally extracted at the time of reserve determination, and total proved reserves as recoverable with reasonable certainty under existing economic and operating conditions (Securities and Exchange Commission, 2021, Sub parts 229.1200 and 229.1300). The criterion of economic viability implies that recorded reserve quantities may vary based on commodity prices, a possibility that our model ignores. For 58% of the firms in our sample, their reserves increase at some point during our sample period. The reason may be increased economic viability, or it may be the result of exploration.

where "Other" consists of any mining firms listed on North American exchanges but incorporated outside of the U.S. and Canada). In terms of the main activity of the companies, the sample is dominated by fuel energy firms, 693 in total, while there are 110 precious-metal producing firms and 28 industrial-metal producers (on average over time). Panel B shows the total number of included firms by year (2001-2018) and by country (US, Canada, and Other). The minimum number of included firms is 176 in 2001 and the maximum is 328 firms in 2014.

1.4.2 Predictor variable construction

To obtain the relevant commodity prices at the firm level, we collect monthly commodity data from Bloomberg for the seven commodities in our sample. These commodity data include spot prices for gold, silver, copper, nickel and zinc, and nearest-to-maturity futures data for the energy fuels, oil and gas. [18] As most of the firms in our sample produce multiple commodities, a firm-level commodity price is calculated as the production-weighted average value of the individual commodity prices. The weights for each commodity are calculated with production quantity and sales price data reported in year $t-1$ since the production data are available only on an annual basis.

Similarly, the variables $q_t x_{t+1}/p_t$ (current value of reserves as a fraction of market equity) and y_t/x_t (production as a fraction of total reserves) are obtained as

[18] Because spot prices are not available from exchanges for the energy fuels it is common to use nearest-to-maturity futures data instead. See, for instance, Litzenberger and Rabinowitz 1995 and Khan et al. 2017. To check that this convention does not distort our results relative to other commodities we also obtain the regression and sorting results when all commodity prices are replaced by nearest-to-maturity futures prices. These results deviate little from our main results and are presented in Web Appendix Wang and Balvers 2021

production-weighted average values since the reserve estimation data and production information are available for each company at the product level instead of at the consolidated company level. We use each firm's market equity at the end of December of year $t-1$ to compute $q_t x_{t+1}/p_t$. Note that x_{t+1} is determined at time t once the current production level is deducted from reserves. Although we abstract from exploration and revaluation of reserve quantities in the model, reappraisals of the value of reserves are included in our empirical measure.

1.4.3 Summary statistics

Table 1.2 provides descriptive statistics of the returns and predictor variables. The effective sample period is from July 2001 to December 2018. Panel A presents the summary statistics for the full sample, aggregated across the seven commodities. The average returns over the sample period are 0.59% a month with a standard deviation of 15.21%. The average monthly interest rate is 0.11% in this period. The average one-month-ahead commodity spot price growth rate is 0.87% with a standard deviation of 1.40%. The average production-to-reserves ratio over the sample period is 0.127 (i.e., current production is on average about 13% of the proven and probable reserves) with a standard deviation of 0.076 . The average value-of-reserves-to-market-equity ratio is 4.65 (which is the product of a leverage term and a reserves-to-firm-value term) with a standard deviation of 5.33. As an indicator of leverage, the average value of debt in ratio to the value of the firm is 0.34 . Panels B-D of Table 1.2 for the same variables provide the statistics separated into precious metals, fuel energy, and industrial metals.[19]

[19]The classification is like that of the Institute for Financial Markets. The three categories we use are the exhaustible resource categories of the seven used by Szymanowska et al. 2014.

1.5 Empirical Results

To estimate Eq. (12) for the forecasts of the stock returns of the mining firms we first obtain the independent variables. At time t the variables are the commodity spot price q_t, the quantity of production (as a fraction of prior reserves), y_t/x_t, the value of reserves per unit of stock market value, $q_t x_{t+1}/p_t$ (note that $x_{t+1} = x_t - y_t$ is known at time t), and our measure for the risk-free rate, r_t. In addition, known at time t, are the commodities price forecast, $q_{t+1} \equiv E_t\left(q_{t+1}\right)$, and the production level forecast, $y_{t+1}/x_{t+1} \equiv E_t\left(y_{t+1}/x_{t+1}\right)$.

1.5.1 The production-to-reserves forecast

Forecast the production level as a fraction of reserves by linearizing equation (14):

$$\frac{y_{t+1}}{x_{t+1}} = a_{0t} + a_{1t}\ln\left(q_{t+1}\right) + a_{2t}\ln\left(q_t\right) + a_{3t}\left(r_t\right) + a_{4t}\left(\frac{y_t}{x_t}\right) + a_{5t}\left(\frac{q_t x_{t+1}}{p_t}\right) \tag{15}$$

Notice that the coefficients change for each period because only past information is used to obtain predicted values at each time. The estimates of equation (15) by OLS at the annual frequency are shown for the full sample only in Table 3, regression (4). The estimation results are presented for both a pooled specification, in which all mining firms are treated equally irrespective of commodity produced or home country (US, Canada, or "Other"), and a panel specification allowing individual firm effects. Separation by commodity is hard because most firms in our sample produce multiple commodities jointly.

(Compared to Szymanowska et al. 2014 we drop the foodbased categories - Meats, Grains, Oilseeds, and Softs)

In both specifications lagged production is significant at the 1% level with a coefficient of 0.91 indicating a high level of persistence in production in the pooled case, and 0.46 in the panel case. The reserves-to-equity variable, acting essentially like a Tobin's Q variable in this context, is also significant at the 1% level, with a coefficient of -0.0016 in the pooling case and -0.0021 in the panel case. It reflects higher production in anticipation of profitable opportunities.

1.5.2 Stock return prediction regressions

Using only past information, stock returns for the following month can be predicted from equation (9). The production and reserve data are annual whereas the other variables - stock returns, commodity prices, and equity value - are at a monthly frequency. Annual production and reserve data from year $t-1$ are used to predict monthly returns from July of year t to June of year $t+1$ (the standard timing convention since Eugene and French 1992). Specifically, we consider the following prediction equation,

$$r_{t+1}^{S} - r_t = b_0 + b_1 \ln\left(q_{t+1}\right) + b_2 \ln\left(q_t\right) + b_3\left(r_t\right) + b_4\left(\frac{y_{t+1}}{x_{t+1}}\right) + b_5\left(\frac{y_t}{x_t}\right) + b_6\left(\frac{q_t x_{t+1}}{p_t}\right) \tag{16}$$

Only variables up to time t (including the right-hand side variables in equation 13) are used to obtain the coefficients to forecast (y_{t+1}/x_{t+1}). The predicted value $r_{t+1}^{S} - r_t$ is derived from information at time t or earlier. The approach may be employed in real time to forecast returns for the month ahead.

The predicted coefficient signs for equation (16) from the Proposition are $b_1 > 0, b_2 < 0$, $b_3 < 0, b_4 < 0, b_5 > 0$, and $b_6 > 0$. In addition, it follows from equation

(12) that $b_1 = -b_2$. This reflects the Hotelling effect: it is the relative price increase that relates to investment returns.

We also have $b_1 = 1$ as tested previously byMiller and Upton 1985, or in a more parsimonious regression that the coefficient on $\ln(q_{t+1}) - \ln(q_t) - r_t$ is equal to 1 . Furthermore, $b_5 \geq -b_4$. This reflects the marginal cost of producing currently versus the next period. By keeping the resource in the ground one period longer, the future costs are reduced as they now are based on a (marginally) larger reserve. If this latter effect is considered negligible, we expect $b_5 = -b_4$.

We check first for multicollinearity in estimating equation (16). In particular, $\frac{y_{t+1}}{x_{t+1}}$ may be highly correlated with $\frac{y_t}{x_t}$, and $\ln(q_{t+1})$ with $\ln(q_t)$. Note first that estimation of $\frac{y_{t+1}}{x_{t+1}}$ employs time-varying firm-specific information and that $\ln(q_{t+1})$ employs time series information of 36 periods. As shown in Panel C of Table 4 the Variance Inflation Factors (VIFs) for each variable are less than 10 (the typical cutoff value for absence of multicollinearity) except for $\ln(q_{t+1})$ and $\ln(q_t)$ which have VIFs of 76.34 and 77.01 , respectively. This implies that the standard errors of the coefficients for $\ln(q_{t+1})$ and $\ln(q_t)$ are, respectively, 8.74 and 8.78 times (square root of VIF) as large as without correlation between these variables themselves and the other regression variables. The coefficient estimates for these variables may be less reliable although this should not affect the overall forecast performance of the regression. To avoid this issue, as well as avoiding overfitting the forecast specification, we provide more parsimonious alternative specifications that combine variables under the assumption that the coefficient restrictions hold.

Panel A of Table 1.4 presents the stock return forecast results of a pooled

OLS regression. Regression 1 qualitatively confirms the simple Hotelling Rule by which the expected commodity price increase implies a higher cost of equity capital (mean stock return). The expected future spot price and the current spot price have a significant, quantitatively similar but opposite, impact around 0.57 on stock returns of the commodity's producer. Regression 2 shows the parsimonious regression with the combined variable $\ln(q_{t+1}) - \ln(q_t) - r_t$. The coefficient is also 0.57 , significantly positive ($t - stat = 11.51$). Panel C shows Chi-squared test results for the coefficient restrictions which cannot reject the restriction that $b_1 = -b_2$. However, the restriction $b_1 = 1$ is rejected statistically, and at 0.57 the coefficient is also economically less than 1 . In the parsimonious specification of regression 2 the coefficient of 0.57 is also significantly smaller than 1 .

Adding production and reserves, Regression 3 shows all coefficients as significant at the 1% level. They have the predicted signs, confirming all six predicted coefficient signs. The spot price effect is numerically again almost identical to the negative of the futures price impact and the chi-squared test in Panel C shows that the null hypothesis $b_1 = -b_2$ cannot be rejected statistically. But again $b_1 = 0.61$ is significantly less than 1 . The predicted future production and current production coefficients are almost identical in absolute value. We have that $b_5 > -b_4$ by a small amount, suggesting that the difference in current and future marginal costs, the resource exhaustion effect as represented by $c_x(y, x)$, is small. The chi-squared test in Panel C cannot reject the hypothesis that $b_5 = -b_4$, confirming our prediction.

Regression 4 in Table 1.4 parsimoniously applies $b_1 = -b_2$ and $b_5 = -b_4$ (if

$c_x(t)$ may be ignored). [20] We thus include only the excess expected future commodity price increase, the production growth rate, and the reserves-to-equity value measure. The results are quantitatively like Regression 3, including that the excess commodity price increase coefficient equals 0.61 and is significantly less than 1.[21]

Focusing on Regression 4 we present the economic importance of the coefficients for predicting stock returns. The price coefficients of around 0.61 in magnitude imply that a one standard deviation increase in the expected commodity price increase (1.40 in Table 1.2) raises monthly stock returns by about 85 basis points. The production coefficient of around -9.2 implies that a one standard deviation boost in output as a fraction of reserves (0.076 in Table 1.2) lowers stock returns by about 70 basis points. Lastly, the reserves-to-equity value coefficient of 0.095 implies that a one standard deviation increase in the value of reserves as a fraction of market equity (5.33 in Table 1.2) would raise the future stock return by about 51 basis points. Accordingly, each of the three variables in the parsimonious regression is economically important in affecting stock returns.

[20] More precisely if $c_x(t) + r_t c_y(t)$ may be ignored, where r is the monthly risk-free rate so very small and $c_x(t)$ and $c_y(t)$ are of opposite sign, with $c_x(t)$ presumably of much smaller magnitude than $c_y(t)$

[21] The R-squares for all pooled regressions in Table 1.4 are from $0.25\% - 0.43\%$ which appears low. However, they are for forecasts based on past variables, for monthly returns, and at the firm level. Predictability of stock portfolio returns at monthly horizons is very low (Fama and French 1988), no matter what variables are used. The reason is that expected returns vary little compared to realized returns at monthly frequencies. Individual firms exhibit even more random variability in returns than portfolios of firms. Zhang (2005, Table 5) finds that zero-investment portfolio returns of high-value stocks and shorting low-value stocks are explained by various portfolio averages of firm characteristics, with the R-squared ranging from 0.16% to 0.71%. Belo et al. (2020) generate much higher R-squares such as around 25% but this is in-sample for the market portfolio and a 5-year horizon. For a 1-year horizon out-of-sample, Belo et al.'s R-squared is only 0.42% for the market portfolio. In comparison, the Rsquared of 0.41% for our main specification is quite high since it is out-of-sample for individual firms and a 1month horizon. We also argue later in the paper that the economic significance of the 0.41%R-squared is sizeable.

The pure Hotelling effect suggests a 1% higher expected return if the expected commodity price increase is 1% higher. The result in regression 4 , however, of 0.61% is significantly less than 1%. Although it has the predicted sign and is significantly different from zero, this is quantitatively lower than expected. A possible explanation is that the driving forces from the Hotelling model work but, in part, are moderated by elements omitted from the model, or caused by measurement error in the forecast variables generating downward biased estimates. For instance, reduced expected commodity price growth may have a diminished impact on the stock return compared to what the model predicts if the firm may exercise its real option to shut down a mine. However, a more direct explanation may be that firms hedge their exposure to commodity price fluctuations. For instance, Acharya et al. 2013 find that, in the period from 2000 to 2010, 88% of fuel-producing firms hedged commodity price risk with derivatives. If firms, say, hedged 39% of their exposure to commodity price risk this would explain a reduction of the theoretical impact from one-to-one to the 0.61 -to-one we find empirically.[22]

Panel B of Table 1.4 presents the results for the panel specification with firm-specific fixed effects.Here we also use the panel results for the predicted production-to-reserves ratio, $\frac{y_{t+1}}{x_{t+1}}$, based on Table 1.3. Overall, the results are much like the pooling case in Panel A. The main difference is for the parsimonious case where the coefficient on excess expected commodity price growth is smaller, equal to 0.34 , again significantly below 1 , but significantly positive at the 1% level. Further, the coefficients on the spot price and expected future spot price, while

[22]Our model is consistent with hedging but has little to say regarding this issue: The degree of hedging commodity price risk by the firm in the model based on maximizing stockholder wealth is indeterminate because investors may always hedge this risk on their own account equally effectively.

again quantitatively similar, are now statistically different from each other.

1.5.3 Commodity differences

Subdividing the mining industry into three main categories - precious metals, industrial metals, and energy fuel (oil and gas) - we further test if the industry category in which the mining firm is classified influences its expected return and predictability. Following the industry category distribution in our sample, we create two industry dummy variables - PreciousMetals and IndustrialMetals. We assign the PreciousMetals variable (set to one) to any observation with combined production value weight of gold and silver of more than 0.5 , and similarly assign the IndustrialMetals dummy variable (set to one) to any firm with combined production value weight above 0.5 in zinc, copper, and nickel to represent the industrial metals category. The dummy coefficients may reflect differences in $\mathrm{Var}\left(y_{t+1}/x_{t+1}\right)$ across the industry categories, which we do not directly capture in our regressions.[23]

Because the production extraction processes and the nature of the commodities market may vary substantially between the three categories, their investment returns may have different relationships to the explanatory variables. To examine this possibility, we expand the category dummy variables to include interactions with the explanatory variables. We limit the specification to only the interactions with the variables in the parsimonious formulation in Table 1.4 (regression 4).

[23]The empirical result with these industry dummies added is a case in-between the pooled and panel results of Table 1.4, Panels A and B, and presented in Appendix. The dummy coefficient is positive significant, at the 1% level for precious metals companies, which means they have higher expected returns than the oil and gas companies, all else equal. Other coefficients are similar to the pooled results in Table 1.4, Panel A.

Table 1.5 (pooled) presents the results with the interaction dummies added to the parsimonious formulation from Panel A in Table 1.4. The interaction dummies capture the effects of structural differences in the cost functions related to variation in extraction technologies among the industry categories. The specification is equivalent to running separate regressions of the parsimonious formulation for the three commodity categories.

The excess expected commodity price growth coefficient increases to 0.78 and the expected production difference parameter more than doubles to -18.3 while the reserves-toequity parameter is roughly unchanged when we add only the interaction dummy variables. When we also add the level dummy variables, the excess expected commodity price growth coefficient becomes 0.85 , the expected production difference parameter is -17.1 , and the reserves-to-equity parameter raises to 0.085 . For both cases (regressions 2 and 3 in Table 5), given the six interaction dummy coefficients, the three for precious metals are significant and the three for industrial metals are not.[24]

Precious metals have a significant and quantitatively large decreased sensitivity to the excess expected commodity price growth, reversing the overall sign. Thus, precious metals' cost of equity capital is negatively related to excess expected commodity price growth which is counter to the simple Hotelling rule and inconsistent with our model. Possibly our simple moving average predictor is inadequate for gold and silver prices. These prices may be hard to predict. One sign

[24]The coefficient restriction tests and multicollinearity checks are consistent with the earlier results for the base case in Table 1.4 and are not tabulated. Specifically, even the expected commodity price growth coefficient of 0.85 is still significantly below 1 . Apart from significant multicollinearity between $\ln(q_{t+1})$ and $\ln(q_t)$ there is no indication of multicollinearity for any of the variables, including the interaction variables.

is that the gold and silver futures premium is not a good predictor of spot prices for gold and silver either. [25] The second interaction variable for precious metals is for the expected production growth difference, which is significantly positive and quantitatively large, and again reverses the negative sign predicted by our model. Conceivably for precious metals there are large inventories of the commodity (essentially any previously mined quantities of gold or silver) that compete with production. The final interaction variable for precious metals is negative. It indicates that for precious metals the impact of the reserves-to-equity ratio on stock returns is diminished, but in total still has the predicted negative sign. For industrial metals, the results are similar as for energy fuels, consistent with the model.

Table 1.5 (fixed effects) allows for firm-specific fixed effects, and accordingly omits the level dummy variables by industry category but retains the interaction dummies. The resulting specification produces a qualitatively similar outcome to the pooled specification; however, the excess expected commodity price growth coefficient is now a bit smaller at 0.61 which is consistent with the results found in the pooling cases without interaction variables.[26]

1.5.4 Prediction errors

To evaluate how well our approach predicts future returns we obtain for each firm in the sample at each month the error from the forecasted return of regression

[25] Chinn and Coibion 2014 find that futures prices for precious metals are poor predictors, whereas futures prices for fuel energy commodities are much better predictors of subsequent prices.

[26] Additional results, dealing with geographic differences of the mining firms, and results focusing on returns on assets instead of equity returns are available in Appendix.

equation (16) and compare it to the subsequent realized return one month later. We use the numbers from a recursive version (using only past information) of the parsimonious regression equation (4) shown in Table 1.4. We calculate the mean-squared forecast error (MSFE) by squaring the errors, averaging over all firms in the sample at each month, then taking the square root.

As a point of reference with other forecast methods, we also predict returns of each firm for one month ahead by their historical means. As demonstrated by Welch and Goyal (2008) and Cenesizoglu and Timmermann (2012) historical means generally forecast future mean returns better than typical forecast variables and statistical models, at least at the aggregate level. With the same approach as above but replacing the forecast from the parsimonious version of our model (using only past data) by the historical mean calculated over the same period (also using only past data). Then calculate the MSFE for the historical means as forecast. The results are displayed in Figure 1.1. We show the 12-month moving average of the monthly MSFE at each point as well as cumulated over time for the forecasts based on our modified Hotelling valuation (blue solid line) against the benchmark MSFE for the forecasts based on the historical means (green dashed line). The prediction errors for our model forecast are substantially lower than for the historical mean forecast (except for the end of the sample).

As an alternative benchmark we also compare the forecast errors for our modified Hotelling Valuation model to a version of the traditional Hotelling Valuation (based on Miller and Upton 1985) in which we omit the production and reserves variables and only use the expected spot rate increase net of the interest rate to forecast the excess stock returns. In Figure 1.1 the MSFE for the traditional

Hotelling Valuation (red dotted line) is worse than for our modified Hotelling Valuation but better than the historical mean-based forecasts.

1.5.5 Portfolio sorting

For a further indication of economic importance and to identify the influence of traditional risk factors we sort at each time all mining firms in our sample by predicted stock returns. [27] To forecast return for time $t+1$, we use the fitted value from the parsimonious version of equation (16), regression 4 in Panel A of Table 1.4, using the coefficients based on each prior observation up to time t along with the predictor variables at t, to sort firms into quintiles. We pool the first 24 time series data points across all firms (4485 data points) to estimate initial coefficient values, then roll forward using an expanding window. Quintile 1 in each month includes the observations (firms) with the 20% lowest predicted returns, and Quintile 5 in each month contains those with the 20% highest predicted returns. Subsequent monthly returns for each quintile are recorded and averaged over time.

For the full sample of 47,852 firm-month observations available for forecasting

[27] The results from equation (16) may call in doubt the model's ability to obtain economically meaningful returns from sorting into quintiles because the R-squared is only 0.0041 (0.41 percent). However, an argument of Cochrane 2001,(p. 447) allows us to estimate the sort of trading strategy returns that we may expect, even with a low R-squared. Since R-squared is the ratio of the explained return variance and total return variance, $R^2 = \sigma^2_{PRED}/\sigma^2_{RET}$ we find $\sigma_{PRED} = \sigma_{RET}\sqrt{R^2} = 15.22 * 0.064 = 0.975$ (numbers from Tables 1.2 and 1.4). Assuming normality, the highest quintile of predicted returns starts at a predicted mean return of $\mu + \sigma_{PRED}x$, with μ the unconditional average return and x given by the point at which the standard normal cumulative distribution is 0.8 : $F(x) = 0.8$ which implies $x = 0.84$. The average predicted return for the top quintile may be determined by using $E(x \mid x > 0.84) = f(0.84)/[1 - F(0.84)] = 0.28/0.20 = 1.40$, where $f(x)$ is the standard normal density. Then, for the predicted return distribution, holding the top and shorting the bottom quintile for one month generates $\mu + 1.40(0.975) - [\mu - 1.40(0.975)] = 2.73$ percent a month. The result is clearly economically significant despite the seemingly low R-squared but is an upper bound because the parameters used for prediction are likely measured with error.

the realized average returns by quintile are in Panel A of Table 1.6. As expected, for a (pseudo) out-of-sample test of the model, these returns increase monotonically (with a minor inversion at Quintiles 3 and 4) from Quintile 1, with the lowest predicted return, to Quintile 5, with the highest predicted return. The difference between the Quintile 5 mean return of 1.08 percent a month and the Quintile 1 mean return of -0.17 percent is 1.24 percent which is economically large and highly significant, with a t-statistic of 3.17 . The compounded annualized mean return difference amounts to an annual return of 16 percent.[28] The result confirms our theoretical prediction and supports the view that mining firms change their risk exposure significantly over time, creating predictable variation in costs of capital.

Theoretically, the return differences between quintiles should be systematic risk premia. The question is if these risk premia are compensation for known risk factors. To check we apply standard risk models to the return series to calculate the risk-adjusted returns. We consider the Fama-French three-factor model, the Fama-French five-factor model, and the Fama-French five-factor model plus the momentum risk factor of Carhart. We also consider the carry (or basis) factor of Yang (2013), Szymanowska et al. (2014), and Bakshi, Gao, and Rossi (2019) added to each of the models to account specifically for systematic commodity price risk.[29] We find that these risk factors generally explain only a small part of the returns. In fact, for the Fama-French three-factor model the alphas are higher than the raw returns (not shown). Table 1.6 presents the risk-adjusted returns based on the

[28]The 1.24 percent return is a reasonable fraction (0.45) of the 2.73 percent return expected (see previous footnote) assuming normal returns and no coefficient estimation error.

[29]We construct the carry factor with data from Bloomberg, using the procedure inBakshi et al. 2019, as the zero-investment return of holding out of 29 different commodity futures the five commodities that are most backwardated and shorting the five commodities that are most in contango.

five-factor model (Fama and French, 2015) and for the five-factor model together with the carry factor. In Panel A the difference in the alphas between Quintiles 5 and 1 is only marginally reduced to 1.09 percent monthly for the five-factor model (t-stat of 2.73) and 1.17 percent monthly for the five-factor plus carry risk model (t-stat of 2.95).

Panel B in Table 1.6 presents the results based on the panel approach with fixed effects for each firm in the sample. We use here the rolling version of regression 4 in Panel B of Table 1.4 to forecast the returns. The results are even clearer in this case: the returns increase perfectly monotonically from Quintile 1 with the lowest predicted return to Quintile 5 with the highest predicted return. The difference between the Quintile 5 mean return of 1.19 percent a month and the Quintile 1 mean return of -0.36 percent is 1.55 percent (t-statistic of 4.67) which implies an annualized return of more than 20 percent.

1.5.6 Benchmarks and industry differences

To provide a perspective for the sorting results we compare them to the return differences from alternative forecasts also using past data: (i) the forecast benchmark of the historical mean returns and (ii) a forecast based on "traditional Hotelling Valuation" using only the expected commodity price increase, both for the pooling and fixed effects, regressions 1 and 4 in Table 1.5. For (i) the difference between Quintile 5 and Quintile 1 is 0.34 percent a month, reduced to 0.19 percent when adjusted for risk. For the pooled case of (ii) the difference is 0.60 percent, reduced to $0.26 - 0.30$ percent after risk adjustment. For the fixed effects case the difference is 0.48 percent, reduced to $0.15 - 0.20$ percent after risk adjustment.

Substantially smaller than for the full model. Details are in Web Appendix of Wang and Balvers 2021.

Examining differences by commodity, we consider forecasts based on regressions 3 and 5 in Table 1.5 and regression 4 of Table 1.4. Accounting for commodity differences increases the sorting returns. Individually, fuel energy and precious metals generate significant return differences among the quintiles, while the differences are less for industrial metals and not statistically significant. See Web Appendix of Wang and Balvers 2021.

1.6 Conclusion

Our model implies and confirms empirically that the expected growth rate of the natural resource price has a positive effect on the expected returns and cost of capital of commodity producing firms. At the same time, firm attributes the reserves-value-to-market-equity ratio and the production-to-reserves ratio change have a positive and negative impact, respectively, on the expected return. The higher expected price of natural resources makes commodity producing firms retain more reserves underground to profit from the higher expected price which makes these firms more sensitive to any systematic risk factors. Firms with lower production growth imply more future risk sensitivity: with given reserves of a commodity producing firm, lower production growth today means higher future production which increases the sensitivity of a commodity-producing firm to future shocks.

The theoretical contribution of the paper is to use the investment-based asset pricing approach to derive results in the Hotelling framework for commodities firms that hold for any equilibrium discount rate. There is no need to specify the systematic risk factors and the discount rate may vary over time. Empirically, it is this theoretical refinement that allows us to test the Hotelling Valuation Principle (HVP) in a new way.

The HPV of Miller and Upton 1985 was initially confirmed but fared worse in subsequent empirical studies. Adelman 1993 finds that the HVP-predicted reserves are only half of measured actual reserves. Identifying the cause of this departure from the theory is difficult. It may be model limitations or measurement and accounting biases. By "differencing" the approach - considering stock returns instead of stock prices - we avoid biases related to the level of measured reserves. We furthermore provide additional implications not available from the level approach.

For North American firms producing precious metals, industrial metals, and fuel energy we find that average stock returns correspond to the HVP. A higher futures premium and higher reserves-to-equity ratios imply higher expected returns, and higher expected production growth rate lowers expected stock returns. The data confirm the predicted coefficient signs for all six of the variables identified by our modified Hotelling model. These restrictions also hold if we subdivide the firms by type of commodity, with the partial exception of precious metals for which we observe discrepancies compared to the model predictions. We assess the quantitative importance of the results, constructing portfolios by sorting the firms according to the predicted returns from our model. The sorting results show quite large average return differences consistent with the model predictions. The degree to which the results follow the predictions of the equilibrium Hotelling model suggests that the average return differences represent differences in equilibrium expected returns.

Theoretically, the source is time variation in the firms' choices of exposure to risk factors. The investment-based asset pricing approach did not need to specify the risk factors, but virtually no part of the average sorting return differences between fifth and first quintiles, amounting to 16 to 20 percent annually, can be explained by standard consumption-based risk factors. Inasmuch as all stocks are affected by the same systematic factors, it appears that one or more key risk factors are still missing in the consumption-based perspective.

The model quantitatively falls short in one respect. The impact of the expected spot growth rate on the excess stock returns is numerically smaller than predicted (about 61% of the predicted value). Although larger than the comparable number obtained from traditional tests of the HPV, around 50% (Adelman 1993), the number is significantly below the expected 100% in all our specifications. Possible reasons are the stylized nature of the model. Future work examining expected stock returns of commodity-producing firms may benefit from adding real options, conjoining the analyst forecasts for future commodity prices suggested by CoCortazar et al. 2019, and considering the impact of hedging strategies.

Hedging commodity prices is common for commodity producers (e.g., Acharya et al. 2013) and may easily, even in the context of our present model, explain the apparent incomplete reaction of stock returns to commodity price changes. If such hedging has become more common over time, it may even explain why the reaction was found to be complete in the earlier work of Miller and Upton 1985. Hedging may also explain the difference between the risk-premia in commodity prices and the risk-premia for the shares of commodity-producing firms. Our analysis points at the risk premium components related to production and reserve attributes of the producers, which are absent for holders of refined commodities, but additionally the stock returns in commodity producing firms contain the effect of hedging commodity price risk which is a component that is suitably excluded in measuring the risk premium for holding refined commodities.

Apart from confirming the HPV from a different perspective, our investment-based analysis provides vital additional insights. First, production factors together matter quantitatively more than the commodity price aspects of the traditional HPV alone. The implication for investment purposes is that holding the shares of commodity-producing firms is not a great substitute for investing in commodities. Second, we find that unfamiliar systematic risk factors, and time variation in firm exposure to these factors, are important determinants of the stock returns of commodity producers.

Future research could explore several extensions to enhance our understanding of the determinants of stock returns for commodity producers. One assumption in our model is that commodity producers issue debt at the risk-free rate. Examining variations in the cost of borrowing could provide clearer insights into how leverage influences the valuation of commodity producing firms. Additionally, while our model assumes that commodity producers will disburse all operating profits, net of interest on the debt and revenue from additional debt issuance to shareholders as dividends, it is worthwhile to study the relationship between dividend payouts and stock returns for commodity producers. Finally, considering market imperfections, such as the influence of large commodity producing companies on commodity market prices (e.g. the influence of large gold mining firms on gold price), would refine our model by accounting for strategic behavior and market power, offering a more comprehensive perspective on the determinants of commodity producer stock returns.

Figure 1.1: **Mean-squared forecast errors (MSFE).**

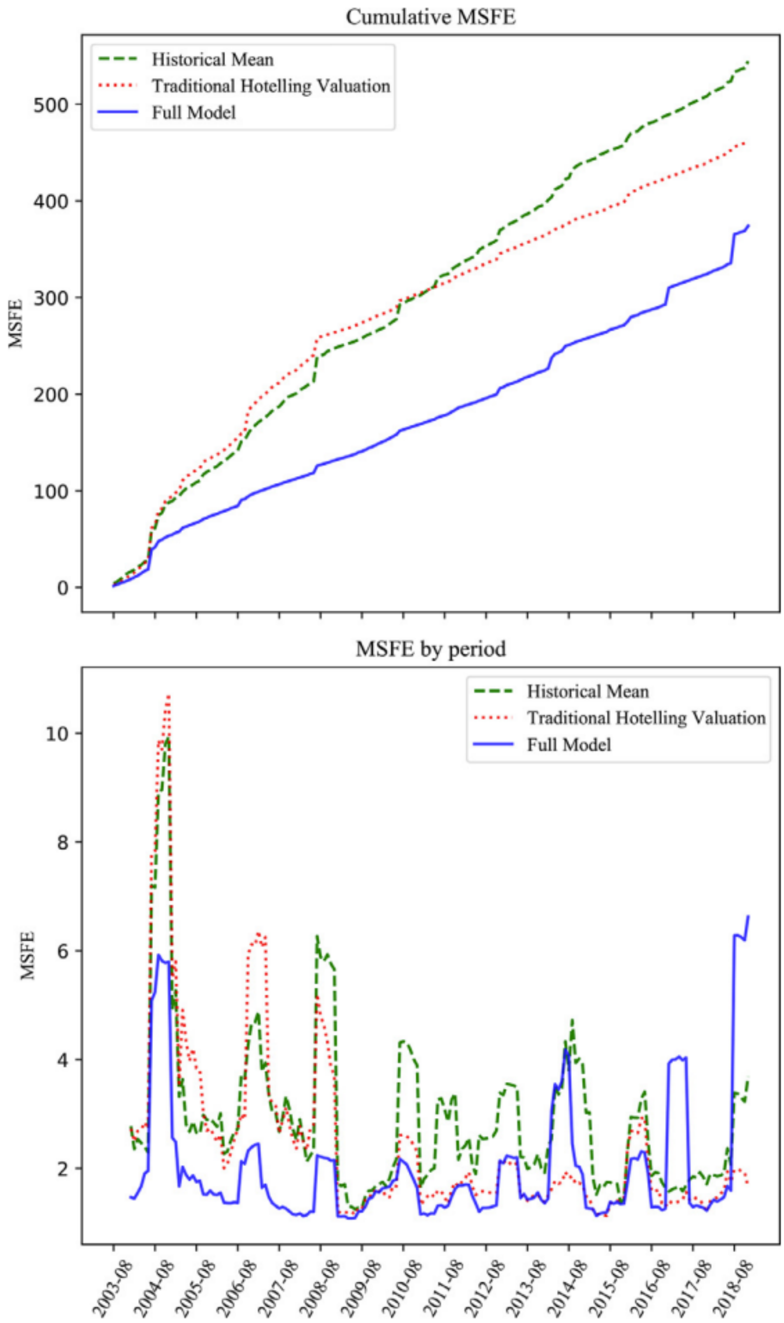

The figure displays the returns of the portfolios, the benchmark and the S&P500 Index from September 2006 to December 2018.

Table 1.1: Sample descriptive statistics by commodity.

We use all mining firms in the Compustat Industry Specific Annual database with available production and reserve data for the following commodities: gold, silver, copper, nickel, zinc, crude oil (WTI) and natural gas. Panel A provides the total number of firms in the sample for each of the commodities separated by country of incorporation. "Other" refers to mining companies incorporated outside of Canada or the U.S. but listed on a U.S. or Canadian stock exchange. Panel B provides the number of mining companies included in the sample aggregated across the commodities by year.

Panel A: Sample Size by Commodity				
Country Commodity	Canada	U.S.	Other	Total
Copper	16	3	1	20
Crude Oil	184	185	33	402
Gold	74	7	13	94
Natural Gas	138	146	7	291
Nickel	3	0	0	3
Silver	13	3	0	16
Zinc	5	0	0	5
Panel B:	Sample Size by	Year		
Country Year	Canada	U.S.	Other	Total
2001	62	104	10	176
2002	79	118	15	212
2003	97	121	18	236
2004	105	121	24	250
2005	114	124	26	264
2006	118	120	22	260
2007	132	120	24	276
2008	141	134	22	297
2009	138	143	23	304
2010	142	140	27	309
2011	132	131	28	291
2012	136	133	32	301
2013	137	145	30	312
2014	147	155	26	328
2015	145	147	26	318
2016	134	136	28	298
2017	121	126	26	273
2018	123	122	28	273

TABLE 1.2: Summary statistics of returns and predictor variables.

The sample period is from July 2001 to December 2018. Panel A presents the summary statistics for the full sample, aggregated across the seven commodities. Stock return is the average return over the sample period in percent per month. Interest rate is the 3-month US T-Bill rate return in percent per month. Spot price growth is the average growth rate in percent per month of the commodity prices determined as the one-month-ahead log spot price minus the log spot price times 100.y/x is the annual production (extraction) of each firm per unit of proven and probable reserves averaged across all years and commodities. qx/p is the value of reserves divided by market equity averaged across all firms and years. b/$(p+b)$ is the average across all firms and period of debt as a fraction of the value of the firm, as a measure of leverage. Panels B, C, and D provide the same statistics separated into precious metals (gold and silver), fuel energy (oil and gas), and industrial metals (copper, nickel, and zinc). std represents the standard deviation of the variable across the sample; $10\%, 50\%$, and 90% indicate the cumulative distribution values across the sample.

	Panel A: Summary Statistics-All Commodities				
	mean	std	10%	50%	90%
stock return	0.591	15.215	-18.688	0.201	19.141
interest rate	0.106	0.128	0.003	0.074	0.319
spot price growth	0.871	1.403	-1.012	0.723	2.747
y/x	0.127	0.076	0.042	0.108	0.243
qx/p	4.654	5.331	0.265	2.167	10.800
$b/(p+b)$	0.340	0.227	0.056	0.303	0.683
	Panel B: Summary Statistics-Precious Metals				
	mean	std	10%	50%	90%
stock return	1.212	16.036	-18.276	-0.602	22.273
spot price growth	0.813	1.012	-0.662	1.054	1.937
y/x	0.105	0.080	0.030	0.076	0.243
qx/p	5.628	6.072	0.613	3.291	14.686
b/$(p+b)$	0.268	0.220	0.047	0.195	0.617
	Panel C: Summary Statistics-Fuel Energy				
	mean	std	10%	50%	90%
stock return	0.408	15.050	-18.917	-0.183	18.286
spot price growth	0.883	1.471	-1.124	0.715	2.836
y/x	0.131	0.074	0.050	0.113	0.244
qx/p	3.660	4.868	0.218	1.972	9.222
b/$(p+b)$	0.352	0.226	0.056	0.322	0.696
	Panel D: Summary Statistics-Industrial Metals				
	mean	std	10%	50%	90%
stock return	2.119	16.051	-17.878	0.577	22.999
spot price growth	1.083	1.402	-0.635	0.846	2.968
y / x	0.088	0.072	0.029	0.062	0.196
qx/p	9.832	8.509	0.507	7.033	22.297
b/$(p+b)$	0.331	0.226	0.088	0.260	0.658

Table 1.3: Results of production prediction regressions.

The production variable is the ratio of the annual extraction quantity by the company per unit of proven and probable reserves, $\widehat{y_{t+1}/x_{t+1}}$. The forecast variables are: the weighted log of the expected future price, $\ln(\widehat{q_{t+1}})$, and the log spot price, $\ln(q_t)$, associated with the commodities produced by the firm (weighted by value); the real risk-free interest rate, r_t; the firm's weighted production of each commodity divided by the proven and probable reserves, y_t/x_t; and the firm's value of the proven and probable reserves divided by its market equity, qx/p. The regressions are: a pooled regression using annual observations with the predictive values lagged by one year (pooled) and a panel regression with firm fixed effects using annual observations with the predictive values lagged by one year (panel). T-stats are given in parentheses. "*" indicates significant at the 10% level, "**" significant at the 5% level, and "***" significant at the 1% level.

	$\widehat{y_{t+1}/X_{t+1}}$	
	Pooled	Panel
const	0.016***	
	(4.22)	-0.005
ln(exp spot price)	-0.010	(−0.87)
	(−0.80)	-0.000
ln(spot price)	0.011	(−0.02)
	(0.98)	0.057
interest rate	0.116	(0.67)
	(1.39)	0.461***
y_t /x_t	0.914***	0.461***
	(57.14)	(19.27)
q/x	−0.0016***	−0.0021***
	(−7.38)	(−7.16)
rsquared	0.459	0.632
no. observations	4528	4528

TABLE 1.4: Stock Return Prediction Regression

OLS regression results across all firms and periods. The dependent variable is the monthly stock return which is predicted monthly from lagged variables. The prediction variables are the log of the weighted expected commodities price for the next month based on the weights of the commodities the firm produces, $\ln(q_{t+1})$; the log of the weighted spot price based on the weights of the commodities the firm produces, $\ln(q_t)$; the real risk-free interest rate, r_t (3-month T-Bill minus inflation rate); the current production level relative to the level of the firm's reserves, y_t/x_t; the forecast of future production relative to the level of reserves based on previous-year variables, $\widehat{y_{t+1}/x_{t+1}}$; the last annual observation of firm reserves valued at current spot prices relative to the current market value of the firm's equity, $q_t x_{t+1}/p_t$. The excess expected spot price, Δ spot, equals $\ln(q_{t+1}) - \ln(q_t) - r_t$. The predicted production difference, Δ prod, equals $\widehat{y_{t+1}/x_{t+1}} - (1 + r_t)(y_t/x_t)$. Panel A shows results for pooled regressions and Panel B for panel regressions with firm-level fixed effects. Panel C provides variation inflation factor (VIF) statistics to detect multicollinearity, and chi-squared statistics to test parameter restrictions. Tstats are in parentheses. Standard errors are Shanken (1992)-adjusted for measurement error in the estimated variables. "*" indicates significant at the 10% level, "**" significant at the 5% level, and "***" significant at the 1% level.

Panel A: Pooled Regression

Dependent Variable: $r^S_{t+1} - r_t$				
	1	2	3	4
constant	-0.082	-0.099	-0.268	-0.411***
	(-0.46)	(-1.20)	(-0.95)	(-3.54)
ln (exp spot price)	0.569		0.608***	
	(11.42)		(9.94)	
ln (spot price)	-0.569***		-0.608***	
	(-11.43)		(-9.96)	
interest rate	-0.735***		-0.681***	
	(-3.27)		(-2.50)	
Δ spot		0.572***		0.611***
		(11.51)		(10.65)
$\widehat{y_{t+1}/x_{t+1}}$			-9.181***	
			(-4.12)	
y_t/x_t			9.984***	
			(4.22)	
Δ prod			-9.183***	
				(-4.42)
qx/p		0.098***	0.095***	
			(3.44)	(3.40)
r-squared	0.0027	0.0025	0.0043	0.0041
number of observations	52337	52337	52337	52337

Panel B: Firm Fixed Effects Panel Regression

Dependent Variable: $\mathrm{r}^{\mathrm{S}}_{\mathrm{t+1}} - \mathrm{r_t}$				
	1	2	3	4
ln (exp spot price)	0.360***		0.421***	
	(6.09)		(4.71)	
ln (spot price)	-0.375***		-438***	
	(-6.32)		(-4.89)	
interest rate	-1.009***		-0.931***	
	(-4.45)		(-2.74)	
Δ spot		0.300***		0.342***
		(5.14)		(4.76)
$\widehat{y_{t+1}/x_{t+1}}$			-17.67***	
			(-5.81)	
y_t/x_t			16.97***	
			(4.51)	
Δ prod			-11.58***	
				(-4.76)
qx/p		0.132***	0.150***	
			(3.44)	(4.38)
r-squared	0.0242	0.0219	0.0265	0.0235
number of observations	52337	52337	52337	52337

Panel C: Multicollinearity and Coefficient Restrictions

Variable	$\ln(q_{t+1})$	$\ln(q_t)$	r_t	$\frac{y_{t+1}}{x_{t+1}}$	$\frac{y_t}{x_t}$	$\frac{q_t x_{t+1}}{p_t}$	Δ spot	Δ prod	$\frac{q_t x_{t+1}}{p_t}$
VIF	76.34	77.01	1.13	4.35	4.27	1.10	1.01	1.01	1.00
			pooled					panel	
Restriction		1	2	3	4	1	2	3	4
$b_I = 1$	χ^2-stat	74.56	74.45	60.92	60.71	117.32	144.87	94.97	127.06
	p-value	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
$b_1 = -b_2$	χ^2-stat	0.001		1.80		113.29		145.96	
	p-value	(0.98)		(0.18)		(0.00)		(0.00)	
$b_4 = -b_5$	χ^2-stat			0.65				0.18	
	p-value			(0.42)				(0.68)	

Table 1.5: Return prediction regression with commodity interaction effects.

Results based on a pooled OLS regression across all firms and periods. The dependent variable is the monthly stock return which is predicted monthly from lagged variables. The prediction variables are the log of the weighted expected commodities price for the next month based on the weights of the commodities the firm produces, $\ln\left(\widehat{q_{t+1}}\right)$; the log of the weighted spot price based on the weights of the commodities the firm produces, $\ln\left(q_t\right)$; the real risk-free interest rate, r_t; the current production level relative to the level of the firm's reserves, y_t/x_t; the forecast of future production relative to the level of reserves based on previous-year variables, $\widehat{y_{t+1}/x_{t+1}}$; the last annual observation of firm reserves valued at current spot prices relative to the current market value of the firm's equity, $q_t x_{t+1}/p_t$. The excess expected spot price, Δ spot, equals $\ln\left(\widehat{q_{t+1}}\right) - \ln\left(q_t\right) - r_t$. The predicted production difference, Δ prod, equals $\widehat{y_{t+1}/x_{t+1}} - \left(1 + r_t\right)\left(y_t/x_t\right)$. PreciousMetals and Industrial Metals are the intercept dummy variables. In addition, the interactions of the commodity dummies with the explanatory variables are included. T-stats are in parentheses. Standard errors are Shanken (1992)-adjusted for measurement error in the estimated variables. "*" indicates significant at the 10% level, "**" significant at the 5% level, and "***" significant at the 1% level.

	Dependent Variable: $\mathrm{r_{t+1} - r_t}$				
	Pooled			Fixed Effects	
constant	-0.411***	-0.331**	-0.663***		
	(-3.54)	(-2.16)	-4.19)		
PreciousMetals			2.602***		
			(4.73)		
IndustrialMetals			0.047		
			(0.06)		
Δ spot	0.611***	0.781***	0.851***	0.342***	0.612***
	(10.65)	(9.80)	(10.96)	(4.76)	(6.35)
Δ prod	-9.183***	-18.28***	-17.08***	-11.58***	-18.83***
	(-4.42)	(-5.96)	(-5.77)	(-4.76)	(-5.54)
qx/p	0.095***	0.060**	0.085***	0.150***	0.131***
	(3.40)	(1.83)	(2.57)	(4.37)	(2.99)
(PreciousMetals) · (Δ spot)		-0.956***	-1.830***		-1.761***
		(-4.10)	(-6.17)		(-5.16)
(PreciousMetals) · (Δ prod)		43.17***	34.87***		35.69***
		(5.22)	(4.27)		(3.63)
(PreciousMetals) · (qx/p)		0.088**	-0.073*		-0.049
		(2.08)	(-1.41)		(-0.63)
(IndustrialMetals) · (Δ spot)		-0.129	0.058		-0.261
		(-0.39)	(0.15)		(-0.60)
(IndustrialMetals) · (Δ prod)		-12.86	-13.016		-19.27*
		(-1.09)	(-1.11)		(-1.41)
(IndustrialMetals) · (qx/p)		0.002	0.010		-0.054
		(0.04)	(-0.15)		(-0.53)
r-squared	0.0041	0.0067	0.0082	0.0235	0.0267
number of observations	52,337	52,337	52,337	52,337	52,337

Table 1.6: Portfolio Sorting Returns

The average returns are shown by quintiles. The quintiles are sorted from low to high by the predicted returns and for each quintile we show the subsequent (one month later) realized return averaged over the pseudo-out-of-sample time periods (July 2003 - December 2018). To forecast the return for time $t+1$, we use the fitted value from equation (16), with dummy variables for the various specifications, from the coefficients based on all prior observations up to time t along with the predictor variables at t, to sort all firms into quintiles. We use the first 24 time series data points to estimate initial coefficients and use an expanding window for subsequent estimation. Quintile 1 in each month includes the observations (firms) with the 20% lowest predicted returns, and Quintile 5 in each month contains the observations with the 20% highest predicted returns. The subsequent monthly returns for each quintile are recorded and averaged and listed as ret 1 through ret 5 for quintiles 1 through 5 , respectively. "tstat" refers to the t-statistic for the test of significance of the return compared to 0 . "alpha1" refers to the riskadjusted return based on the five-factor model of Fama and French (2015). "alpha2" refers to the riskadjusted return using the five-factor model of Fama and French (2015) plus the carry factor specific for commodity price risk based on Bakshi, Guo, and Rossi (2019). Panels A and B present the results for all firms sorted based on the predictions from Tables 4 A(4) and 4 B(4), respectively.

Panel A: Portfolio Sorting - Full Sample (Pooled)

	ret 1	ret 2	ret 3	ret 4	ret 5	ret 5-1
mean	-0.165	0.093	0.621	0.499	1.078	1.243
tstat	-0.287	0.161	1.054	0.803	1.767	3.173
alpha1	-0.810	-0.649	-0.258	-0.407	0.277	1.088
tstat	-1.629	-1.359	-0.535	-0.757	0.493	2.731
alpha2	-1.007	-0.817	-0.410	-0.590	0.165	1.172
tstat	-2.069	-1.735	-0.859	-1.113	0.294	2.946

Panel B: Portfolio Sorting - Full Sample (Fixed Effects)

	ret 1	ret 2	ret 3	ret 4	ret 5	ret 5-1
mean	-0.358	0.285	0.398	0.605	1.191	1.549
tstat	-0.617	0.505	0.673	1.009	1.968	4.677
alpha1	-1.073	-0.428	-0.481	-0.308	0.437	1.510
tstat	-2.178	-0.905	-0.991	-0.595	0.790	4.449
alpha2	-1.259	-0.613	-0.666	-0.469	0.341	1.599
tstat	-2.601	-1.324	-1.400	-0.914	0.615	4.737

Chapter 2

Carbon Pricing and the Commodity Risk Premium

2.1 Introduction

The field of climate finance has rapidly expanded in recent years, with a significant focus on carbon risk due to the primary role of human carbon emissions in global warming. To mitigate greenhouse gas (GHG) emissions and prevent catastrophic climate outcomes, numerous jurisdictions have implemented carbon pricing mechanisms, such as carbon taxes and cap-and-trade. As of December 2021, carbon taxes have been implemented in 14 [1] out of the 31 high-income OECD countries, including Canada, France, Japan, Sweden, Switzerland, and the United Kingdom Yunis and Aliakbari (2020). Additionally, seven regions, including the EU, China,

[1]The fourteen countries that implemented carbon taxes are Canada, Chile, Denmark, Finland, France, Iceland, Ireland, Japan, Norway, Portugal, Slovenia, Sweden, Switzerland, and the United Kingdom.

Korea, Australia, New Zealand, California (US), and Quebec (Canada), have established cap-and-trade systems.

Carbon pricing is recognized as an efficient and cost-effective strategy to reduce GHG emissions and is expected to become increasingly prevalent. Consequently, understanding how investors perceive and manage carbon pricing risk is crucial for advancing climate finance research. This paper conducts an empirical analysis to investigate whether carbon pricing risk explains the cross-section of commodity futures returns, motivated by three key aspects of the existing literature.

Firstly, the current climate finance literature overlooks the impact of carbon pricing as a risk factor, focusing instead on firm-level carbon-related characteristics such as carbon emissions and ESG ratings. Bolton and Kacperczyk (2021) demonstrate that U.S. stocks with higher total carbon emissions earn higher returns and that institutional investors screen out firms with high Scope 1 emission intensity. Hsu et al. (2023) use toxic emission intensity from the Toxic Release Inventory (TRI) database and find a spread between brown and green U.S. stocks. Choi et al. (2020), using global stock data, show that carbon-intensive firms underperform during periods of abnormally warm weather, and that public attention to climate change, measured by Google search volume, increases during these warmer periods. Based on the above, it is evident that previous studies on climate finance have seldom considered carbon pricing. The most recent paper that comes close to carbon pricing risk is Azlen et al. (2022), where carbon market is explored as an asset class and a prospective carbon risk premium is identified. However, the carbon price as a significant transition risk affecting other assets still remained unexplored.

Secondly, while extensive research has examined the pricing of climate change risk across various asset classes such as stocks, options, bonds, and so on, the commodity market remains underexplored. Given the growing investor interest in climate change risk, increasing attention has been directed towards studying the pricing of climate-related risks in various financial markets. In the stock market, Görgen et al. (2020) create a carbon factor based on ESG rating variations to capture firms' sensitivity to the low-carbon transition, while Barnett (2019) uses event study analysis to show that stocks with high climate policy risk exposure suffer significant declines in returns following policy announcements. In the options market, Ilhan et al. (2021) find that the cost of option protection is higher for carbon-intensive firms, particularly when public attention to climate risk is heightened. Similarly, Huynh and Xia (2021) demonstrate that bonds with positive covariance with a climate news index earn lower returns. In the real estate market, Giglio et al. (2021) explore the effects of climate change on property valuation by studying the pricing of the risk of sea level rise in four coastal US states: Florida, New Jersey, North Carolina, and South Carolina. Despite substantial progress in understanding the relationship between climate change and asset prices, the commodity futures market has not been sufficiently studied in this context.

Thirdly, various studies have proposed hedging strategies against climate change risk using financial assets, yet few address the investment constraints faced by sustainable investors. Engle et al. (2020) develop a climate risk measure using textual analysis and construct a mimicking portfolio based on firms' E-scores [2]. Alekseev

[2] E-score is an objective measurement or evaluation of a given company, fund, or security's performance with respect to Environmental issues. It measures whether the organization is operating as a steward of the environment and covers sustainability issues such as greenhouse gases (GHG), loss of biodiversity, carbon emissions and pollution. https://corporatefinanceinstitute.com/resources/esg/esg-score/

et al. (2022) propose a quantity-based approach that leverages mutual fund managers' trading responses to extreme heat events to build a climate hedge portfolio using stocks. While these methodologies show promise, they do not consider the potential conflicts between hedging demand and the divestment requirements of sustainable investments in specific "brown" industries. In this paper, I adopt the quantity-based approach of Alekseev et al. (2022) to construct a hedging portfolio using commodity futures based on carbon pricing risk loadings. Commodity futures offer a distinct advantage as they do not finance the production of underlying commodities or directly affect their supply, unlike investments in the equities or bonds of carbon-intensive firms. Therefore, sustainable investment metrics such as carbon emissions intensity or ESG ratings are not directly applicable to commodity futures investment (Danielsen (2020)). By constructing a hedging portfolio with commodity futures, investors can manage climate change risk without increasing their carbon footprint.

In this paper, I examine the impact of carbon pricing risk on commodity futures. Carbon pricing, employed as a regulatory measure to reduce greenhouse gas (GHG) emissions, is implemented through various mechanisms worldwide. Specifically, this study refers to the price of tradable carbon emission allowances within a cap-and-trade system, known for its flexibility and cost-effectiveness in curbing GHG pollution. Under this system, a cap limits the total amount of certain greenhouse gases that firms can emit, with the cap progressively decreasing to ensure overall emissions reduction. Participating firms can trade emissions allowances among themselves, ensuring these allowances hold economic value due to the cap on their total number.

For this research, I use the price of emission allowances from the EU ETS (European Union Emission Trading Scheme) as a proxy for carbon pricing. Launched in 2005, the EU ETS is the largest and oldest carbon market in the world, covering over 11,000 facilities across 31 countries and accounting for 45% of the EU's total GHG emissions [3]. The prices of carbon allowances from the EU ETS have been extensively studied and serve as a primary reference for carbon pricing analysis. Recent research, such as Tan et al. (2020), explores the interconnectedness of the European carbon market with a broad spectrum of markets, including equities, bonds, and commodities. These literatures [4]motivate my investigation into carbon pricing risk in commodity markets, an area less explored despite extensive research on carbon pricing's economic implications. Understanding the dynamic links between carbon prices and commodity markets can also provide insights into the effectiveness of the cap-and-trade system in curbing GHG pollution.

Unlike stocks or bonds that reflect climate risks through characteristics like E-scores, commodities lack such individual-level features. However, carbon pricing can significantly impact commodity production costs, influencing their market dynamics indirectly through covariance with climate change risks. In the following sections, I analyze how carbon pricing shocks, akin to negative investment shocks in financial theory, affect commodity futures returns, drawing on methodologies from Kogan et al. (2009) and Yang (2013). Empirical findings across a comprehensive sample of 35 commodities confirm that carbon pricing risk is priced in commodity futures markets. Specifically, the risk premium associated with carbon pricing

[3]https://climate.ec.europa.eu/system/files/2016-12/factsheet$_{e}ts_{e}n$.pdf

[4]In addition to the paper mentioned, there are also Kim and Koo (2010), Chevallier (2011), Sousa et al. (2014), Yu et al. (2015), Koch et al. (2014) , Zhang and Sun (2016), , Wen et al. (2017), Ji et al. (2018), Wang and Guo (2018), Bai and Okullo (2023), and Ahmed et al. (2024)

shocks is statistically significant and economically meaningful: a one-standard-deviation increase in carbon pricing risk correlates with a 0.085 percent increase in daily excess returns of commodity futures.

This paper contributes to several areas of the literature. Firstly, it identifies a new risk factor in the commodity futures market, adding to studies on commodity risk premia. Secondly, it expands climate finance research by demonstrating that carbon pricing risk affects commodity futures returns. Finally, it suggests a novel hedging strategy using commodity futures, providing a viable alternative for sustainable investors seeking to hedge climate change risk without compromising their investment principles.

The rest of this paper is organized as follows. Section 2.2 presents the hypothesis development based on an investment-based model for commodity producers. Section 2.3 describes the sample and data, followed by the empirical results and economic significance analysis in Section 2.4. Hedging portfolio test results are discussed in Section 2.4. Section 2.5 concludes the paper.

2.2 Hypothesis Development

The development of the hypotheses is influenced by Yang (2013), whose investment-based asset pricing model explains how shocks to investment affect risk premiums in commodity futures. In Yang's framework, commodity futures risk premiums are affected by investment shocks, aggregate demand shocks, and idiosyncratic demand shocks. An investment shock represents uncertainty affecting the real investment of all commodity producers. A positive shock indicates technological advancements in capital production, leading to increased future capital (higher commodity supply) for the same investment level. Conversely, a negative investment shock reduces future capital produced for the same amount of investment, thereby decreasing commodity supply. In the context of carbon pricing, an unexpected increase in the carbon price serves as a negative investment shock in the commodity production because of its negative impact on the efficiency of commodity investment.

My hypotheses build on this investment-based model by considering how carbon pricing uncertainty impacts the efficiency of commodity investment. An increase in carbon pricing will raise investment costs not only for high-carbon-emitting commodities like oil and gas but also for other low-carbon-emitting commodities through higher electricity and transportation prices (Bolton and Kacperczyk (2021)). Thus, an unexpected rise in the carbon price (carbon pricing shock) acts as a negative investment shock affecting all commodity producers' investment decisions. Commodities with current high investment levels are more vulnerable to carbon pricing shocks due to their substantial new capital installations.

Yang (2013) suggests that risk premiums on positive investment shocks are negative, indicating investors accept lower returns on portfolios with such shocks. Given carbon pricing as a negative investment shock, commodity futures are expected to exhibit a positive risk premium in response to carbon pricing shocks. This expectation arises because a negative investment shock reduces investment efficiency, thereby decreasing future capital (future supply) relative to current investment levels, leading to higher commodity prices. The positive risk premium on carbon pricing shocks suggests that investors demand higher returns on commodity futures exposed to higher carbon pricing risks. Based on this reasoning, I propose the following hypothesis regarding carbon pricing risk:

Hypothesis 1: The carbon pricing risk has a significant and positive risk premium in the commodity futures market.

Given that carbon pricing risk affects commodity futures risk premiums through the investment channel, commodities characterized by substantial investment (i.e., crude oil, natural gas, etc.) are likely to be more vulnerable to this carbon pricing risk due to their extensive capital requirements, intensifying the impact of carbon pricing uncertainties. Furthermore, variations in carbon emissions during production contribute to different sensitivities to carbon pricing risk across commodities. Those with lower carbon emission such as wheat, corn, and sugar, are anticipated to exhibit lower vulnerability to carbon pricing shocks. From this analysis, the following hypothesis emerges:

Hypothesis 2: Commodities within the energy and metal mining sectors will demonstrate higher risk premiums associated with carbon pricing risk, whereas commodities within the grain, soft, and livestock sectors will exhibit lower risk premiums.

The growing societal awareness and concern regarding climate change are significantly influencing investor behavior. The exponential growth in ESG investments indicates heightened investor awareness of climate change risks. According to the Global Sustainable Investment Alliance [5], global sustainable investment reached $35.3 trillion in 2020, a 15% increase over two years. Concurrently, there has been a notable increase in climate-focused financial products, such as green

[5] https://www.gsi-alliance.org/wp-content/uploads/2021/08/GSIR-20201.pdf

bonds and climate indices. These developments reflect investors' growing focus on climate-related risks, including carbon pricing risk. Previous research on carbon risk (Bolton and Kacperczyk (2021), Huynh and Xia (2021),Ilhan et al. (2021)) indicates that carbon risk premia increase as investors' awareness about climate change risk grows. Consequently, it is anticipated that carbon risk premia are expected to rise over time as investor awareness and concerns about climate change risk grow. To assess whether the carbon pricing risk premium in the commodity futures market exhibits an increasing trend, the following hypothesis will be examined:

Hypothesis 3: The risk premium associated with carbon pricing risk in the commodity futures market increases over time.

Finally, I examine whether the commodity futures' risk loading on carbon pricing risk varies depending on any specific commodity characteristics. The literature on commodity futures pricing can be viewed as consisting of two strands: the theory of storage (Brennan (1976); Kaldor (1939); Working (1949)) and the theory of normal backwardation (Hicks (1975); Keynes (1930)). This study examines two key characteristics—basis and hedging pressure—derived from these theories, which play pivotal roles in understanding commodity futures pricing dynamics. Building on Hypotheses 1 and 2, which link risk exposure to commodity production investment, any influence of these characteristics on carbon risk sensitivity is expected to operate through the investment channel.

The first commodity characteristic, basis, represents difference between the short-term futures price and the long-term futures price, providing insights into commodity futures risk premia (Bakshi et al. (2019),Boons and Prado (2019), Gorton et al. (2013), Szymanowska et al. (2014), Yang (2013)). According to the theory of storage, holders of commodity inventories receive a "convenience yield" as an implicit benefit of keeping commodities in storage, and this convenience yield is closely tied to the basis (Görgen et al. (2020)).

Theoretically, the basis can be decomposed as two parts: the expected spot price change $P_t - E[P_{j,T}]$ plus the risk premium $-\frac{cov(M_t,P_{j,T})}{E[M_T]}$ (eq. 1)[6]. A high basis, or equivalently a lower long-term futures price, can arise because of either a lower expected spot price or a higher risk premium (more negative covariance between the stochastic discount factor (SDF) and spot price, given the positive expected futures excess returns most of the time).

$$Basis_t = P_t - F_{t,T} = P_t - E[P_{j,T}] - \frac{cov(M_t, P_{j,T})}{E[M_T]}$$

Commodity futures with a high basis, as suggested by Yang (2013), tend to exhibit higher investment rate, thus are more sensitive to any investment shocks. At the same time, a high investment rate predicts high future capital employed and, hence, a low expected spot price ($E[P_{j,T})$ due to increased future supply. Both implying that commodities with a high investment rate are expected to show a higher basis. Therefore, high-basis commodities are expected to have high risk loadings on carbon pricing risk. The other strand of commodity futures theoretical literature, normal backwardation, posits that hedgers (both commodity producers and commodity inventory holders) manage future price risk by taking short positions in the commodity futures market. To induce speculators to take the opposite

[6] $F_{t,T} = \frac{E(M_T,P_{j,T})}{E[M_T]} = E[P_{j,T}] - \frac{cov(M_t,P_{j,T})}{E[M_T]}$

long positions, commodity futures prices are set at a discount to expected future spot prices at maturity, generating a return for speculators' risk bearing. The theory of normal backwardation links hedging pressure, measured by the relative size of short positions taken by commodity producers or inventory holders, to the risk premium of commodity futures. For instance, De Roon et al. (2000) show that hedging pressure significantly affects commodity futures returns. As hedging pressure measures the commodity producers' need to hedge future price risk, higher hedging pressure, indicating greater hedging needs by commodity producer, is associated with higher future production and investment. Aligning with Hypothesis 1, commodities with higher investment and more future capital tend to have higher risk sensitivity to carbon pricing risk. Therefore, commodities experiencing higher hedging pressure are expected to have higher risk sensitivities to carbon pricing risk. Therefore, commodities with higher basis and hedging pressures are hypothesized to demonstrate stronger risk loadings on carbon pricing risk. The analysis of the two commodity-specific characteristics leads to the following hypothesis:

Hypothesis 4: The commodity's risk loading on the carbon pricing risk,β^{carbon}, positively correlates with basis and hedging pressure.

2.3 Data

2.3.1 Carbon pricing

"Carbon pricing is an instrument that captures the external costs of greenhouse gas (GHG) emissions—the costs of emissions that the public pays for, such as damage to crops, health care costs from heat waves and droughts, and loss of property from flooding and sea level rise—and ties them to their sources through a price, usually in the form of a price on the carbon dioxide (CO2) emitted." [7] According to the World Bank, more than 40 national jurisdictions have implemented carbon pricing initiatives. The two most common forms of carbon pricing are emission trading systems (ETS) and carbon taxes. As of December 2021, 38 national and regional jurisdictions have implemented or scheduled the implementation of ETS as their carbon pricing initiative.

An emission trading system (ETS) is a mechanism where emitters can trade emission allowances to meet their emission targets. By creating a supply and demand market for emission units, an ETS establishes a market price for GHG emissions. A main type of ETS [8])is the cap-and-trade system, where a cap is set on emissions from installations or factories within the ETS, and emission allowances are initially distributed for free or through auctions. The cap is reduced over time to ensure that the total amount of emissions falls.

The carbon price data used in this paper is the emission allowance price series from the European Union (EU) ETS, which is a cap-and-trade system implemented by the EU in 2005. The European Union Allowance (EUA), which permits the emission of one metric ton of CO2 under the EU ETS, has both spot and futures markets in exchanges such as the European Climate Exchange, ICE, and the European Energy Exchange. The EU ETS has undergone three phases: the first phase ran from 2005 to 2007, the second from 2008 to 2012, and the third

[7] https://carbonpricingdashboard.worldbank.org/what-carbon-pricing

[8] There are two main types of trading systems: "Cap-and-trade systems" and "baseline-and-credit systems." Under a baseline-and-credit system, there is no fixed limit on emissions, but polluters that reduce their emissions more than they otherwise are obliged to can earn 'credits' that they sell to others who need them to comply with regulations they are subject to. (https://www.oecd.org/env/tools-evaluation/emissiontradingsystems.html

phase started in 2013 and ended in 2020. The EUA daily price data used in this paper covers phase 2 and 3 and is collected from Bloomberg (ticker: MO1 Comity), spanning from December 31, 2008, to December 31, 2020. Figure 2.1 shows the carbon price data used over the sample period.

2.3.2 Commodity futures data

The commodity futures price dataset consists of 35 commodities across 5 categories: energy, grain, livestock, metal, and soft commodities. This sample is comparable to those used in studies by (Gorton et al. (2013), Hong and Yogo (2010), Sakkas and Tessaromatis (2020), Szymanowska et al. (2014), Bakshi et al. (2019)). Following the common practice in these studies, the daily and monthly excess returns for commodity futures are calculated as follows:

$$R^e_{i,t+1,T} = \frac{F_{i,t+1,T}}{F_{i,t,T}} - 1$$

Where $F_{i,t,T}$ is the price at time t of the contract with maturity T, and $F_{i,t+1,T}$ is the price of the same futures contract at time $t + 1$. Guided by Bakshi et al. (2019), $F_{i,t,T}$ is chosen as the second nearest contract at the end of month t in order to avoid any impact on price caused by physical delivery while ensuring its first notice day is after the end of month $t + 1$. For each commodity i, the futures returns are calculated based on a roll-over strategy, where an investor maintains a long position in the second-nearest futures contract. All the commodity futures price data is collected from Bloomberg.

To conduct the empirical tests, the 35 commodities are sorted based on their sorting characteristics at the end of time t, and the factor portfolio returns are calculated over the subsequent time $t + 1$. The commodities are then sorted into quintiles, and the portfolio returns are computed as the equal-weighted average of the excess returns of the commodities included in each portfolio. The excess return of the "average" portfolio, which is the equal-weight excess return of all 35 commodities in the sample, is calculated as the market factor. Following Gorton et al. (2013), Hong and Yogo (2010), Sakkas and Tessaromatis (2020), Szymanowska et al. (2014), and Bakshi et al. (2019), I sort the 35 commodities based on five sorting characteristics: Basis, Momentum, Basis-Momentum, Value, and Volatility[9].

The summary statistics for 35 commodities in our sample are reported in Table 2.1. The average annualized excess returns over the period from December 2008 to December 2020 range from 22.5% for natural gas to 30.4% for palladium. Out of the 35 commodities, 23 have positive average returns and 12 have negative average returns. 23 out of 35 commodities have a Sharpe ratio less than 0.25, consistent with previous findings (Bakshi et al. (2019),Sakkas and Tessaromatis (2020)). Table 2.2 reports the summary statistics of the portfolio sorted by the commodity characteristics (basis, momentum, basis momentum, value, and volatility). The high-minus-low portfolio based on basis shows an annualized return of 12.3% with a t-statistic of 2.44, consistent with prior Bakshi et al. (2019), Gorton et al. (2013), Sakkas and Tessaromatis (2020), and Yang (2013). The high-minus-low portfolio based on basis-momentum exhibits a positive annualized return of 14.01%. The high-minus-low portfolios based on momentum, value, and volatility do not show statistically significant returns in this sample, possibly due to the use of daily instead of monthly returns for portfolio construction.

[9]Please refer to the definition for the five characteristics in Appendix

Panel A of Table 2.3 provides the summary statistics on the carbon price sample data from EU ETS. The carbon price data is the settlement price of the generic 1st futures of EUA from December 2008 to December 2020 [10]. The carbon price is converted to USD using the exchange rate from Bloomberg. The average carbon price over the sample period is $14.6 with minimum $3.54 on April 17, 2013, and the maximum $40.63 on December 28, 2020. The ADF test with trend for carbon price level shows a strong evidence of unit root as the test statistics indicate that the null hypothesis of the unit root cannot be rejected. I use the change rate of carbon price level as the carbon pricing shock and the second line of panel A in Table 2.3 reports the statistics of daily carbon pricing shock with the mean of 0.065 percent and standard deviation of 3.3 percent. The statistics for the ADF test strongly reject the null hypothesis that the carbon pricing shock has a unit root. Panel B of Table 2.3 reports the full sample unconditional correlation between the commodity portfolios returns and the carbon pricing shock. For the 25 portfolios constructed based on the five commodity characteristics (basis, momentum, basis-momentum, volatility, and value) the carbon pricing shock is positively correlated. The correlation between the carbon pricing shock and the high-minus-low portfolios returns are not significant.

[10] I exclude the first phase of EU ETS and the first year of phase 2 (2008) since the trading system collapsed in early 2008 and carbon price drop to zero due to the policy adjustment instead of market fundamentals. This adjustment is consistent with Tan et al. (2020)

2.4 Empirical Results

2.4.1 Cross-sectional tests: carbon pricing shock as a risk factor

The Fama and MacBeth (1973) methodology is employed to assess the impact of carbon pricing shock on commodity futures returns across multiple factor models. This approach allows for the estimation of factor sensitivities and risk premia, which are crucial for addressing Hypotheses 2 and 3.

Based on the existing studies on the commodity market, there is a lack of consensus on the factors that can explain the average commodity futures returns. Yang (2013) proposes a two-factor model with a basis factor along with the commodity market factor [11]. Bakshi et al. (2019) suggest that the momentum factor contains additional information in explaining the commodity returns cross-sectionally. Boons and Prado (2019) substantiate that the basis-momentum factor is a priced risk, and its risk premium represents a reward for bearing commodity market volatility risk.

In this paper, I consider three specifications of commodity models: two-factor model (market, basis), three-factor model (market, basis, momentum) and four-factor model (market, basis, momentum, basis-momentum). These are the most common systematic risk factors that have been found priced in the prior literature on commodity asset pricing (Sakkas and Tessaromatis (2020)). To these factor models, I add the carbon pricing shock as an additional factor. As stated

[11]The commodity Market factor refers to the return of the equal-weighted portfolio using all commodities in the sample (Bakshi et al. (2019), Yang (2013)

in hypothesis 1, a positive risk premium for carbon pricing risk is expected to compensate for the additional risk.

Employing the two-pass methodology, I estimate the factor sensitivities at the end of each t using the 24-month daily data up to $t-1$ in the first pass. In the second pass, I obtain the daily risk premia for the risk factors using the estimators from the first pass [12] . Table 2.4 presents the cross-sectional results with 25 commodity characteristics-sorted portfolios (basis, momentum, basis-momentum, volatility, value) as test assets. The second pass results of two-factor, three-factor, and four-factor models, and the augmented models with carbon pricing risk added are reported in Table 2.4. The results reveal that the risk premium for carbon pricing risk is significantly positive in all three factor model specifications, which confirms Hypothesis 1. The premium estimates for carbon pricing risk range from 0.06516 percent (two-factor plus carbon pricing risk) to 0.08560 percent (four-factor plus carbon pricing risk).

2.4.2 Sensitivity to carbon pricing risk

To test Hypothesis 2, I estimate the factor loadings for each individual commodity to examine if the carbon pricing shock has an uneven impact across different commodities. Table 2.5 reports the factor loadings from the time series regressions of all 35 commodities over the whole sample period from December 31, 2008, to December 31, 2020. Overall, 8 of the 35 commodities have significantly positive loadings on the carbon pricing shock, most of which fall into the energy sector. Brent Crude Oil, Crude Oil, Gasoil, Gasoline, Heating Oil, Copper, and

[12] The betas obtained from the first pass are rescaled to be on the same scale.

Aluminum have significantly positive sensitivity to carbon pricing risk. Intuitively, these commodities have high exposure to carbon pricing risk either because they are industries with heavy investment or because their production processes emit relatively substantial amounts of carbon emissions (e.g., copper mining). These industries also have high exposure to climate change risk and temperature shocks, as Balvers et al. (2017) suggest.

Most metal mining commodities, including Nickel, Platinum, Palladium, Silver, Tin, and Zinc, also have positive sensitivity to carbon pricing risk, although not significantly so. The risk loadings on the carbon pricing shock are negative for most agricultural commodities. Corn, Kansas Wheat, Oats, Sugar, Wheat, Soybeans, and Wheat have significant negative sensitivity to carbon pricing risk. The positive sensitivity of energy and metal commodities, coupled with the negative sensitivity of agricultural commodities, confirms the uneven impact described in Hypothesis 2.

In addition, based on the carbon pricing risk loadings shown in Table 2.5 and the summary statistics in Table 2.1, it is revealed that high-beta commodities tend to have higher (excess) returns than low-beta commodities. Specifically, sorting the individual commodities in the sample by their carbon pricing risk loadings in Table 2.5, the mean annual return of the top quartile is 3.4% higher than the mean annual return of the lowest quartile, which confirms Hypothesis 1 again.

2.4.3 Carbon pricing risk premium

To determine if the risk premium of carbon pricing risk is increasing over time (Hypothesis 3), I repeat the Fama and MacBeth (1973) procedure and obtain the risk premium time series estimated in the second pass regression. The risk premium at each time t is cross-sectionally estimated using the first-pass estimation based on the previous 24-month daily data. Consistent with the prediction of Hypothesis 3, Figure 2.2 shows that the moving average of the 5-year risk premium displays an upward trend. The risk premium of the carbon pricing shock is time-varying and increasing over time.

2.4.4 Commodity specific determinants of carbon pricing beta

I next test Hypothesis 4, the relationship between commodity-specific characteristics and carbon pricing risk loading. I regress the basis and hedging pressure of individual commodity futures on their risk loading at the next period t+1 on carbon pricing risk $(^{carbon}_{i,t+1})$. For each commodity futures on each day t, I estimate the carbon pricing risk beta from daily rolling regressions of commodity futures excess returns on the four-factor model plus carbon pricing shock over a 24-month window. Table 2.6 reports the estimation results. Columns 1 and 2 present the regression results for the commodity basis, while Columns 3 and 4 present the results for commodity hedging pressure. The hedging pressure is calculated as the relative size of short positions of commercial traders, $\frac{number of short hedge positions - number of long hedge positions}{total number of hedge positions}$ using data on large traders'[13] positions from CFTC (Commodity Futures Trading Commission).

[13]Large traders refer to commodity futures market participants who are subject to reporting requirements. The reporting levels are set by the Commodity Futures Trading Commission (CFTC). The current reporting levels can be found in CFTC Regulation 15.03(b). https://www.ecfr.gov/current/title-17/chapter-I/part-15/section-15.03

The results show that the coefficient on the basis is significantly negative in both the pooled regression (column 1) and the panel regression with time and firm/commodity fixed effects added (column 2). This indicates that the commodity basis is negatively associated with carbon pricing risk exposure, which contradicts the first half of Hypothesis 4. I infer the hypothesis of a positive relationship between basis and future carbon pricing risk sensitivity based on the relationship between basis and investment rate proposed by Yang (2013) using simulated data for 1000 years. The contradictory results here are likely due to the sample period length, 11 years, which may not be long enough to display the relationship between basis and investment rate implied by the investment-based model. The negative coefficient on the basis in Table 2.6 suggests that the basis, as a price measure, might not be an effective indicator of the investment of commodity producers during the sample period in this paper, and high-basis commodities tend to be less sensitive to the carbon pricing shock.

Columns 3 and 4 report significantly positive coefficients on hedging pressure in both pooled and panel regressions, confirming that commodities with higher hedging pressure tend to have higher carbon pricing exposure, supporting the second part of Hypothesis 4. The sample size for testing hedging pressure is smaller than that for the basis because CFTC data on large traders' positions does not cover all the commodities in the sample.

2.4.5 Commodity Hedging Portfolios for climate change risk

Given the growing awareness of the economic and financial risks associated with climate change, there is an increasing need among investors to hedge against these risks. In this section, I propose an innovative approach to construct a hedge portfolio using commodity futures to effectively mitigate climate change risk.

Following Alekseev et al. (2022), I utilize the quantity-based method to construct hedging portfolios. This approach allows us to leverage cross-sectional information on the carbon beta of commodity futures returns, which reflects their exposure to climate change risk. An advantage of this method is its independence from the requirement of long time series data. In contrast, conventional mimicking portfolio approaches for hedging are sensitive to the availability of time-series data and may be less effective with shorter sample periods. Adopting the quantity-based methodology proposed by Alekseev et al. (2022), I construct the hedging portfolio based on commodities' sensitivities to carbon pricing risk. The excess returns of the hedging portfolio are calculated as follows:

$$CarbonPortfolio_t = \sum_j \beta_{j,t}^{carbon} r_{j,t}$$

Monthly data is used to construct the hedging portfolio to align with the frequency of the climate news sample (the hedging target). For commodity j in month t, the carbon beta $_{j,t}^{carbon}$ is estimated using the previous 60 monthly observations of commodity futures returns and carbon pricing shocks. Each component of the portfolio represents an excess return; hence no scaling of weights is necessary.

One of the primary challenges in constructing a portfolio to hedge against climate risk is the absence of a definitive hedge target. Climate change is a multi-faceted issue encompassing physical risks such as sea level rise and transition risks like regulatory changes stemming from climate concerns in the future. Different investors may perceive and prioritize these risks differently, and their realization may not occur simultaneously. To address the challenge of hedging climate change risk, Engle et al. (2020) demonstrated the feasibility of using text analysis techniques to construct a series of news indices that serve as potential hedging targets for future climate risks. Following their approach, various climate risk news indices have been developed by researchers using different textual analysis techniques. In this paper, I employ several climate risk news indices as hedging targets to evaluate the performance of different portfolios. The focus of this study is not to ascertain the accuracy or appropriateness of these climate risk news indices as measurements of climate change risk.

Following Engle et al. (2020), I adopt the innovation in climate news series as indicated by the AR(1) model for use as hedge targets. The climate news series used in this section consist of two categories. The first category includes the Wall Street Journal (WSJ) climate news index and the Crimson Hexagon Negative News (CHNEG) climate news index, developed by Engle et al. (2020). These indices are monthly and cover the period from July 2008 to June 2017. To mitigate potential biases arising from reliance on a single information source, I also incorporate other climate news series from Faccini et al. (2021). They created four climate news series based on a wide array of articles from Thomas Reuters. Faccini et al. (2021)

constructed indices related to international climate summits, global warming, natural disasters, and US climate policy, respectively. The series on US climate policy and international summits capture transitional risks, while those on global warming and natural disasters aim to reflect physical risks. These news measures are available at a daily frequency from 2000 to November 2019, aggregated to monthly frequency by averaging.

I compare the hedging performance of our commodity hedging portfolios with two alternative approaches: narrative-based approaches and mimicking portfolio approaches, as outlined by Alekseev et al. (2022) and Engle et al. (2020). For the narrative-based approach, the first strategy involves taking positions in all US-listed stocks covered by the MSCI database. Stocks are ranked monthly by their E-score [14] , which is then demeaned and rescaled to range from -0.5 to +0.5. This ranking determines the stock positions in the portfolio at each time point. Another narrative-based strategy uses ETFs: long PBD (the Invesco Global Clean Energy ETF), short XLE (the Energy Select Sector SPDR Fund), and a stranded asset portfolio based on the weights of 0.3 XLE + 0.7 KOL (VanEck Vectors Coal) – SPY (SDPR S&P 500 ETF), as described by Jung et al. (2021).

The mimicking portfolio approach projects the climate risk factor, CC_t, onto a set of excess returns of portfolios, r_t:

$$CC_t = \epsilon + w^{'} r_t + e_t$$

[14] An overall environmental score for each firm is to subtract the total scores in the negative environmental subcategories from the total scores in positive environmental subcategories Engle et al. (2020), Hong and Kostovetsky (2012)

Here, the portfolio weights w are estimated in each month t using 60-month rolling window. The excess return of the hedging portfolio using mimicking portfolio approach is calculated as $h_t^{CC} = w' r_t$. Three sets of portfolios are considered to construct the mimicking portfolios (Alekseev et al. (2022)): first, using the three Fama-French factors (Mkt-rf, SMB, and HML); second, adding the two ETFs, PBD and XLE, to the Fama-French factors; and lastly, using the 30 industry-based portfolios collected from the Fama-French website.

Table 2.7 reports the hedging performance of the commodity portfolio and various alternative portfolios. I evaluate the hedging performance of different portfolios in the testing period from 2015 to 2019, given the data availability on carbon pricing. A five-year rolling window of monthly data is used to estimate the portfolio weights of the commodity hedging portfolios and the mimicking portfolios. Following Alekseev et al. (2022) and Engle et al. (2020), I compare the out-of-sample correlation between the hedging portfolios' returns and the climate news index AR(1) innovations. Each row of Table 2.7 reports the correlations of the corresponding hedging portfolio with the corresponding target, and each column represents a different climate news index. These indices are built on textual analysis from different information sources, and all indicate negative climate news (higher climate risk) with higher numbers. Therefore, a positive correlation of a hedge portfolio with the climate news AR(1) innovation demonstrates a successful hedge.

The first row of Table 2.7, "commodity," presents the hedging performance of the commodity portfolio constructed based on the estimated $_t^{Carbon}$. The excess return of the commodity hedging portfolio yields a positive out-of-sample correlation with all the climate news measures except the climate news index for "International Summits." The commodity portfolio produces a correlation of 0.217 with CHNEG (Crimson Hexagon Negative News), which is the same as the hedging performance of the equity portfolio specifically built for hedging the risk indicated by the CHNEG index by Engle et al. (2020). The high correlation of 0.206 between the commodity portfolio and the "Natural Disasters" climate news indicates a strong hedging ability of commodities against the physical risks of climate change. A main reason for this higher relationship with physical risk is that physical risks, including extreme weather events, rising temperatures, and changing precipitation patterns, can lead to significant disruptions in commodity production and transportation. Regarding the relationship between the commodity portfolio and transitional risk (specifically indicated by "US Climate Policy" and "International Summits" in Table 2.7), although the correlation of the commodity portfolio with transitional risk is not as high as with physical risk, it still shows a positive and significant correlation with the "US Climate Policy" index, implying the commodity portfolio provides an effective hedge for transitional risks as well in addition to physical risks. As transitional risks tend to influence the commodity market more gradually compared to the abrupt changes caused by physical risks, an increasing correlation between the commodity portfolio and transitional risk is expected in the future.

The second panel, "Narrative-approach," of Table 2.7 reports the performance of the narrative approach portfolios. The first row describes the performance of the narrative portfolio constructed using the E-score from MSCI. This E-score portfolio is unable to provide successful hedges for all the climate targets, as it has positive correlations with Climate Policy and International Summits only, and these correlations are of quite small magnitude. Longing PBD, a clean energy fund, provides a successful hedge against the first four climate news indices but fails for both the WSJ and CHNEG climate news indices. Shorting XLE, a fund of polluting firms, fails to provide successful hedges for most climate targets. The third row of panel b, "short stranded," is strong in hedging Global Warming and Natural Disasters climate targets but fails in hedging the WSJ climate index.

The last panel of Table 2.7 reports the out-of-sample correlation of the mimicking portfolios built specifically for each climate target. Unlike the narrative approach that produces portfolios staying constant along the rows, the rows in this panel of mimicking portfolios show in each cell the portfolio specifically built for the respective climate news series. These mimicking portfolios show various performances in hedging their climate targets. The portfolio using the three Fama-French factors to hedge against the WSJ climate index yields a high correlation of 0.204 with its climate news target, but the hedging performance is reduced by adding the two ETFs to the Fama-French three factors, decreasing the correlation to 0.106. The remaining mimicking portfolios fail to provide strong hedging performance as their out-of-sample correlations with the targets are close to zero. Overall, the results of the hedging performance show that the commodity hedging portfolio suggested in this paper delivers an effective hedge for climate news AR(1) innovation, which is superior to the alternative portfolios using other financial assets like ETFs and stocks.

2.5 Conclusion

In this paper, I investigate the impact of carbon pricing risk on commodity futures excess returns. Leveraging the investment-based asset pricing model for commodity producers proposed by Yang (2013), I hypothesize that carbon pricing risk carries a positive risk premium due to its negative effect on production costs. Using the EU ETS emission allowance prices as a proxy for carbon pricing, my analysis confirms that carbon pricing risk is a significant factor in commodity futures markets. The risk premium for carbon pricing is consistently positive across various commodity futures factor models, ranging from 0.065 percent to 0.085 percent daily.

Furthermore, the moving average of the 5-year carbon pricing risk premium, estimated using the Fama and MacBeth (1973) method, shows an upward trend, supporting the hypothesis that the risk premium is increasing over time. This trend aligns with the growing recognition of climate change risks in the financial literature (Balvers et al. (2017), Giglio et al. (2021), Ilhan et al. (2021)). Analyzing individual commodity futures reveals that commodities with high sensitivity to carbon pricing, such as crude oil, gasoil, and copper, exhibit positive loadings on carbon pricing risk. Additionally, the study finds that commodity-specific characteristics, such as basis and hedging pressure, significantly influence carbon pricing risk sensitivity.

Building on these insights, I construct a climate change hedge portfolio using commodity futures based on their carbon pricing risk loadings ($(_t^{Carbon})$). This commodity hedging portfolio demonstrates strong performance compared to alternative hedging strategies that utilize other financial assets.

The results of this research extend the scope of climate finance to include commodities, providing new avenues for managing climate change risks. These findings offer valuable implications for policymakers and investors, highlighting the importance of incorporating carbon pricing risk into their risk management and investment strategies. By understanding the role of carbon pricing in commodity futures markets, stakeholders can better navigate the financial challenges posed by climate change.

Future research may enhance the findings in this paper in several ways. Firstly, using commodity futures contracts other than the second-nearest expiry contracts, especially those with longer terms to maturity, might yield stronger empirical results regarding the carbon pricing risk premium. Additionally, measuring the basis using longer maturity commodity futures contracts might produce different results for the relationship between basis and the risk sensitivity. Secondly, investigating the channels through which carbon pricing risk affects different commodities would be an insightful extension. Considering that risk sensitivity to the carbon pricing risk premium varies across commodities, splitting commodities by different characteristics, such as liquidity, industry, or emission levels, could provide insightful results. Thirdly, as more emission trading systems are launched globally, considering a global carbon price will be important for future research. This extension could offer valuable insights into the broader implications of carbon pricing on global commodity markets. Finally, Arlinghaus 2015 found that carbon pricing in the EU could increase production costs for energy-intensive industries by approximately 5-8%. As our results show that the average of the annualized carbon pricing risk premium for all the commodities covered in the sample is more than 10%, this implies the presence of significant factors beyond the direct increase in production costs due to carbon pricing risk. Exploration of the nature of the indirect contributing factors would be worthwhile.

Table 2.1: **Summary statistics of commodity futures returns**

This table reports the summary statistics of the 35 commodities daily futures excess returns for the period from 2008:12 to 2020:12. Mean is the average excess return. Std is the standard deviation. SR is the Sharpe Ratio. Mean, Std and SR are annualized and in percentage. All the commodity futures price data are collected from Bloomberg.

Sector	Commodity	N	Mean	Std	SR
energy	Brent Crude Oil	3113	2.201	34.415	6.395
	Crude oil	3113	-0.908	38.878	-2.334
	Gasoil	3113	-0.601	30.316	-1.982
	Gasoline	3113	11.938	36.231	32.951
	Heating oil	3113	0.317	31.112	1.018
	Natural Gas	3113	-22.472	41.608	-54.010
grain	Corn	3113	-1.666	25.625	-6.503
	Kansas wheat	3113	-7.627	27.647	-27.588
	Oats	3113	7.283	28.294	25.741
	Rough rice	3113	-6.348	20.760	-30.578
	Soybean meal	3113	15.786	23.562	66.998
	Soybean oil	3113	-0.522	20.325	-2.566
	Soybeans	3113	7.847	20.414	38.440
	Wheat	3113	-7.476	28.971	-25.805
livestock	Feeder cattle	3113	1.485	17.022	8.726
	Lean hogs	3113	-6.337	28.734	-22.054
	Live cattle	3113	-0.465	16.196	-2.869
metal	Aluminum	3113	-0.777	20.404	-3.809
	Gold	3113	6.790	16.785	40.452
	Lead	3113	8.286	29.006	28.567
	Nickel	3113	7.258	31.550	23.004
	North American Copper	3113	9.715	24.162	40.206
	Palladium	3113	30.433	31.921	95.339
	Platinum	3113	2.767	23.198	11.927
	Silver	3113	10.968	31.129	35.234
	Tin	3113	10.278	24.481	41.984
	Zinc	3113	9.175	27.181	33.755
soft	Cocoa	3113	1.869	26.302	7.108
	Coffee	3113	-2.976	30.733	-9.682
	Cotton	3113	8.053	24.383	33.027
	Ethanol	3113	12.073	24.489	49.299
	Lumber	3113	10.586	31.405	33.708
	Milk	3113	4.661	21.683	21.496
	Orange juice	3113	5.683	30.429	18.677
	Sugar	3113	0.328	30.179	1.088
All		3113	3.656	27.129	14.439

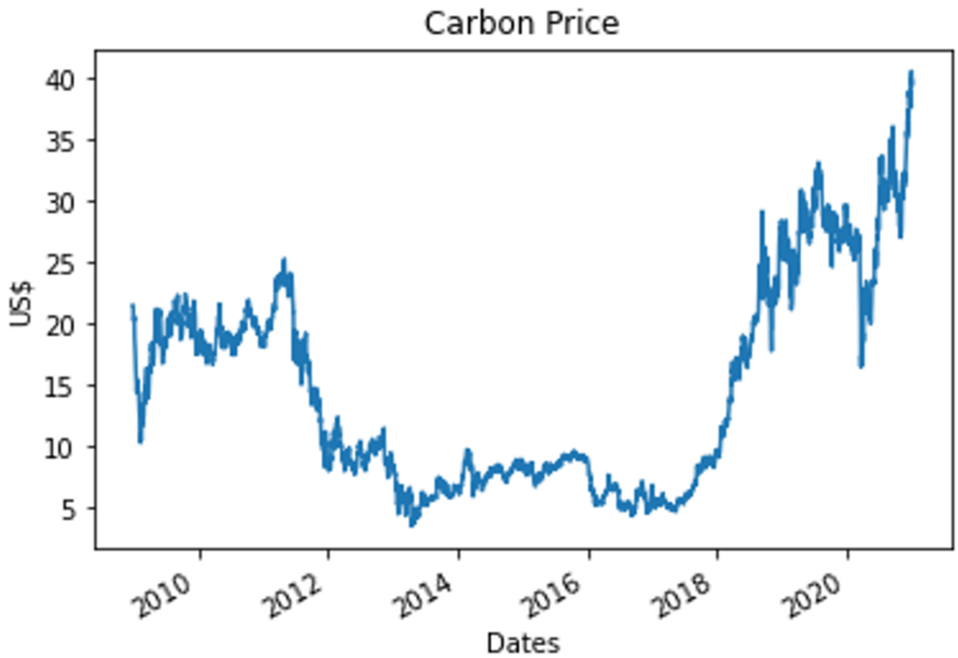

Figure 2.1: **Carbon Price**
The figure plots the carbon price, Carbon Emissions Allowance price (EUA), over the sample period (2008:12 to 2020:12). The EUA daily price data is collected from Bloomberg (ticker: MO1 Comdty).

Table 2.2: **Summary statistics of portfolios sorted by commodity characteristics**

This table presents the descriptive statistics of the commodity futures returns sorted by five characteristics of commodity futures: basis, momentum, basis-momentum, value, and volatility. Five portfolios are constructed based on each of the sorting characteristics. The portfolio returns in each quintile are calculated as equally average of all commodities daily futures returns within a portfolio. The HML portfolio returns are the returns of highest portfolio (High) minus the returns of the lowest portfolios (Low). The portfolio means, standard deviation and Sharpe Ratio are annualized.

	Low	P2	P3	P4	High	HML
Basis						
Mean	-3.072	-3.160	-2.942	11.274	8.893	12.342
Std	16.442	15.535	16.165	16.073	14.716	15.993
Sharpe						
ratio	-18.682	-20.340	-18.201	70.144	60.432	77.176
t-stats	-0.637	-0.694	-0.620	2.232	1.944	2.444
Momentum						
Mean	6.371	-2.011	0.377	1.395	4.172	-2.068
Std	19.384	15.394	14.068	14.612	15.816	19.824
Sharpe						
ratio	32.870	-13.064	2.680	9.544	26.376	-10.433
t-stats	1.070	-0.443	0.090	0.318	0.868	-0.354
Basis-Momentum						
Mean	-4.834	3.230	-0.322	4.009	8.503	14.011
Std	15.593	15.467	15.938	16.924	15.450	16.181
Sharpe	-31.002	20.882	-2.020	23.687	55.034	86.590
ratio	-1.067	0.690	-0.068	0.780	1.773	2.721
t-stats						
Value	-1.950	-2.242	0.130	7.335	7.086	-8.440
Mean	-15.722	14.678	16.074	15.386	17.216	17.941
Std	15					
Sharpe	-12.400	-15.273	0.807	47.672	41.159	-47.043
ratio	-0.420	-0.519	0.027	1.545	1.335	-1.650
t-stats						
Volatility						
Mean	1.125	4.067	-2.566	6.388	1.133	0.008
Std	15.558	15.950	15.200	16.659	14.632	13.280
Sharpe						
ratio	7.231	25.498	-16.881	38.345	7.743	0.059
t-stats	0.241	0.839	-0.574	1.248	0.258	0.002

Table 2.3: **Summary statistics of carbon price**

This table reports the summary statistics of the carbon price. The price of Carbon Emission Allowance price is used as proxy of carbon price level following Choi et al. (2020). The generic 1st EUA (EU allowance) futures data is collected from Bloomberg (Bloomberg symbol:MO1 Cody). Panel A reports the descriptive statistics, the first-order autoregressive coefficient (AR(1)) and the test statistics for an ADF (augmented Dickey-Fuller (Dickey and Fuller, 1979)) test with trend for the null hypothesis of a unit root. Std. dev., is the standard deviation, skew., is the skewness and kurt., is the kurtosis. Carbon pricing shock is defined as the change rate of carbon price since the carbon price series is integrated of order $1, I(1)$. The sample period starts from 2008:12 to 2020:12. Panel B reports the unconditional correlation between carbon pricing shock and excess returns of commodity portfolios. "*, **, ***" indicate significant at the $10\%, 5\%$ and 1% level, respectively.

Panel A: descriptive statistics of carbon price

	mean	Std. dev.	skew.	kurt.	AR(1)	ADF test with trend
Carbon Pricing level	14.604	8.605	0.749	-0.027	0.995***	0.638
Carbon pricing shock	0.065	3.3	0.090	11.008	-0.005	−14.483***

Panel B: unconditional correlation between portfolios returns and carbon pricing shock

	Basis	Momentum	Basis Mom	Value	Volatility
Low	0.192***	0.213***	0.172***	0.170***	0.209***
P2	0.202***	0.236***	0.196***	0.173***	0.189***
P3	0.211***	0.208***	0.219***	0.198***	0.203***
P4	0.206***	0.161***	0.209***	0.189***	0.201***
High	0.170***	0.156***	0.187***	0.238***	0.202***
HML	-0.05	−0.09***	0.018	0.099***	-0.026

Table 2.4: **Cross-section Asset pricing tests**

This table reports average slopes, and sample sizes for the Fama and MacBeth 1973 two-pass OLS regression with 25 characteristics-sorted commodity portfolios as the test assets from December 2008 to December 2020. The risk premium associated with each factor is reported with t-statistics given in parentheses. The t-statistics are adjusted using a Newey-West correction.Twenty-four months of daily data are used to estimate the factor sensitivities in the first pass, and these sensitivities are used to obtain risk premia in the second pass. The models estimated are as follows:

$$2F : r_{i,t} = \alpha_i + m_i AVG_t + c_i CARRY_t + \epsilon_{it}$$

$$2F + Carbon : r_{i,t} = \alpha_i + m_i AVG_t + c_i CARRY_t + carbon_i CarbonShock_t + \epsilon_{i,t}$$

$$3F : r_{i,t} = \alpha_i + m_i AVG_t + c_i CARRY_t + mom_i MOM_t + \epsilon_{i,t}$$

$$3F{+}Carbon : r_{i,t} = \alpha_i{+}m_i AVG_t{+}c_i CARRY_t{+}mom_i MOM_t{+}carbon_i CarbonShock_t{+}\epsilon_{i,t}$$

$$4F : r_{i,t} = \alpha_i + m_i AVG_t + c_i CARRY_t + mom_i MOM_t + bm_i Basis_M om_t + \epsilon_{i,t}$$

$$4F{+}Carbon{:}r_{i,t}{=}\alpha_i{+}m_i AVG_t{+}c_i CARRY_t{+}mom_i MOM_t{+}carbon_i CarbonShock_t{+}carbon_i CarbonShock_t{+}\epsilon_{i,t}$$

Where $r_{i,t}$ is the daily excess return in percentage on commodity I at time t and AVG_t is the average excess return of all commodities, CARRY is the difference between the returns on equal weighted portfolios of commodities with highest basis and lowest basis. MOM is the difference between the returns on equal weighted portfolios of commodity with highest past returns and lowest past returns. $Basis_M om$ is the difference between the returns on equal weighted portfolios of commodity with highest basis-momentum and lowest basis-momentum where basis-momentum (Boons and Prado (2019)) is defined as the difference between the first- and second- nearest futures strategies. Carbon pricing shock is the daily change rate of carbon price series.

	2F	2F+ Carbon	**3F**	3F+ Carbon	**4F**	4F+ Carbon
CONST	-0.0008	0.0149	-0.0061	0.0029	-0.0097	-0.0010
	(−0.0369)	(0.6291)	(−0.2245)	(0.1079)	(−0.3520)	(−0.0357)
AVG	-0.0023	-0.0180	0.0030	-0.0061	0.0066	-0.0021
	(−0.0902)	(−0.6996)	(0.1017)	(−0.2107)	(0.2231)	(−0.0725)
CARRY	0.0340	0.0321	0.0363	0.0345	0.0233	0.0223
	(1.4192)	(1.3643)	(1.5283)	(1.4761)	(0.9866)	(0.9524)
MOM			0.0019	0.0011	-0.0025	-0.0020
			(0.0710)	(0.0412)	(−0.0930)	(−0.0747)
Basis$_M om$					0.0493	0.0480
					(2.3828)	(2.3729)
CarbonPrice shock		0.0652		0.0747		0.0856
		(2.0864)		(2.3305)		(2.6549)
N	2608	2608	2608	2608	2608	2608

Table 2.5: **Factor loading of individual commodity**

The factor loading of the individual commodity is inferred from the four-factor model with carbon price shock: $r_{i,t} = \alpha_i + m_i AVG_t + c_i CARRY_t + mom_i MOM_t + carbon_i CarbonShock_t + carbon_i CarbonShock_t + \epsilon_{i,t}$. The significant factor loadings on carbon price shock at the 5% significance level are in bold.

Sector	Commodity	CONST	AVG	CARRY	MOM	Basis Mom	Carbon $_pricing$	R2
energy	Brent Crude Oil	0.000	1.663	-0.064	-0.408	0.152	**0.021**	0.548
	Crude oil	0.000	1.815	-0.181	-0.474	0.246	**0.027**	0.554
	Gasoil	0.000	1.250	-0.042	-0.277	0.123	**0.042**	0.405
	Gasoline	0.000	1.576	-0.171	-0.360	0.361	**0.022**	0.478
	Heating oil	0.000	1.531	-0.025	-0.345	0.077	**0.017**	0.539
	Natural Gas	-0.001	0.782	-0.090	-0.203	-0.487	-0.005	0.134
	Corn	0.000	1.197	-0.004	0.109	-0.131	**−0.024**	0.313
	Kansas wheat	0.000	1.203	-0.047	0.124	-0.370	**−0.017**	0.310
	Oats	0.000	1.073	0.040	0.110	-0.119	**−0.018**	0.213
	Rough rice	0.000	0.518	-0.077	0.080	-0.102	**−0.009**	0.101
	Soybean meal	0.000	0.885	0.153	0.071	0.005	-0.023	0.222
	Soybean oil	0.000	1.006	0.031	-0.006	0.022	-0.011	0.371
	Soybeans	0.000	1.001	0.100	0.052	0.001	**−0.018**	0.366
	Wheat	0.000	1.281	-0.076	0.144	-0.341	**−0.017**	0.315
livestock	Feeder cattle	0.000	0.208	0.048	-0.037	0.011	-0.003	0.028
	Lean hogs	0.000	0.574	0.200	-0.074	-0.290	-0.004	0.085
	Live cattle	0.000	0.350	0.041	-0.018	-0.011	-0.005	0.076
metal	Aluminum	0.000	0.895	0.049	0.061	0.011	**0.013**	0.311
	Gold	0.000	0.538	-0.005	0.114	0.027	-0.003	0.165
	Lead	0.000	1.262	0.054	0.128	0.093	-0.001	0.311
	Nickel	0.000	1.430	0.037	0.105	0.051	0.008	0.330
	North American Copper	0.000	1.246	0.035	0.043	0.157	**0.017**	0.463
	Palladium	0.001	1.375	-0.069	0.240	0.203	0.014	0.324
	Platinum	0.000	1.007	-0.017	0.121	0.090	0.008	0.301
	Silver	0.000	1.340	-0.009	0.245	0.102	0.002	0.298
	Tin	0.000	0.967	0.019	0.058	0.199	0.006	0.273
	Zinc	0.000	1.262	0.063	0.117	0.097	0.002	0.356
	soft	0.000	0.608	0.034	0.003	0.036	**0.018**	0.095
	Cocoa	0.000	0.983	-0.121	0.084	-0.242	-0.012	0.174
	Coffee	0.000	0.827	0.080	0.038	0.059	0.005	0.190
	Cotton	0.000	0.996	-0.104	0.055	0.088	**−0.015**	0.266
	Ethanol	0.000	0.633	-0.077	0.040	0.049	0.009	0.066
	Lumber	0.000	0.230	-0.016	0.019	0.059	-0.010	0.019
	Milk	0.000	0.490	0.106	0.020	-0.067	0.003	0.042
	Orange juice	0.000	0.996	0.106	0.023	-0.157	**−0.037**	0.163
	Sugar							

TABLE 2.6: **Carbon pricing shock beta and commodity characteristics**

This table reports the results for the regression tests of the characteristics determinants of the carbon price shock beta. The dependent variable is β_{carbon}, which is estimated using 35 individual commodities as test asset in the first pass regression. All independent variables are lagged. The basis for a commodity is calculated as the difference between the first- and second-nearest prices of futures, $Basis_{i,t} = \frac{log(F_{i,t,1})-log(F_{i,t,2})}{(D_2-D_1}$, where $F_{i,t,1}$ is the futures price of commodity i on day t for the first nearby contract for which data are available and D_1 is the number of days to maturity on this contract. $F_{i,t,2}$ is the futures price of commodity i on day t for the second nearby contract for which data are available and D_1 is the number of days to maturity on this contract. Following De Roon et al. (2000) Hedging pressure for commodity is caluated as $\frac{numberofshorthedgepositions-numberoflonghedgepositions}{totalnumberofhedgepositions}$. The data for hedging pressure calculation is semimonthly data of the positions of large traders collected from CFTC (Commodity Futures Trading Commission). "*, **, ***" indicate significant at the 10%, 5% and 1% level, respectively.

	β_{carbon}	β_{carbon}	β_{carbon}	β_{carbon}
const	0.000	0.000	-0.0050	-0.008
	(-0.6118)	(-0.6049)	(-42.169)	(-29.558)
$Basis_{t-1}$	-0.7553^{***}	-1.187^{***}		
	(-6.1519)	(-5.7520)		
HP_{t-1}			0.0170^{***}	0.0053^{***}
			(37.215)	(8.7293)
Time FE	No	Yes	No	Yes
Entity FE	No	Yes	No	Yes
R2	0.0004	0.3167	0.0201	0.2920
N	91349	91349	67363	67363

Table 2.7: **Climate Hedge Performance**

This table shows the monthly correlation between various climate change hedge portfolios and the AR(1) innovations in a broad range of climate change news index from 2015 to 2019. Each row presents the correlations for one hedge portfolio return and positive correlations are highlighted in bold. Commodity, the target portfolio, represents the hedge portfolio built based on the carbon pricing risk exposure using commodity futures. The remaining rows report the correlations between different benchmark portfolios created by following Alekseev et al. (2022) and Engle et al. (2020). US climate policy, International summits, Global warming and Natural disasters are the climate new indices created by Faccini et al. (2021) based on a broad of global news article from Thomas Reuters. WSJ and CHNEG are the news indices constructed by Engle et al. (2020). All the climate indices are created to indicate negative climate news by higher number; hence the positive correlation indicate a successful hedge.

	US climate policy	International summits	Global warming	Natural disasters	WSJ	CHNEG
Commodity	**0.052**	-0.036	**0.014**	**0.206**	**0.081**	**0.217**
Narrative Method						
esg_weighted	**0.010**	**0.010**	-0.090	-0.067	-0.301	-0.295
long_PBD	**0.048**	**0.104**	**0.084**	**0.186**	-0.027	-0.049
short_XLE	-0.113	-0.032	-0.103	-0.028	-0.111	**0.074**
short_stranded	**0.050**	-0.030	**0.154**	**0.296**	-0.014	**0.122**
Mimicking Portfolio						
FF3 factor	-0.082	**0.011**	-0.006	**0.021**	**0.204**	-0.274
ETF	-0.068	**0.009**	-0.016	**0.075**	**0.106**	-0.147
Industry	**0.007**	**0.028**	-0.052	-0.019	**0.083**	**0.036**

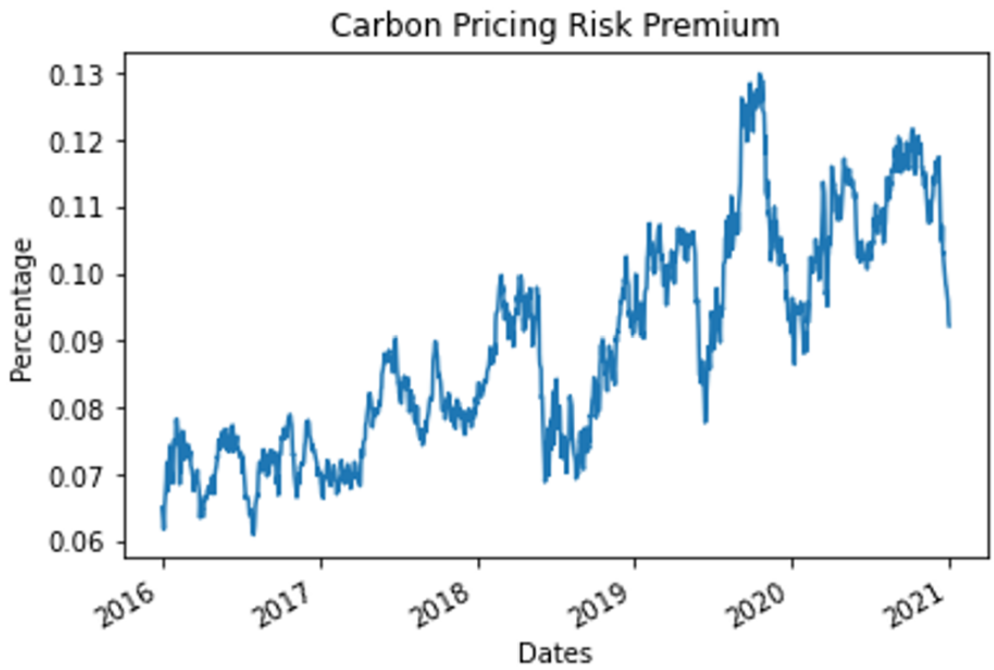

FIGURE 2.2: **Carbon pricing risk premium**
shows the 5-year moving-average for the carbon pricing risk premium for the 4F+Carbon model as described in Table 2.4

Chapter 3

Are analysts' forecasts reliable? A machine learning-based analysis of the target price accuracy

3.1 Introduction

Security analysts' forecasts have garnered considerable attention from both investors and academics. The role of security analysts (or "sell-side analysts") in forecasting the earnings and stock prices of publicly traded companies is crucial to the capital market. As information providers, security analysts decide which companies to follow, the number of reports to issue for these companies, and the investment recommendations to make. As pointed out by prior studies (e.g., Bradshaw et al. (2013)), the forecasts and recommendations issued by security analysts significantly affect the capital market from at least three perspectives. First, investors, especially those with limited time or ability to analyze individual securities or stocks, often rely on analysts' reports to make investment decisions. Second, security analysts can influence the management of publicly traded companies through their interactions with corporate managers. These managers must understand the information needs of security analysts and their processing methods to aim for positive forecasts. Finally, regulators and academics show keen interest in the work of security analysts, who are pivotal contributors to the capital market's information flow. Their research extensively utilizes financial statement analyses and stock recommendations issued by analysts.

An equity research report issued by security analysts typically includes three elements: earnings forecasts, stock recommendations and target price forecasts. Compared to recommendations (e.g., buy, hold, sell) which are discrete, target prices provide investors with well-defined horizons and specific investment signals regarding firm values. Although the credibility and usefulness of target prices have previously been criticized[1], a significant number of research efforts continue to concentrate on the topic of target prices. Numerous studies demonstrate that target prices offer independent investment value distinct from earnings forecasts and stock recommendations. Investors can determine expected returns over a specific horizon (e.g., 12-month) by relying solely on target prices, without requiring other forecast information such as earnings forecasts and investment recommendations (Brav and Lehavy (2003); Asquith et al. (2005); Da and Schaumburg (2011); Gleason et al. (2013)). However, the accuracy of target price forecasts is questionable,

[1]Morgenson (2001): "Price targets are hazardous to investors' wealth," https://www.nytimes.com/2001/08/05/business/market-watch-price-targets-are-hazardous-to-investors-wealth.html

as some studies indicate that these forecasts are biased (Ottaviani and Sørensen (2006); Bonini et al. (2010)). Regarding the causes of analysts' bias, Lee et al. (2008) suggest that analysts and managers adopt a bounded rationality premise, which contrasts with the rational expectation hypothesis in relation to analyst forecasts. They find that neglecting the business cycle contributes to analysts' bounded rational behavior. Additionally, analysts are also found to exaggerate target prices to cater to investors (Chen et al. (2016)).

Additionally, some studies have directly pointed out that target prices are not as accurate as expected. For instance, Bradshaw et al. (2013) find that fewer than half of target prices were met, both at the end of and during the forecast period. They conclude that analysts do not demonstrate persistent differential ability in forecasting target prices. Therefore, it is crucial for investors to identify accurate target prices, or in other words, to predict whether a target price will be met within the next 12 to 18 months. However, few studies have explored this area, which primarily motivated our research to predict the accuracy of target prices.

Prior studies on target prices have identified many important factors, beyond analysts' ability, that affect their accuracy. Some studies find that market sentiment has a significant impact on analyst optimism and target price accuracy (Qian (2009); Clarkson et al. (2020); Buxbaum et al. (2023)). Target prices are found to be closer to intrinsic values when investor sentiment is low (Buxbaum et al. (2023)). Bilinski et al. (2013) find that country-specific characteristics, such as the origin of the legal system, cultural traits, and IFRS regulation, can explain variations in target price accuracy. Furthermore, firm-specific indicators like fundamental information and past stock prices are also significant. Kerl (2011) finds that target price accuracy is negatively associated with analyst optimism, volatility, and price-to-book ratio and positively associated with size. Cheng et al. (2019) suggest that strong corporate governance is positively associated with target price accuracy. Clarkson et al. (2020) find that the 52-week high price has a significant influence on forecast errors. Palley et al. (2023) document a negative relationship between forecasted and actual returns for those stocks with high target price dispersion, which they attribute to the delayed update of target prices by some analysts after bad news. He and Li (2024) find that media coverage increases the business risk (earnings volatility) and information risk (bid-ask spread) of firms, thereby enlarging the forecast error. Finally, analyst-specific characteristics are also pivotal. Bradshaw et al. (2013) observe that target price optimism is positively related to analysts' conflicts of interest. Similarly, Frankel et al. (2006) confirm that the analyst reports are more informative when their potential brokerage profits are higher. Green et al. (2014) argue that conference-hosting brokers provide more informative and accurate forecasts. In summary, the determinants of target price accuracy are multifaceted, including the state of the capital market, industry-specifics, firm-specific information, and analyst-specific factors. This paper aims to precisely predict target price accuracy by incorporating all known determinants and possible factors, based on market-level, firm-level, and analyst-level information sets.

Among prior studies on analyst forecasts, our study on predicting target price accuracy is closely related to previous research on forecast error. The traditional approach to examine whether analysts' forecasts of earnings are accurate typically involves the estimation of ordinary least squares regressions, which implicitly assumes that analysts strive to minimize their mean squared forecast errors. However, several studies have questioned this assumption, proposing alternative viewpoints on the loss function used to examine analysts' forecast errors (Lambert (2004); Markov and Tan (2006); Clatworthy et al. (2012)). Our focus in this paper is not on identifying the loss function that best represents analysts' incentives. Instead, the controversy over the statistical methods used to examine analyst forecast errors has motivated us to employ more flexible techniques beyond the traditional approach (i.e., linear regression) that relies on strong assumptions about analysts' forecasts. To incorporate all influential factors on target price for better prediction performance, we naturally consider using machine learning techniques, which are adept at handling high dimensionality and complex dependencies among predictors. As a branch of artificial intelligence, machine learning techniques are advantageous because they can identify patterns too complex for human detection, make predictions based on larger datasets, and adapt to changes in these datasets. The application of machine learning in stock market prediction dates back to the early 2000s, exemplified by Jasic and Wood (2004), who developed an artificial neural network to predict daily stock market index returns. Over the past two decades, the increasing availability of data, reductions in data storage costs, and advancements in computer processing speeds have made machine learning applications increasingly popular in finance (Karolyi and Van Nieuwerburgh (2020)). Our paper connects to the growing body of literature on machine learning in finance by using these tools to predict the accuracy of target prices. By leveraging machine learning techniques, we aim to uncover complex behavioral patterns of analysts when forecasting stock prices.

Our paper extends prior research on target prices in three ways. Our first contribution lies in predicting the accuracy of target prices issued by sell-side analysts. While target prices have garnered significant attention from academics, most previous studies explore factors influencing analysts' target prices. However, forecasting target price accuracy is also crucial. First, analysts' target prices are not as accurate as expected. In our sample, only 36.33% of the target prices are met. Second, for investors who base their investment decisions on target prices, predicting the accuracy of target prices is of immense importance.

Our second innovation involves employing machine learning techniques for our predictions, incorporating all available influential factors on target prices to make the most accurate predictions. Specifically, we use ensemble learning methods. Ensemble learning is a technique that combines several models to improve predictive performance. It can automatically select the most relevant features for prediction, in contrast to traditional econometric methods, which require manual selection of independent variables. Unlike traditional econometric methods that focus on a single "best" model, ensemble methods consider an inventory of many models and "average" them to produce the final prediction. Thus, we rely on ensemble learning to incorporate valuable information for better prediction performance.

Finally, to assess the investment benefits of our predictions, we establish long-short portfolios based on the prediction results of target price achievement and examine their out-of-sample returns. These equally weighted portfolios demonstrate the economic benefit of our approach.

The rest of the paper is arranged as follows. Section 3.2 presents the sample and research methods followed by the empirical results and economic significance analysis in Section 3.3. Section 3.4 concludes the paper.

3.2 Data and methodology

3.2.1 Target price

Target prices[2] of stocks listed in the U.S. in the period between 1999[3] and 2021 are obtained from the I/B/E/S (Institutional Brokers' Estimate System) database. Among different horizons for target price, we consider only 12-month target prices. Since the analyst report date is discrete and daily data are not feasible to create portfolios, we downsample the data from daily to monthly. Following Bradshaw et al. (2019) and Pursiainen (2022), we exclude observations for which the target price to current price ratio is below 0.7, above 4, or equal to 1. We also exclude observations without 12-month-ahead closing prices and those with missing values or extreme values of the features. Table 3.1 shows the details of the sample size. In our sample, 83.03% of the observations have a target price above the current price.

A target price reflects an analyst's degree of optimism towards a stock. An informative measure of this optimism is the distance between the target price and the current price (Bradshaw et al. (2019)). Some studies (e.g., Da and Schaumburg (2011), Da et al. (2016), Bradshaw et al. (2019), Hao and Skinner (2022)) introduce

[2] According to Hao and Skinner (2022), most analyst reports state that the forecasted 12-month dividends have been discounted in the 12-month target prices. The target prices do not include dividends to be paid within a year. Moreover, the dividend payout ratio has been included in the firm characteristics.

[3] The I/B/E/S database provides target price data dating back to July 1999.

the concept of implied return of the target price. A higher implied return indicates greater optimism regarding the stock price. In this paper, we define the implied return of target price as

$$TPP^{j}_{i,t,t+12} = ln(\frac{TP^{j}_{i,t+12}}{P_{i,t}}), \tag{3.1}$$

where $TP^{j}_{i,t+12}$ is the target price of stock i issued by analyst j with 12-month forecast horizon $t+12$, $P_{i,t}$ is the current price of stock i. At the end of 12-month forecast horizon, $t+12$, we obtain the target price forecast error and determine whether the target price is achieved. Table 3.2 presents the statistics of the implied returns of target prices and the actual returns.

To better illustrate the fact that analysts' target price forecasts are not always accurate and that naïve investment strategies blindly following these target prices are likely to result in losses, we plot the relationship between the implied returns of target prices and the actual realized returns at the end of the forecast period in Figure 3.1. This histogram displays the implied returns based on target price observations in our sample and the realized returns at the end of the forecast period for the predicted stocks. We use color brightness to represent frequency: the brighter the color, the higher the frequency of occurrence. While most of the implied returns of target prices are contained within the 0 to 0.2 range, actual returns seem to be randomly distributed across the -0.2 to 0.3 range. Figure 3.1 reveals two important facts about target price forecasts. First, many optimistic forecasts with positive implied returns ultimately result in negative realized returns. Therefore, if investors blindly trust these analysts' forecasts and invest the stocks with optimistic forecasts accordingly, they may eventually face financial losses. Second, some stocks with neutral target price forecasts, indicating zero implied return, exhibit positive realized returns at the end of the forecast period. This inaccuracy in target price forecasts can cause investors to miss out on investment opportunities with positive returns. The information shown in Figure 3.1 underscores observations underscore the significance of our work in predicting the accuracy of target price forecasts.

In this paper, we use target price forecast error and target price achievement as proxies to measure the accuracy of target prices. Target price forecast error, defined as the deviation of the "realized" price (the stock price 12 months after the target price announcement) from the target price, has been widely used to assess target price accuracy in previous studies (e.g., Buxbaum et al. (2023); Dechow and You (2020); Chen et al. (2016); Bonini et al. (2010); Demirakos et al. (2010); Bradshaw et al. (2013); Kerl (2011); Bilinski et al. (2013)). Target price forecast error can be written as

$$\ln(P_{i,t+12}/TP^{j}_{i,t+12}) = R_{i,t,t+12} - TPP^{j}_{i,t,t+12} = T(X^{j}_{i,t}, \beta_1) + \epsilon_{1t}, \tag{3.2}$$

where $R_{i,t,t+12}$ is the logarithmic return of stock i, $TPP^{j}_{i,t,t+12}$ is the logarithmic implied return of stock i based on the target price issued by analyst j, $X^{j}_{i,t}$ is a series of features at t, $T(\cdot)$ is the ensemble method of decision trees for regression.

In addition, we introduce absolute target price forecast error to examine the exact amount by which the target price is off. Absolute target price forecast error is defined as

$$|\ln(P_{i,t+12}/TP^{j}_{i,t+12})| = |R_{i,t,t+12} - TPP^{j}_{i,t,t+12}| = T(X^{j}_{i,t}, \beta_2) + \epsilon_{2t}. \tag{3.3}$$

We denote the forecast error, as outlined in Eq. (3.2), as Error I, and the absolute forecast error, as in Eq. (3.3), as Error II. Figure 3.2 illustrates the distribution of the forecast errors. The negative skewness and high kurtosis of Error I indicate that a significant number of forecasts ended with stock prices substantially lower than the predicted target prices. This implies that most analysts' target price forecasts are too optimistic, overpredicting stock prices at the end of the forecast horizon. Table 3.3 exhibits a statistical summary of the forecast errors. For observations where the target price exceeds the current price (positive implied return), the negative mean implies that most of the "positive" target price forecasts are not realized at the end of the forecast horizon. This supports prior studies concluding that analysts are overly optimistic in their forecasts. Conversely, for observations with a target price below the current price (negative implied return), the positive mean error suggests analysts tend to overestimate the decline in stock prices, highlighting a potential bias in downward price predictions. Additionally, the statistics for Error II show that the magnitude of the target price forecast error averages above 30%, which further confirms the inaccuracy of target price predictions.

In addition to target price forecast errors, target price achievement is another measurement of the accuracy. Target price achievement is defined by a dummy variable with the value of one if, for target prices that are higher (or lower) than the current price, the 12-month-ahead closing price is higher (or lower) than the target price, and zero otherwise. In our sample, 36.33% of the target prices are achieved at the end of the forecast period. Figure 3.3 presents the target price achievement rate by year in our sample, along with the average achievement rate for each calendar month. The achievement rate tends to decrease during financial crises, such as the dot-com bubble, the global financial crisis, and the COVID-19 pandemic. However, it does not exhibit a consistent pattern across different calendar months. Figure 3.4 shows the target price achievement rates (the number of achieved target prices divided by the total number of target prices) and the next twelve months (NTM) return of the S&P 500 Index by month from 1999 to 2021. The achievement rates and the S&P 500 Index return present an extremely high correlation, suggesting that, on average, whether a target price can be achieved mainly depends on the market performance.

Figure 3.5 illustrates the relationship between the target price and forecast accuracy. We find that the larger the deviation of the target price from the current price, the greater the forecast error and the lower the achievement rate.

3.2.2 Ensemble methods, sampling and hyperparameters

To predict the target price accuracy, we employ ensemble methods of decision trees, which use a combination of decision trees to improve prediction accuracy. Decision trees are widely used non-parametric supervised learning models that essentially learn a hierarchy of if/else questions, leading to a decision. Random forest and gradient boosting are two popular ensemble methods of decision trees. Random forest is a robust algorithm that provides relatively fast predictions compared to other complex algorithms, such as deep learning and support vector machines. Gradient boosting is known for its high accuracy. The main difference between random forest and gradient boosting is that random forest selects trees randomly, while gradient boosting is an iterative method that builds a sequence of decision trees, where each tree tries to correct the error of the previous one. Both methods are flexible for classification and regression problems and have advantages over traditional econometric methods such as ordinary least squares (OLS). They can capture non-linear relationships between the independent and dependent variables and work well in handling large datasets, missing data, outliers, and a mixture of categorical and numerical data. Ensemble learning methods can automatically select the most relevant features for prediction, while traditional econometric methods require the users to manually select the independent variables. In this paper, we implement regression models to predict forecast error and classification models to predict target price achievement. The algorithms for the ensemble methods are detailed in Appendix C1. We report accuracy rates and feature importances derived from Mean Decrease in Impurity (MDI) in Section 3.3.

Addressing overfitting is crucial. Ensemble methods inherently act as a form of regularization, making them less prone to overfitting. Random forest uses bootstrapping and feature randomness to reduce variance, while gradient boosting uses sequential learning and can be adjusted with shrinkage to prevent overfitting. For both regression and classification models, we randomly split the dataset into training and test sets to ensure evaluation against an unseen dataset, thereby assessing the model's generalization capability. The size of the test set was set to 0.25.

Throughout the implementation of the random forest and gradient boosting models, hyperparameters are diligently optimized to enhance predictive accuracy and model robustness. Hyperparameter tuning also helps address overfitting by optimizing the model's parameters to find the right balance between underfitting and overfitting, ensuring the model is neither too complex nor too simplistic for the underlying data pattern. This selection involves a systematic search and fine-tuning process, considering factors such as maximum depth, number of estimators, number of features and learning rate, yielding 48 candidate models for random forest and 27 for gradient boosting. We implement 3-fold cross-validation, splitting the training set randomly into 3 subsets, to select the optimal hyperparameters. Specifically, each candidate model undergoes three rounds of training and evaluation, using two subsets for training and the remaining one for evaluation. Details on the selection and tuning of hyperparameters are presented in Table 3.4.

3.2.3 Market-, firm-, and analyst-level features

Table 3.5 lists the 56 features selected in our models, including 11 market-level features, 41 firm-level features, and 4 analyst-level features. In analyzing firm-level features, we primarily follow the methodologies outlined in the studies by Chen et al. (2023) and Kaniel et al. (2023). We select 41 variables derived from accounting data in the Compustat database or stock prices in the Center for Research in Security Prices (CRSP). Variables updated annually are refreshed each June in alignment with the Fama-French convention, while those that change monthly are updated at the end of each month for use in the subsequent month. We normalize all firm-level features across the dataset to fall within a range of -0.5 to 0.5, based on each stock's ranking in relation to the specific feature. Analyst-level data are sourced from the I/B/E/S database. Additionally, we collect market-level data, including the S&P 500 Index returns from CRSP, the VIX Index from the Chicago Board Options Exchange (CBOE) website, bond yields from the Federal Reserve's official website, and the sentiment index from Jeffrey Wurgler's website.

3.3 Empirical results

3.3.1 Target price forecast error

The regressions for Eq. (3.2) are estimated with ensemble methods. Table 3.6 shows the coefficient of determination of the prediction (R^2), Mean Absolute Error (MAE), and Root Mean Squared Error (RMSE) for the regressions. Gradient boosting outperforms random forest in prediction accuracy, yielding higher R^2 values and lower MAE and RMSE for Error I. Conversely, random forest excels in predicting Error II. Figure 3.6 presents the feature importances of the models. We use different colors to indicate the categories of features. The firm/stock-specific predictor variables are shown in blue, while the market-specific and analyst-specific predictor variables are indicated in green and orange, respectively. The upper panel displays the feature importance for predicting Error I, and the lower panel illustrates the feature importance for predicting Error II. Notably, the implied return of target price is the most important feature. Intuitively, target price forecasts with high implied returns are more likely to be assigned to growth stocks that demonstrate high growth potential. Such growth stocks typically carry greater risks, or high volatility, which makes it more challenging for analysts to make accurate target price forecasts. Closeness to past year high (rel2high), which is the ratio of the stock price at the end of the previous calendar month and the highest daily price in the past year, exhibits high importance scores in the models. Its significant and substantial importance in predicting target price forecast errors aligns with the findings of George and Hwang (2007). They argue that when good (bad) news pushes a stock's price near (far from) a reference point (e.g., the 52-week high), investors become hesitant to bid the price higher (lower), even if the information justifies it. This behavior is consistent with the "adjustment and anchoring bias" where traders use the 52-week high as an anchor while adjusting stock values in response to new information. Beyond the implied return and closeness to past year high, the other features do not show as much importance. Market-level features, such as the sentiment index, TED spread and 60-day VIX, are more pivotal in gradient boosting compared to random forest.

3.3.2 Target price forecast achievement

The same features are utilized for predicting the achievement of target prices using classification models. Table 3.7 details the models' prediction accuracy. The random forest classification model achieves an accuracy rate of 88.7% on the test set. In other words, 88.7% of the out-of-sample predictions are correct at the end of the 12-month period. The gradient boosting classification model achieves an accuracy rate of 89.0% on the test set. These high accuracy rates underscore the significant predictability of target price achievement by the end of the forecast period. A small gap between training and testing performance indicates good generalization. The balanced accuracy score, presented in the last row, addresses the issues caused by the imbalanced datasets. It is defined as the average of sensitivity and specificity[4], measuring the average accuracy achieved across both minority

[4]Sensitivity: the "true positive rate," or the percentage of positive cases the model can detect; specificity: the "true negative rate," or the percentage of negative cases the model can detect.

and majority classes. The balanced accuracy score is 86.6% for random forest and 87.3% for gradient boosting, indicating a commendable balance between sensitivity and specificity. Figure 3.7 displays the feature importances of the models. Consistent with the regression models for forecast errors, the implied return of target price is the most crucial feature in the classification models. Likewise, closeness to past year high (rel2high) has high importance scores. However, market-level features demonstrate lower importance in the random forest classification model.

For robustness check, we present confusion matrices in Figure 3.8. The main purpose of the confusion matrices is to provide a detailed breakdown of the models' performance, indicating the number of instances that are correctly or incorrectly classified for each class. This approach provides a more thorough evaluation than merely using accuracy scores, especially in scenarios with unbalanced achievements. Both the random forest and gradient boosting models are highly robust in identifying target prices that cannot be achieved at the end of 12 months, with accuracy rates above 90%. While they are somewhat less precise in pinpointing achievable target prices, the accuracy rates are still around 80%. These outcomes further endorse the models' effectiveness. For buy-side analysts, the models yield critical insights into the feasibility of target prices. Investors, in turn, can use these models to sidestep unattainable target prices and make more informed investment decisions.

3.3.3 Portfolio

To explore the economic significance of our prediction for investment, we apply the prediction results to create equally weighted long-short portfolios. Keeping the chronological order, we estimate the classification models on an expanding window. The fundamental concept of constructing the hypothetical zero-cost portfolios involves selecting specific stocks each month, based on their predicted target price achievements as determined by the classification models. The returns generated from the portfolios are essentially abnormal returns.

Initially, we run the target price achievement prediction model monthly to examine whether the target price is "achieved" or "unachieved." The classification model considers observations from the beginning of the sample period and ensures the inclusion of at least five years of data. Consequently, the portfolios begin in June 2005. We use the same optimal hyperparameters[5] and features as the classification models.

Subsequently, we rank the stocks according to their implied returns. For a stock with a positive implied return, if the target price is predicted to be achieved, we assign a "buy" label; conversely, if it's predicted not to be achieved, we assign a "sell" label. For a stock with a target price below its current price, if the target price is predicted to be achieved, we assign a "sell" label; conversely, if it's predicted not to be achieved, we assign a "buy" label. For a stock with multiple target prices in a month, only the highest (lowest) target price is considered for the "buy" ("sell") label.

Lastly, we build portfolios by longing the top 10% of stocks with the highest target price implied returns from the stocks with "buy" labels and shorting the

[5]The number of estimators "n_estimators" is reduced from 500 to 150 to mitigate the computation time.

bottom 10% of stocks with the lowest implied returns from the stocks with "sell" labels.[6]. The holding period for each portfolio is one year.

Additionally, we establish an equally weighted long-short benchmark based solely on the implied returns of target prices. Each month, the stocks are sorted by their implied returns. For each stock with multiple target prices in a month, if it has a positive (negative) implied return, we retain the highest (lowest) target price. Subsequently, we construct the benchmark portfolio each month by longing the top 10% of stocks with the highest target price implied returns and shorting the bottom 10% of stocks with the lowest implied returns. The annualized one-month Treasury bill rates from the Kenneth French Data Library are utilized as the risk-free rate to calculate excess returns.

Figure 3.9 depicts the performance of the portfolios and the benchmark from 2005 to 2021. Notably, these portfolios achieved exceptionally high returns during the financial crisis of 2008, whereas they experienced significantly negative returns during the COVID-19 pandemic. Figure 3.10 illustrates the cumulative returns of both the model-based portfolios and the benchmark, showing that the model-based portfolios significantly outperformed the benchmark. Table 3.6 reports the mean and volatility of portfolio returns, the average number of stocks, and the Sharpe ratios of the portfolios. The portfolios based on the random forest model show an average return of 2.94% and a volatility of 13.75%, whereas those based on the gradient boosting model yield an average return of 4.05% and a volatility of 20.98%. Both portfolios exhibit a Sharpe ratio of approximately 0.13, markedly

[6]For example, if there are 2024 target prices in a certain month, we include 202 stocks in both the long and short segments.

surpassing the benchmark's Sharpe ratio of 0.03. These results further corroborate the effectiveness of the classification models in guiding investment decisions.

3.4 Conclusion

Investors who base their decisions on inaccurate target prices may experience lower returns or even losses, as evidenced by our analysis of the accuracy of 12-month target price forecasts spanning from 1999 to 2021. Our study reveals that, on average, closing prices at the end of the forecast horizon fall below target prices by 14.07%, and the absolute forecast error is 36.63%. Additionally, only 36.33% of the target prices are met at the end of the forecast horizon. We show that whether a target price is achieved is strongly affected by market performance. Analysts consistently demonstrate limited ability to forecast stock prices accurately. This highlights the need for investors to be cautious when relying solely on target prices for investment decisions.

Nevertheless, target prices in analyst reports can still inform investment decisions. Our study utilizes ensemble methods in machine learning to predict target price accuracy. An important advantage of ensemble methods of decision trees is that they can capture complex nonlinear relationship among a large set of variables without significant risk of overfitting. We use forecast error and target price achievement as proxies for target price accuracy. We consider market-level, firm-level, and analyst-level features in predicting the accuracy. Our model allows us to identify the key drivers in predicting the accuracy of target prices. We show that the implied return of target price is the most important factor that affects the prediction of target price accuracy. Other important features include closeness to the past year high and market-level features such as the sentiment index and volatility. The regression models demonstrate good predictive power for target price forecast errors, while the classification models demonstrate strong predictive ability for target price achievement. Moreover, our results have direct practical benefits for portfolio management through predicting the achievability of target prices. The portfolios based on the prediction of target price achievement outperform the benchmark and provide significant returns, underscoring the value of these techniques in investment decision-making.

The application of machine learning to investment is a relatively recent development in the field of finance. However, it has rapidly gained popularity in recent years due to the explosion of available data and the development of sophisticated algorithms that can quickly and accurately process large amounts of information. Machine learning techniques are considered superior to traditional methods in terms of accuracy and flexibility. This study not only affirms the potential of machine learning in reducing bias and improving risk management but also paves the way for future research to explore and refine these methods further.

Several extensions could further our understanding of analyst forecast accuracy. One potential extension is to focus on the significance of contributions from analyst-specific characteristics in determining if analyst experience significantly impacts target price accuracy. Additionally, constructing portfolios by double sorting based on forecast bias and forecast precision could yield interesting insights. This method might reveal how the interplay between bias and precision affects investment performance, providing a nuanced understanding of analysts' forecast reliability. Moreover, with technological advancements, more advanced machine learning techniques, such as neural networks that require large computational capacity, are also worth exploring. These techniques could potentially improve the accuracy of predictions and offer deeper insights into the factors influencing analyst forecast misses.

Table 3.1: **Sample Selection**

The table reports the reduction in the sample size due to our data requirements. The target prices with 12-month forecast horizon are from I/B/E/S. The sample period is between January 2000 and December 2018.

All I/B/E/S target prices with 12-month forecast horizon issued between January 2000 and December 2018	1,157,659
Less: Observations with missing stock prices or features	(382,598)
Less: Observations with stock prices lower than one dollar	(792)
Less: Observations with extreme firm size, PB ratios, dividend payout ratios, the ratios of intangibles to total asset	(25,046)
Less: Observations with target prices above 400% or below 70% of or equal to current prices	(24,594)
Final sample	724,635

Table 3.2: **Summary statistics of implied returns of target prices and actual returns.**

This table summarizes descriptive statistics, including mean, standard deviation, and quartiles (Q1, median, Q3), for the implied returns of target prices, $TPP^{j}_{i,t,t+12}$, and the actual returns, $R_{i,t,t+12}$.

	Implied Return	Actual Return
Mean	0.1572	0.0165
Std	0.2010	0.4752
Q1	0.0415	-0.1764
Median	0.1335	0.0691
Q3	0.2408	0.2752

Table 3.3: **Summary statistics of target price forecast errors.**

The table exhibits the mean, standard deviation, skewness, and kurtosis of Error I and Error II for full sample and observations with positive and negative implied return of target price.

	Error I			Error II		
Summary Statistics	All	$TP^j_{i,t+12} > P_{i,t}$	$TP^j_{i,t+12} < P_{i,t}$	All	$TP^j_{i,t+12} > P_{i,t}$	$TP^j_{i,t+12} < P_{i,t}$
Mean	-0.1407	-0.1954	0.1271	0.3663	0.3704	0.3459
Std	0.5258	0.5223	0.4554	0.4025	0.4169	0.3224
Skewness	-1.4843	-1.6489	-0.9208	2.9269	2.9251	2.4971
Kurtosis	6.8780	7.3632	6.1874	15.3263	14.9835	12.4208

Table 3.4: **Selection of tuning hyperparameters.**

This table presents the set of tuning hyperparameters. There are 48 candidate models for random forest and 27 for gradient boosting. We use a 3-fold cross-validation splitting strategy for the training set. The optimal hyperparameters are selected for the corresponding models.

Notation	Hyperparameters	Candidates	Optimal
	Random Forest Regression		
n_estimators	Number of trees	100, 300, 500	500
max_features	Number of features	auto, 0.1, 0.2, 0.5	0.1
max_depth	Maximum depth	10, 15, 25, None	None
	Random Forest Classification		
n_estimators	Number of trees	100, 300, 500	500
max_features	Number of features	auto, 0.1, 0.2, 0.5	0.2
max_depth	Maximum depth	10, 15, 25, None	None
	Gradient Boosting Regression		
n_estimators	Number of trees	100, 300, 500	500
learning_rate	Learning rate	0.01, 0.1, 0.2	0.2
max_depth	Maximum depth	5, 10, 15	10
	Gradient Boosting Classification		
n_estimators	Number of trees	100, 300, 500	500
learning_rate	Learning rate	0.01, 0.1, 0.2	0.2
max_depth	Maximum depth	5, 10, 15	15

Table 3.5: **List of Features.**

This table describes features used in the models. We obtain analyst-level data from I/B/E/S, stock prices and market-level data from CRSP, and firm-level characteristics from Compustat.

Feature	Detail
Market-level features	
vix20	Average value of VIX in the previous 20 trading days
vix60	Average value of VIX in the previous 60 trading days
SENTO	Orthogonalized sentiment index by Baker and Wurgler (2006)
sp1m	1-month return of S&P500 Index
sp3m	3-month return of S&P500 Index
usd	US Dollar Index
ted	The difference between the three-month Treasury bill and the three-month LIBOR in USD
y_2	2-year treasury rate
y_5	5-year treasury rate
y_10	10-year treasury rate
y_diff	The difference between 2-year and 10-year treasury rates
Firm-level features	
SIZE	Firm size measured by the natural logarithm of total asset
ROA	Return on asset
turnover	63-day moving average volume / share outstanding
PB	Price-to-booking ratio
DIV	Dividend payout ratio
INTAN	Intangible asset to total asset ratio
ind_X	Industry dummy variable where X is the industry classified by the SIC system
1m_logret	1-month log return of the stock
2m_logret	2-month log return of the stock
3m_logret	3-month log return of the stock
6m_logret	6-month log return of the stock
vol_12m_back	1-year volatility of the stock return
prc_max_12mback	52-week high of the stock
prc_min_12mback	52-week low of the stock
Analyst-level features	
F_#Ana	The number of analysts following the stock
A_#Fm	The number of stocks the analyst follows
Ana_exp	The analyst's working experience in year
Ana_exp_fm	The number of years that the analyst has covered the stock
Ana_expo_exp	The number of more years that the analyst covers this firm compared with the average of the other analysts
logtpp	The logarithmic ratio of target price to current price
buy	Dummy variable with value of one if the target price is greater than the current price, and zero otherwise

Table 3.6: **Regression Results for the Target Price Errors.**

This table shows the coefficient of determination of the prediction, mean absolute error and root mean squared error for the random forest and gradient boosting regressions.

	Error I		Error II	
	Random Forest	Gradient Boosting	Random Forest	Gradient Boosting
R^2	0.369	0.570	0.341	0.516
MAE	0.2914	0.2422	0.2187	0.1976
RMSE	0.4198	0.3464	0.3257	0.2885

Table 3.7: **Prediction accuracy of target price achievement.**

This table shows the prediction accuracy on the training set and the test set and the balanced accuracy score for the random forest and gradient boosting classification models. We consider 56 features which combine firm/stock-specific, market-specific and analyst-specific predictor variables and use the optimal hyperparameters in the models.

	Random Forest	Gradient Boosting
Accuracy on Training Set	1.000	1.000
Accuracy on Test Set	0.887	0.890
Balanced Accuracy Score	0.866	0.873

Table 3.8: **Performance of equally weighted long-short portfolios.**

The table presents the average number of stocks, mean, volatility and Sharpe ratios of the benchmark and equally weighted long-short portfolios that use the predictions of target price achievement with random forest and gradient boosting classification models.

	Stocks	Mean (%)	Std (%)	t-stat	SR
Random Forest	290	2.94	13.75	3.02***	0.13
Gradient Boosting	291	4.05	20.98	2.72***	0.14
Benchmark	282	1.75	21.18	1.17	0.03

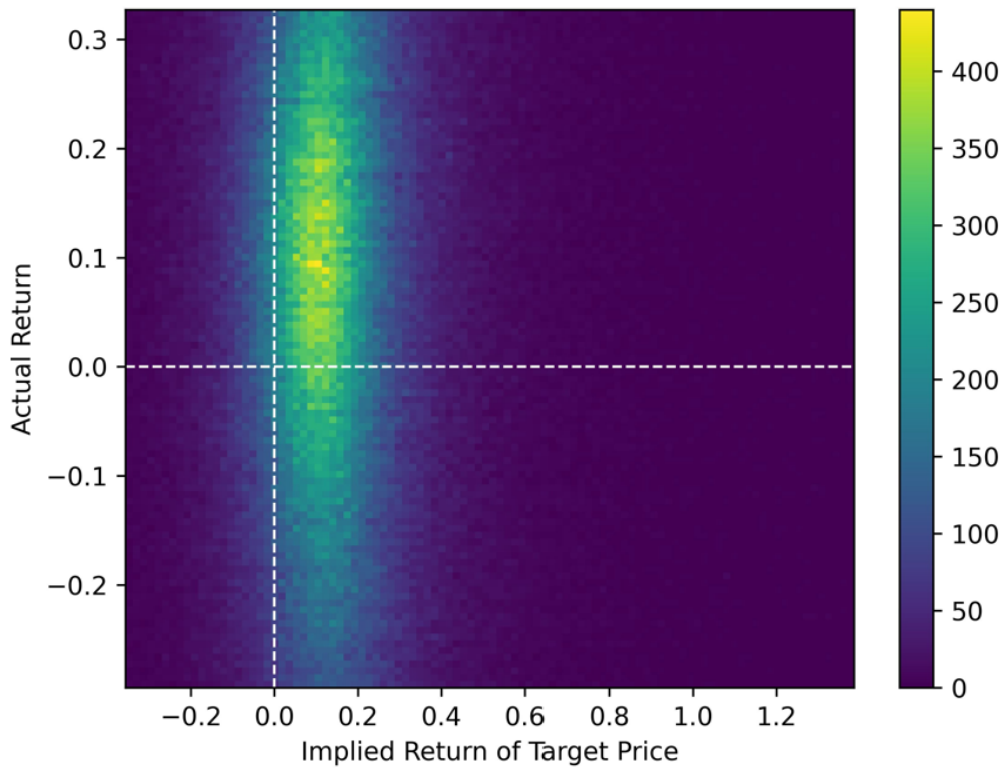

FIGURE 3.1: **Implied returns of target prices and actual returns.**

The figure displays the histogram of implied returns of target prices, $TPP^{j}_{i,t,t+12}$, and actual returns, $R_{i,t,t+12}$. The brighter the color, the higher the frequency of occurrence.

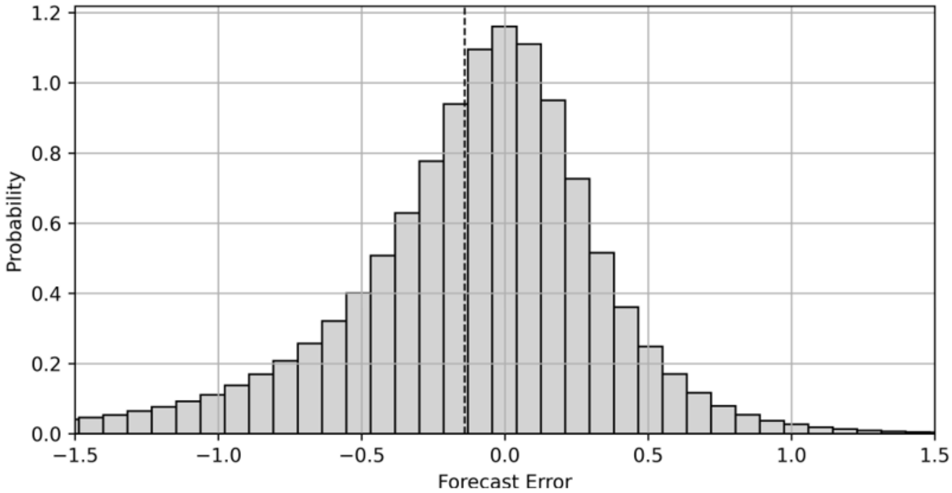

(a) Error I

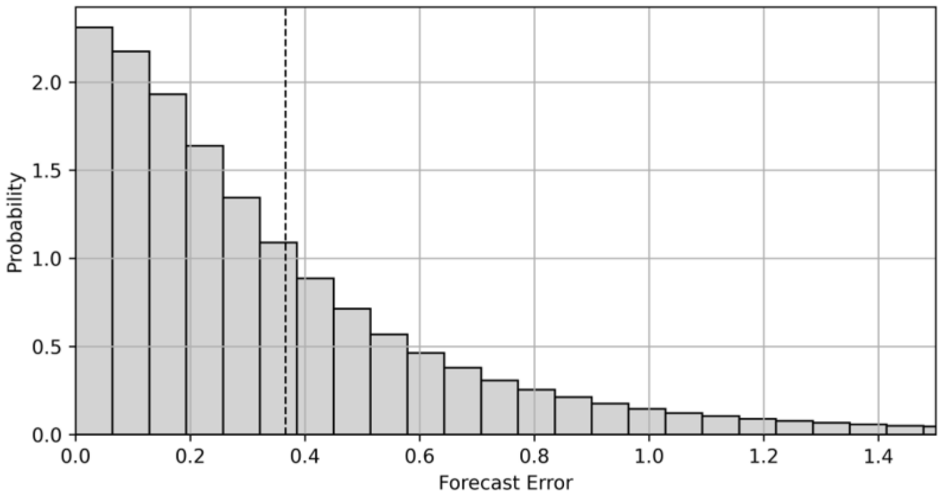

(b) Error II

Figure 3.2: **Histogram of target price forecast errors.**
The figure exhibits the histogram of the target price forecast errors (Error I and Error II). The dashed line shows the mean errors.

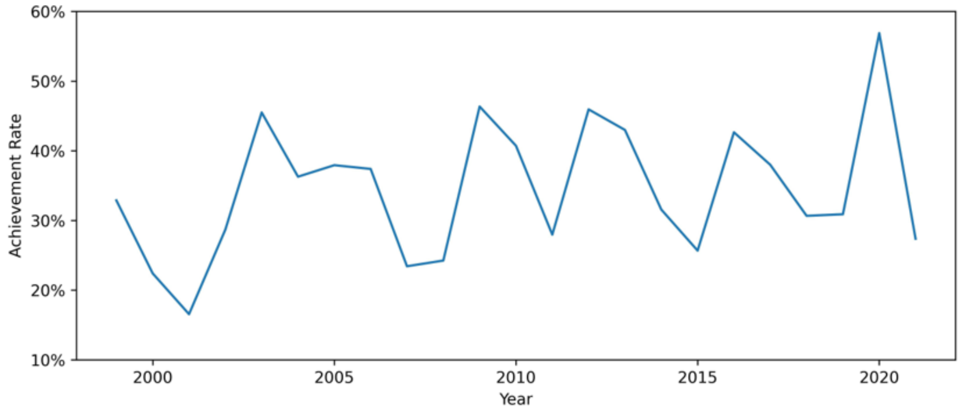

(a) Achievement Rate by Year

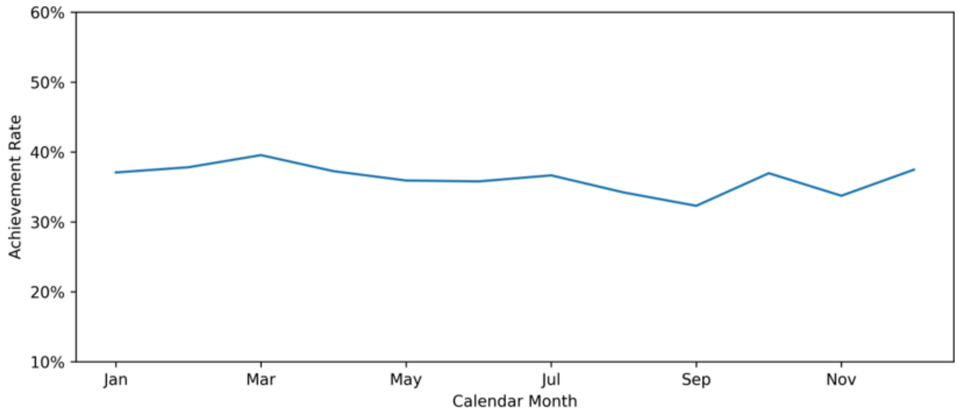

(b) Achievement Rate for Each Calendar Month

Figure 3.3: **Achievement rate of target prices by year and calendar month.**
The figure shows the annual achievement rate of target prices between 1999 and 2021, along with the average achievement rate for each calendar month.

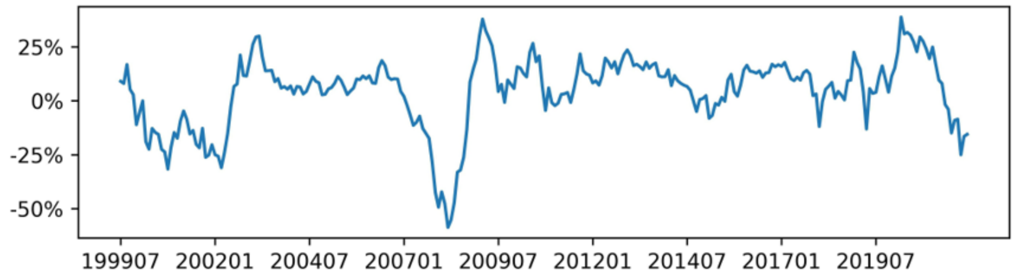

(a) SP500 Index Return

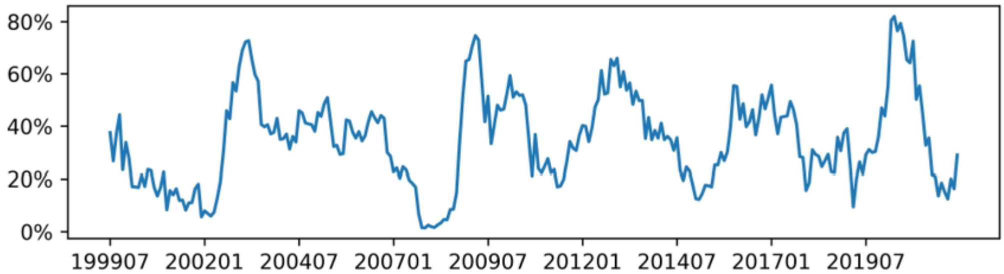

(b) Achievement Rate ($TP^{j}_{i,t+12} > P_{i,t}$)

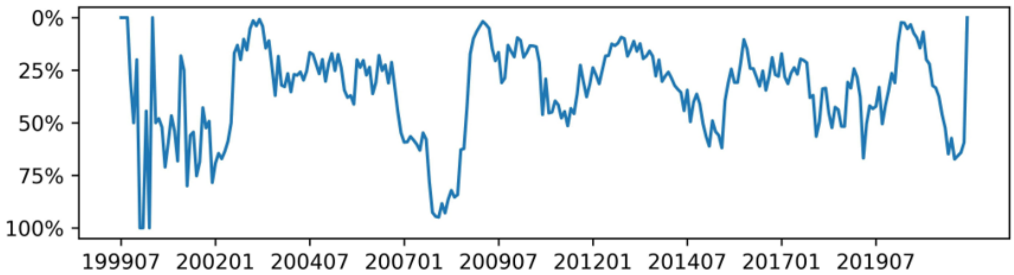

(c) Achievement Rate ($TP^{j}_{i,t+12} < P_{i,t}$)

FIGURE 3.4: **Market returns and target price achievement rate.**

The figure displays the S&P 500 Index returns and the target price achievement rates for observations with positive and negative implied return of target price between 1999 and 2021.

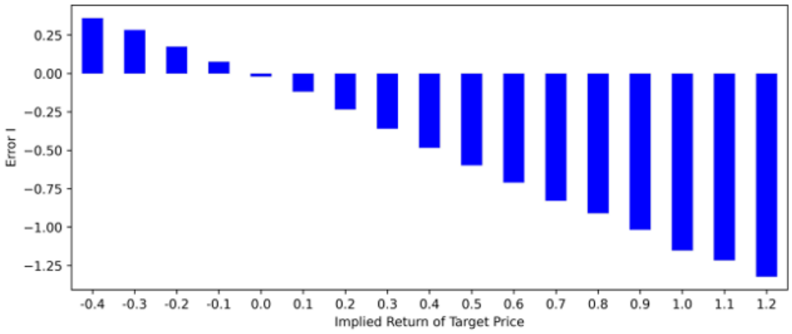

(a) Implied Return vs Error I

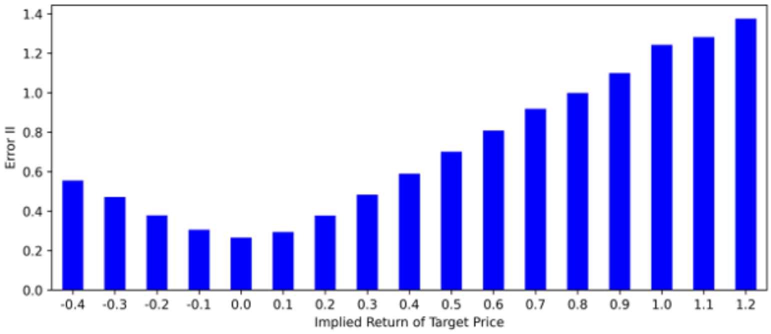

(b) Implied Return vs Error II

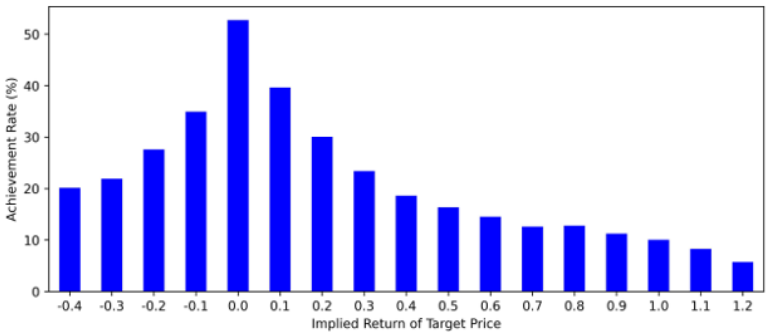

(c) Implied Return vs Achievement Rate

FIGURE 3.5: **Relationship between implied return of target price and forecast accuracy.**
The figure displays the variation of Error I, Error II and target price achievement with implied return of target price, $TPP^j_{i,t,t+12}$.

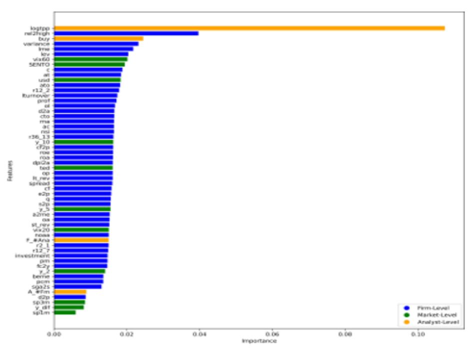

(a) Random Forest (Error I)

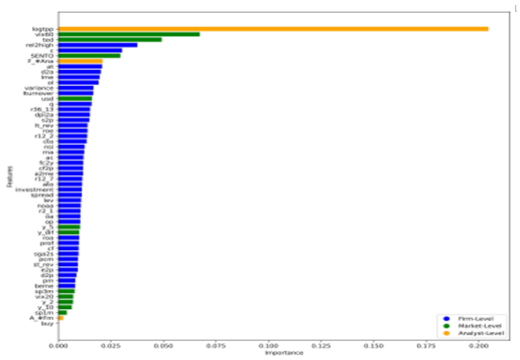

(b) Gradient Boosting (Error I)

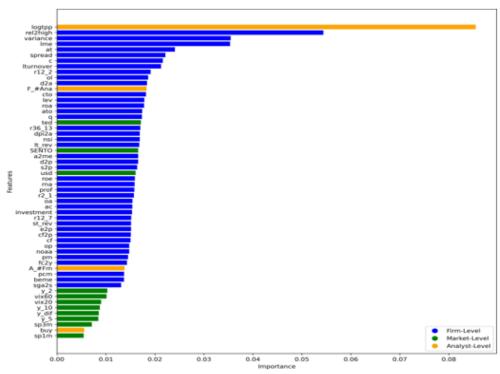

(c) Random Forest (Error II)

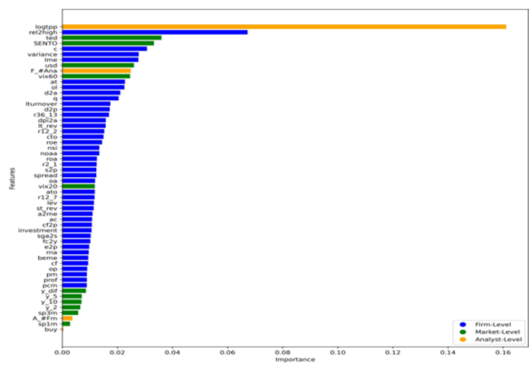

(d) Gradient Boosting (Error II)

FIGURE 3.6:
The figure presents the importance ranking of the 56 features in terms of Mean Decrease in Impurity (MDI) for the target price forecast error models. The firm/stock-specific predictor variables are highlighted in blue, and the market-specific and analyst-specific predictor variables are marked in green and orange, respectively.

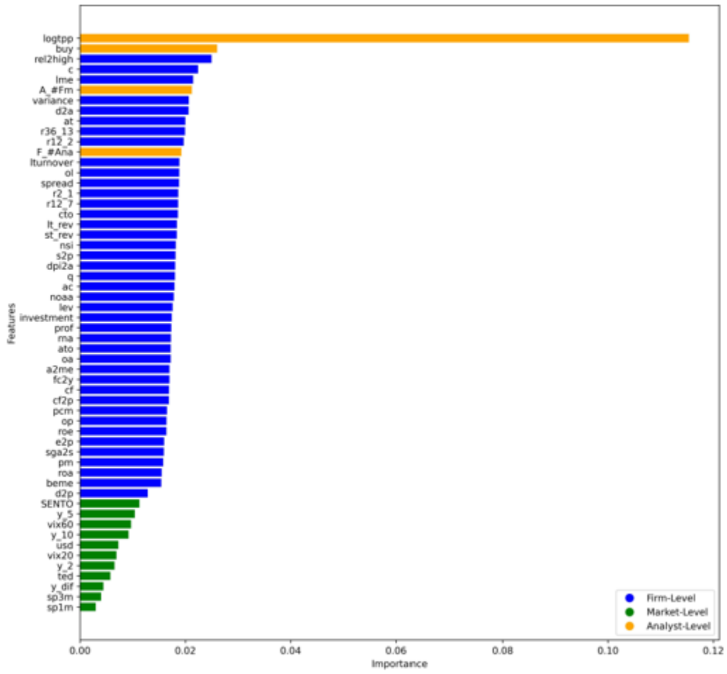

(a) Random Forest

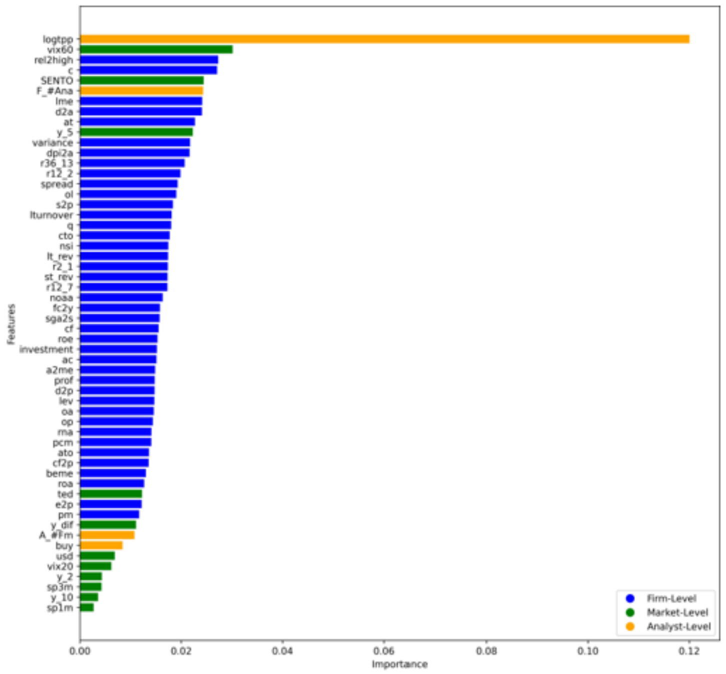

(b) Gradient Boosting

Figure 3.7: **Feature importances for the target price achievement models.**

The figure presents the importance ranking of the 56 features in terms of Mean Decrease in Impurity (MDI) for the target price achievement models. The firm/stock-specific predictor variables are highlighted in blue, and the market-specific and analyst-specific predictor variables are marked in green and orange, respectively. The upper (lower) panel displays the feature importance for predicting Error I (II).

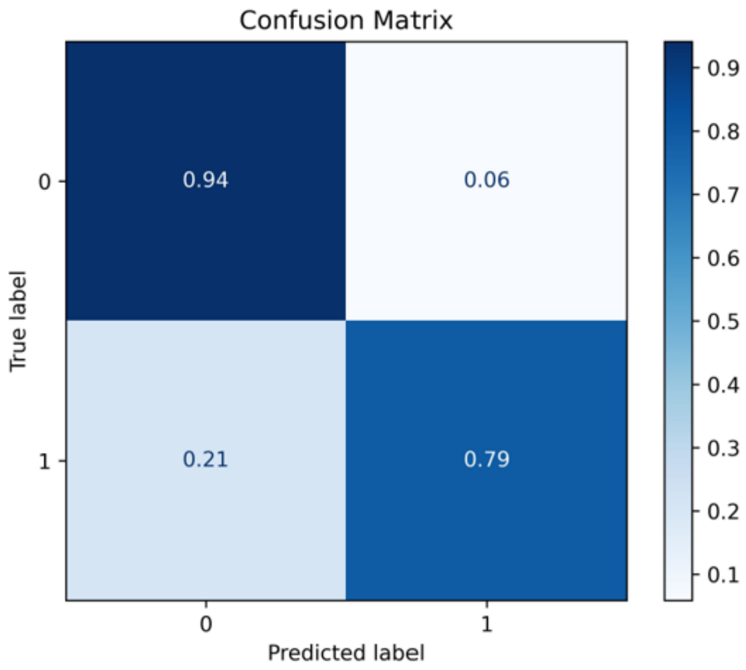

(a) Random Forest

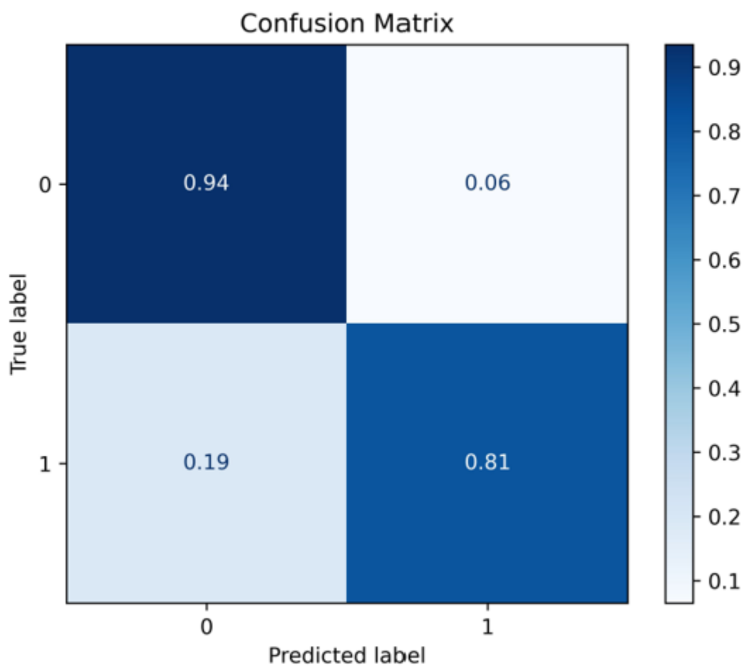

(b) Gradient Boosting

Figure 3.8: **Confusion matrix for the target price achievement models.**

The figure displays the confusion matrix for the target price forecast achievement models, indicating the accuracy of predictions per class (true positive, true negative, false positive, and false negative).

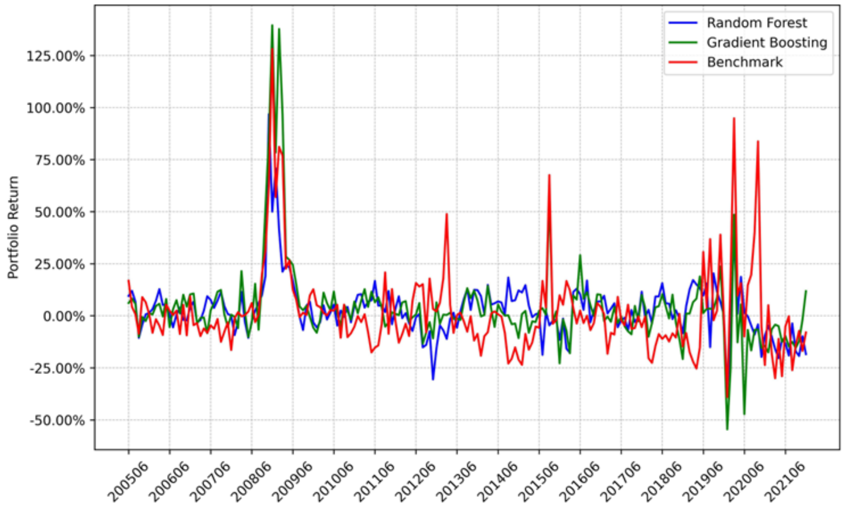

FIGURE 3.9: **Performance of equally weighted long-short portfolios and benchmark (2005-2021).**

The figure displays the annual returns of long-short portfolios, which are based on predictions of target price achievement using random forest and gradient boosting classification models, for each month over the period from 2005 to 2021. Keeping the chronological order, the model is estimated on an expanding window. We select the top (bottom) 10% of stocks by implied return for inclusion in the long (short) segment of the benchmark.

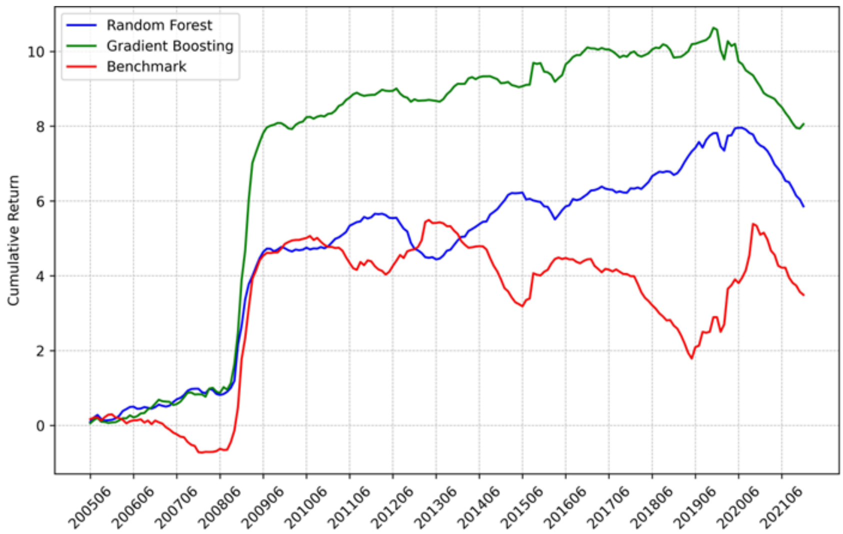

Figure 3.10: **Cumulative returns of equally weighted long-short portfolios and benchmark (2005-2021).**
The figure displays the cumulative returns of long-short portfolios, which are based on predictions of target price achievement using random forest and gradient boosting classification models, over the period from 2005 to 2021. Keeping the chronological order, the model is estimated on an expanding window. We select the top (bottom) 10% of stocks by implied return for inclusion in the long (short) segment of the benchmark.

Conclusion

In conclusion, this thesis contributes to the fields of commodity finance, climate finance, and behavioral finance by offering new methodologies, empirical insights, and practical implications for investors and policymakers.

The first paper presents a model confirming that the expected growth rate of natural resource prices positively affects the expected stock returns for commodity-producing firms. It highlights how firm-specific attributes influence expected returns, using an investment-based asset pricing approach to provide a refined test of the Hotelling Valuation Principle (HVP). The findings indicate the presence of unidentified systematic risk factors, suggesting areas for further research into the risk premia of commodity-producing firms, considering for instance real options and hedging strategies.

The second paper examines the impact of carbon pricing risk in the commodity markets, finding that this risk carries a significant positive risk premium. It underscores the importance of carbon pricing in financial markets and demonstrates that commodities with high sensitivity to carbon pricing risk exhibit positive loadings on this risk factor. By constructing a climate change hedge portfolio based on carbon pricing risk loadings, this study offers a practical method for managing climate change risks, contributing to the broader discourse on climate finance.

In the third chapter, the accuracy of target price forecasts by sell-side analysts is analyzed, and it reveals that target prices tend to overestimate future stock prices, significantly impacted by market performance. Utilizing ensemble methods of machine learning, we can well predict the accuracy of target prices. These models also identify key drivers of target price accuracy and show that portfolios based on these predictions outperform benchmarks, highlighting the potential of machine learning in improving investment decision-making and risk management.

Bibliography

Acharya, V. V., Lochstoer, L. A., and Ramadorai, T. (2013). Limits to arbitrage and hedging: Evidence from commodity markets. *Journal of Financial Economics* 109(2), 441–465.

Adelman, M. A. (1990). Mineral depletion, with special reference to petroleum. *The Review of Economics and Statistics*, 1–10.

Adelman, M. A. (1993). Modelling world oil supply. *The Energy Journal* 14(1), 1–32.

Ahmed, R., Bouri, E., Hosseini, S., and Hussain Shahzad, S. J. (2024). Spillover in higher-order moments across carbon and energy markets: A portfolio view. *European Financial Management*.

Alekseev, G., Giglio, S., Maingi, Q., Selgrad, J., and Stroebel, J. (2022). A quantity-based approach to constructing climate risk hedge portfolios. *Available at SSRN 4283192*.

Anderson, S. T., Kellogg, R., and Salant, S. W. (2018). Hotelling under pressure. *Journal of Political Economy* 126(3), 984–1026.

Arlinghaus, J. (2015). Impacts of carbon prices on indicators of competitiveness: A review of empirical findings-environment working paper no. 87. *OECD Environment Working Papers* (87), 1.

Asquith, P., Mikhail, M. B., and Au, A. S. (2005). Information content of equity analyst reports. *Journal of financial economics* 75(2), 245–282.

Azlen, M., Gostlow, G., and Child, A. (2022). The carbon risk premium. *The Journal of Alternative Investments*.

Bai, Y. and Okullo, S. J. (2023). Drivers and pass-through of the EU ETS price: Evidence from the power sector. *Energy Economics* 123, 106698.

Baker, M. and Wurgler, J. (2006). Investor sentiment and the cross-section of stock returns. *The journal of Finance* 61(4), 1645–1680.

Bakshi, G., Gao, X., and Rossi, A. G. (2019). Understanding the sources of risk underlying the cross section of commodity returns. *Management Science* 65(2), 619–641.

Balvers, R., Du, D., and Zhao, X. (2017). Temperature shocks and the cost of equity capital: Implications for climate change perceptions. *Journal of Banking & Finance* 77, 18–34.

Balvers, R. J., Cosimano, T. F., and McDonald, B. (1990). Predicting stock returns in an efficient market. *The Journal of Finance* 45(4), 1109–1128.

Barnett, M. D. (2019). A run on oil: Climate policy, stranded assets, and asset prices. PhD thesis. The University of Chicago.

Baur, D. G. (2014). Gold mining companies and the price of gold. *Review of Financial Economics* 23(4), 174–181.

Beckmann, J., Belke, A., and Czudaj, R. (2014). Does global liquidity drive commodity prices? *Journal of Banking & Finance* 48, 224–234.

Belo, F., Donangelo, A., Lin, X., and Luo, D. (2023). What drives firms' hiring decisions? An asset pricing perspective. *The Review of Financial Studies* 36(9), 3825–3860.

Berk, J. B., Green, R. C., and Naik, V. (1999). Optimal investment, growth options, and security returns. *The Journal of finance* 54(5), 1553–1607.

Bianchi, R. J., Drew, M. E., and Fan, J. H. (2016). Commodities momentum: A behavioral perspective. *Journal of Banking & Finance* 72, 133–150.

Bilinski, P., Lyssimachou, D., and Walker, M. (2013). Target price accuracy: International evidence. *The accounting review* 88(3), 825–851.

Bolton, P. and Kacperczyk, M. (2021). Do investors care about carbon risk? *Journal of Financial Economics* 142(2), 517–549. ISSN: 0304-405X.

Bonini, S., Zanetti, L., Bianchini, R., and Salvi, A. (2010). Target price accuracy in equity research. *Journal of Business Finance & Accounting* 37(9-10), 1177–1217.

Boons, M. and Prado, M. P. (2019). Basis-momentum. *The Journal of Finance* 74(1), 239–279.

Bradshaw, M. T., Brown, L. D., and Huang, K. (2013). Do sell-side analysts exhibit differential target price forecasting ability? *Review of Accounting Studies* 18(4), 930–955.

Bradshaw, M. T., Huang, A. G., and Tan, H. (2019). The effects of analyst-country institutions on biased research: Evidence from target prices. *Journal of Accounting Research* 57(1), 85–120.

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

Brennan, M. J. (1976). The supply of storage. *The economics of futures trading*, 100–107.

Brock, W. A. (1982). Asset prices in a production economy. In: *The economics of information and uncertainty*. University of Chicago Press, 1–46.

Buxbaum, M., Schultze, W., and Tiras, S. L. (2023). Do analysts' target prices stabilize the stock market? *Review of Quantitative Finance and Accounting*, 1–54.

Cairns, R. D. and Davis, G. A. (2001). Adelman's rule and the petroleum firm. *The Energy Journal* 22(3), 31–54.

Cairns, R. D. and Van Quyen, N. (1998). Optimal exploration for and exploitation of heterogeneous mineral deposits. *Journal of Environmental Economics and Management* 35(2), 164–189.

Chen, L., Pelger, M., and Zhu, J. (2023). Deep learning in asset pricing. *Management Science*.

Chen, A., Chang, C., Cheng, L., and Tu, H. (2016). Do analysts cater to investor beliefs via target prices. *International Review of Economics & Finance* 44, 232–252.

Cheng, L., Su, Y., Yan, Z., and Zhao, Y. (2019). Corporate governance and target price accuracy. *International Review of Financial Analysis* 64, 93–101.

Chevallier, J. (2011). A model of carbon price interactions with macroeconomic and energy dynamics. *Energy economics* 33(6), 1295–1312.

Chinn, M. D. and Coibion, O. (2014). The predictive content of commodity futures. *Journal of Futures Markets* 34(7), 607–636.

Choi, D., Gao, Z., and Jiang, W. (Feb. 2020). Attention to Global Warming. *The Review of Financial Studies* 33(3), 1112–1145. ISSN: 0893-9454.

Clarkson, P., Nekrasov, A., Simon, A., and Tutticci, I. (2020). Target price forecasts: The roles of the 52-week high price and recent investor sentiment. *Journal of Business Finance & Accounting* 47(9-10), 1365–1399.

Clatworthy, M. A., Peel, D. A., and Pope, P. F. (2012). Are analysts' loss functions asymmetric? *Journal of Forecasting* 31(8), 736–756.

Cochrane, J. H. (1991). Production-based asset pricing and the link between stock returns and economic fluctuations. *The Journal of Finance* 46(1), 209–237.

Cochrane, J. H. (1996). A cross-sectional test of an investment-based asset pricing model. *Journal of Political Economy* 104(3), 572–621.

Cochrane, J. H. (2001). *Asset pricing*. Princeton university press.

Cooper, I. A. and Nyborg, K. G. (2006). The value of tax shields IS equal to the present value of tax shields. *Journal of Financial Economics* 81(1), 215–225.

Cortazar, G., Millard, C., Ortega, H., and Schwartz, E. S. (2019). Commodity price forecasts, futures prices, and pricing models. *Management science* 65(9), 4141–4155.

Cox, J. C., Ingersoll Jr, J. E., and Ross, S. A. (1985). An intertemporal general equilibrium model of asset prices. *Econometrica: Journal of the Econometric Society*, 363–384.

Da, Z., Hong, K. P., and Lee, S. (2016). What drives target price forecasts and their investment value? *Journal of Business Finance & Accounting* 43(3-4), 487–510.

Da, Z., Huang, D., and Yun, H. (2017). Industrial electricity usage and stock returns. *Journal of Financial and Quantitative Analysis* 52(1), 37–69.

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

Danielsen, G. B. (2020). Can a Responsible Investor Invest in Commodity Futures? *Global Commodities Applied Research Digest*, 88–96.

Dar, A. B., Bhanja, N., and Paul, M. (2019). Do gold mining stocks behave like gold or equities? Evidence from the UK and the US. *International Review of Economics & Finance* 59, 369–384.

Daskalaki, C., Kostakis, A., and Skiadopoulos, G. (2014). Are there common factors in individual commodity futures returns? *Journal of Banking & Finance* 40, 346–363.

De Roon, F. A., Nijman, T. E., and Veld, C. (2000). Hedging pressure effects in futures markets. *The Journal of Finance* 55(3), 1437–1456.

Dechow, P. M. and You, H. (2020). Understanding the determinants of analyst target price implied returns. *The Accounting Review* 95(6), 125–149.

Demirakos, E. G., Strong, N. C., and Walker, M. (2010). Does valuation model choice affect target price accuracy? *European Accounting Review* 19(1), 35–72.

Ellis, G. M. and Halvorsen, R. (2002). Estimation of market power in a nonrenewable resource industry. *Journal of Political Economy* 110(4), 883–899.

Engle, R. F., Giglio, S., Kelly, B., Lee, H., and Stroebel, J. (2020). Hedging climate change news. *The Review of Financial Studies* 33(3), 1184–1216.

Eugene, F. and French, K. (1992). The cross-section of expected stock returns. *Journal of Finance* 47(2), 427–465.

Faccini, R., Matin, R., and Skiadopoulos, G. (2021). *Are climate change risks priced in the us stock market?* Tech. rep. Danmarks Nationalbank Working Papers.

Falls, N. and Wilson, O. C. (n.d.). Reexamine the Hotelling Valuation Principle: Empirical Evidence from Canadian Oil and Gas Royalty Trusts ().

Fama, E. F. and French, K. R. (2015). A five-factor asset pricing model. *Journal of financial economics* 116(1), 1–22.

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

Frankel, R., Kothari, S., and Weber, J. (2006). Determinants of the informativeness of analyst research. *Journal of Accounting and Economics* 41(1-2), 29–54.

Gaudet, G. and Khadr, A. M. (1991). The evolution of natural resource prices under stochastic investment opportunities: an intertemporal asset-pricing approach. *International Economic Review*, 441–455.

George, T. J. and Hwang, C.-y. (2007). Long-term return reversals: Overreaction or taxes? *The Journal of finance* 62(6), 2865–2896.

Giglio, S., Maggiori, M., Rao, K., Stroebel, J., and Weber, A. (2021). Climate change and long-run discount rates: Evidence from real estate. *The Review of Financial Studies* 34(8), 3527–3571.

Gleason, C. A., Bruce Johnson, W., and Li, H. (2013). Valuation model use and the price target performance of sell-side equity analysts. *Contemporary Accounting Research* 30(1), 80–115.

Görgen, M., Jacob, A., Nerlinger, M., Riordan, R., Rohleder, M., and Wilkens, M. (2020). Carbon Risk. *Available at SSRN*.

Gorton, G. B., Hayashi, F., and Rouwenhorst, K. G. (2013). The fundamentals of commodity futures returns. *Review of Finance* 17(1), 35–105.

Green, T. C., Jame, R., Markov, S., and Subasi, M. (2014). Access to management and the informativeness of analyst research. *Journal of Financial Economics* 114(2), 239–255.

Hao, J. and Skinner, J. (2022). Analyst target price and dividend forecasts and expected stock returns. *Journal of Asset Management*, 1–13.

Harris, R. S. and Pringle, J. J. (1985). Risk-adjusted discount rates-extensions from the average-risk case. *Journal of Financial Research* 8(3), 237–244.

Hayashi, F. (1982). Tobin's marginal q and average q: A neoclassical interpretation. *Econometrica: Journal of the Econometric Society*, 213–224.

He, G. and Li, A. Z. (2024). Does media coverage of firms' environment, social, and governance (ESG) incidents affect analyst coverage and forecasts? A risk perspective. *International Review of Financial Analysis* 94, 103289.

Hicks, J. R. (1975). Value and capital: An inquiry into some fundamental principles of economic theory. *OUP Catalogue*.

Hong, H. and Kostovetsky, L. (2012). Red and blue investing: Values and finance. *Journal of financial economics* 103(1), 1–19.

Hong, H. and Yogo, M. (2010). Commodity market interest and asset return predictability. *SSRN eLibrary* 3(5), 6.

Hotelling, H. (1931). The economics of exhaustible resources. *Journal of political Economy* 39(2), 137–175.

Hou, K., Xue, C., and Zhang, L. (2015). Digesting anomalies: An investment approach. *The Review of Financial Studies* 28(3), 650–705.

Hsu, P.-H., Li, K., and Tsou, C.-Y. (2023). The Pollution Premium. *The Journal of Finance* 78(3), 1343–1392.

Huang, D. and Kilic, M. (2019). Gold, platinum, and expected stock returns. *Journal of Financial Economics* 132(3), 50–75.

Huynh, T. D. and Xia, Y. (2021). Climate change news risk and corporate bond returns. *Journal of Financial and Quantitative Analysis* 56(6), 1985–2009.

Ilhan, E., Sautner, Z., and Vilkov, G. (2021). Carbon tail risk. *The Review of Financial Studies* 34(3), 1540–1571.

Jacobs, H. and Müller, S. (2020). Anomalies across the globe: Once public, no longer existent? *Journal of Financial Economics* 135(1), 213–230.

Jasic, T. and Wood, D. (2004). The profitability of daily stock market indices trades based on neural network predictions: Case study for the S&P 500, the DAX, the TOPIX and the FTSE in the period 1965–1999. *Applied Financial Economics* 14(4), 285–297.

Ji, C.-J., Hu, Y.-J., and Tang, B.-J. (2018). Research on carbon market price mechanism and influencing factors: a literature review. *Natural Hazards* 92, 761–782.

Jung, H., Engle, R. F., and Berner, R. B. (2021). *Climate stress testing.* Tech. rep. Staff Report.

Kaldor, N. (1939). Speculation and economic stability. *The Review of Economic Studies* 7(1), 1–27.

Kaniel, R., Lin, Z., Pelger, M., and Van Nieuwerburgh, S. (2023). Machine-learning the skill of mutual fund managers. *Journal of Financial Economics* 150(1), 94–138.

Karolyi, G. A. and Van Nieuwerburgh, S. (2020). New methods for the cross-section of returns. *The Review of Financial Studies* 33(5), 1879–1890.

Kerl, A. G. (2011). Target price accuracy. *Business Research* 4(1), 74–96.

Keynes, J. M. (1930). A Treatise on Money, Volume II: The Applied Theory of Money.

Khan, S., Khokher, Z., and Simin, T. (2017). The information content of a nonlinear macro-finance model for commodity prices. *The Review of Financial Studies* 30(8), 2818–2850.

Kim, H. S. and Koo, W. W. (2010). Factors affecting the carbon allowance market in the US. *Energy Policy* 38(4), 1879–1884.

Koch, N., Fuss, S., Grosjean, G., and Edenhofer, O. (2014). Causes of the EU ETS price drop: Recession, CDM, renewable policies or a bit of everything?—New evidence. *Energy Policy* 73, 676–685.

Kogan, L., Livdan, D., and Yaron, A. (2009). Oil futures prices in a production economy with investment constraints. *The Journal of Finance* 64(3), 1345–1375.

Koijen, R. S., Moskowitz, T. J., Pedersen, L. H., and Vrugt, E. B. (2018). Carry. *Journal of Financial Economics* 127(2), 197–225.

Lambert, R. A. (2004). Discussion of analysts' treatment of non-recurring items in street earnings and loss function assumptions in rational expectations tests on financial analysts' earnings forecasts. *Journal of Accounting and Economics* 38, 205–222.

Lee, B., O'Brien, J., and Sivaramakrishnan, K. (2008). An analysis of financial analysts' optimism in long-term growth forecasts. *The Journal of Behavioral Finance* 9(3), 171–184.

Lin, C.-Y. C. and Wagner, G. (2007). Steady-state growth in a Hotelling model of resource extraction. *Journal of Environmental Economics and Management* 54(1), 68–83.

Lin, X. (2012). Endogenous technological progress and the cross-section of stock returns. *Journal of Financial Economics* 103(2), 411–427.

Lin, X. and Zhang, L. (2013). The investment manifesto. *Journal of Monetary Economics* 60(3), 351–366.

Linnainmaa, J. T. and Roberts, M. R. (2018). The history of the cross-section of stock returns. *The Review of Financial Studies* 31(7), 2606–2649.

Litzenberger, R. H. and Rabinowitz, N. (1995). Backwardation in oil futures markets: Theory and empirical evidence. *The journal of Finance* 50(5), 1517–1545.

Liu, L. X., Whited, T. M., and Zhang, L. (2009). Investment-based expected stock returns. *Journal of Political Economy* 117(6), 1105–1139.

Livernois, J. (2009). On the empirical significance of the Hotelling rule. *Review of Environmental Economics and policy.*

Markov, S. and Tan, M. Y. (2006). Loss function asymmetry and forecast optimality: Evidence from individual analysts' forecasts.

McDonald, S. L. (1994). Why the Principle Over-Predicts Actual Values*. *The Energy Journal* 15(3), 1–17.

Miller, M. H. and Upton, C. W. (1985). A test of the Hotelling valuation principle. *Journal of Political Economy* 93(1), 1–25.

Morgenson, G. (2001). Price targets are hazardous to investors' wealth. *The New York Times.*

Ottaviani, M. and Sørensen, P. N. (2006). The strategy of professional forecasting. *Journal of Financial Economics* 81(2), 441–466.

Palley, A., Steffen, T. D., and Zhang, F. (2023). The effect of dispersion on the informativeness of analyst target prices. *Available at SSRN 3467800.*

Pindyck, R. S. (1978). The optimal exploration and production of nonrenewable resources. *Journal of political economy* 86(5), 841–861.

Pindyck, R. S. and Rotemberg, J. J. (1990). The excess co-movement of commodity prices. *The Economic Journal* 100(403), 1173–1189.

Pursiainen, V. (2022). Cultural biases in equity analysis. *The Journal of Finance* 77(1), 163–211.

Qian, H. (2009). Time variation in analyst optimism: An investor sentiment explanation. *The Journal of Behavioral Finance* 10(3), 182–193.

Ratti, R. A. and Vespignani, J. L. (2016). Oil prices and global factor macroeconomic variables. *Energy Economics* 59, 198–212.

Restoy, F. and Rockinger, G. M. (1994). On stock market returns and returns on investment. *The Journal of Finance* 49(2), 543–556.

Sakkas, A. and Tessaromatis, N. (2020). Factor based commodity investing. *Journal of Banking & Finance* 115, 105807.

Schwerhoff, G., Stuermer, M., et al. (2019). Non-renewable resources, extraction technology and endogenous growth. *Fed. Reserve Bank Dallas Work. Pap* 2015, 6–8.

Slade, M. E. (1982). Trends in natural-resource commodity prices: an analysis of the time domain. *Journal of Environmental Economics and Management* 9(2), 122–137.

Slade, M. E. (1984). Tax policy and the supply of exhaustible resources: Theory and practice. *Land Economics* 60(2), 133–147.

Slade, M. E. and Thille, H. (1997). Hotelling confronts CAPM: a test of the theory of exhaustible resources. *Canadian Journal of Economics*, 685–708.

Slade, M. E. and Thille, H. (2009). Whither Hotelling: tests of the theory of exhaustible resources. *Annu. Rev. Resour. Econ.* 1(1), 239–260.

Sousa, R., Aguiar-Conraria, L., and Soares, M. J. (2014). Carbon financial markets: a time–frequency analysis of CO2 prices. *Physica A: Statistical Mechanics and Its Applications* 414, 118–127.

Szymanowska, M., De Roon, F., Nijman, T., and Van Den Goorbergh, R. (2014). An anatomy of commodity futures risk premia. *The Journal of Finance* 69(1), 453–482.

Tan, X., Sirichand, K., Vivian, A., and Wang, X. (2020). How connected is the carbon market to energy and financial markets? A systematic analysis of spillovers and dynamics. *Energy Economics* 90, 104870.

Tobin, J. (1969). A general equilibrium approach to monetary theory. *Journal of money, credit and banking* 1(1), 15–29.

Tufano, P. (1998). The determinants of stock price exposure: Financial engineering and the gold mining industry. *The Journal of Finance* 53(3), 1015–1052.

Varian, H. R. (1992). *Microeconomic analysis*. WW Norton & Company.

Wang, Q. and Balvers, R. (2021). Determinants and predictability of commodity producer returns. *Journal of Banking & Finance* 133, 106278.

Wang, Y. and Guo, Z. (2018). The dynamic spillover between carbon and energy markets: new evidence. *Energy* 149, 24–33.

Wen, X., Bouri, E., and Roubaud, D. (2017). Can energy commodity futures add to the value of carbon assets? *Economic Modelling* 62, 194–206.

Working, H. (1949). The theory of price of storage. *The American Economic Review* 39(6), 1254–1262.

Yang, F. (2013). Investment shocks and the commodity basis spread. *Journal of Financial Economics* 110(1), 164–184.

Yu, L., Li, J., Tang, L., and Wang, S. (2015). Linear and nonlinear Granger causality investigation between carbon market and crude oil market: A multi-scale approach. *Energy Economics* 51, 300–311.

Yunis, J. and Aliakbari, E. (2020). *Carbon Pricing in High-Income OECD Countries.*

Zapata, H. O., Detre, J. D., and Hanabuchi, T. (2012). Historical Performance of Commodity and StockMarkets. *Journal of agricultural and applied Economics* 44(3), 339–357.

Zhang, L. (2005). The value premium. *The Journal of Finance* 60(1), 67–103.

Zhang, L. (2017). The investment CAPM. *European Financial Management* 23(4), 545–603.

Zhang, Y. (2015). The securitization of gold and its potential impact on gold stocks. *Journal of Banking & Finance* 58, 309–326.

Zhang, Y.-J. and Sun, Y.-F. (2016). The dynamic volatility spillover between European carbon trading market and fossil energy market. *Journal of Cleaner Production* 112, 2654–2663.

Appendix A

Chapter 1 Appendix

The model in the text consists of the Bellman equation:

$$V\left(s_t\right) = \text{Max}_{y_t,b_{t+1}}\left\{d_t + E_t\left[m_{t+1}V\left(s_{t+1}\right)\right]\right\}, \tag{A1}$$

with constraints:

$$d_t = q_t y_t - c\left(y_t, x_t\right) - r_{t-1}b_t + \left(b_{t+1} - b_t\right), \quad x_{t+1} = x_t - y_t, \quad q_{t+1} = h\left(q_t, \varepsilon_{t+1}\right) \tag{A2}$$

The first-order condition for y_t implies

$$q_t = c_y\left(y_t, x_t\right) + E_t\left[m_{t+1}V_x\left(x_{t+1}, q_{t+1}, b_{t+1}\right)\right] \tag{A3}$$

From the envelope theorem

$$V_x\left(s_t\right) = -c_x\left(y_t, x_t\right) + E_t\left[m_{t+1}V_x\left(s_{t+1}\right)\right] \tag{A4}$$

Combining (A3) and (A4) yields

$$V_x(s_t) = q_t - c_y(y_t, x_t) - c_x(y_t, x_t) \tag{A5}$$

Updating (A5) by one period and substituting into (A3) yields:

$$E_t\left[m_{t+1}\left(\frac{q_{t+1} - c_y(y_{t+1}, x_{t+1}) - c_x(y_{t+1}, x_{t+1})}{q_t - c_y(y_t, x_t)}\right)\right] = 1 \tag{A6}$$

The first-order condition for the choice of debt b_{t+1} generates:

$$E_t\left[m_{t+1}(1 + r_t)\right] = 1 \tag{A7}$$

which implies that $1 + r_t = 1/E_t(m_{t+1})$. This is simply an equilibrium condition resulting from the fact that the debt is riskless. It follows that the model does not pin down the level of debt, b_{t+1}, and the firm's capital structure. We include debt in the model to allow us to conceptualize differences between returns on assets and stock returns and relate them empirically to the proper variables. Given the equilibrium condition for stock returns, $E_t\left[m_{t+1}\left(1 + r^S_{t+1}\right)\right] = 1$ and the definition of stock returns as $1 + r^S_{t+1} = (p_{t+1} + d_{t+1})/p_t$ we have that

$$p_t = E_t\left[m_{t+1}(p_{t+1} + d_{t+1})\right]. \tag{A8}$$

This implies that $p_t = E_t\left[\sum_{j=1}^{\infty} m_{t+j} d_{t+j}\right]$, which means that p_t is the ex-dividend market equity of the firm (normalizing the number of outstanding shares to one) so that $V(s_t) = d_t + p_t$ by comparison with equation (2). Given (A2),

$$p_t = E_t \left\{ m_{t+1} \left[p_{t+1} + q_{t+1} y_{t+1} - c\left(y_{t+1}, x_{t+1}\right) - r_t b_{t+1} + \left(b_{t+2} - b_{t+1}\right) \right] \right\}. \tag{A9}$$

Using the method of undetermined coefficients, guess that stock prices are proportional to the reserves and outstanding debt and then confirm that this guess is justified. For all t :

$$p_t = F_t x_{t+1} + G_t b_{t+1}. \tag{A10}$$

Substitute into (A9) and use (A3) as well as the property of the homogeneous cost function that $c(y_t, x_t) = c_y(y_t/x_t)\, y_t + c_x(y_t/x_t)\, x_t$. Then:

$$\begin{aligned} F_t x_{t+1} + G_t b_{t+1} = E_t \{ m_{t+1} [& F_{t+1} x_{t+2} + G_{t+1} b_{t+2} + q_{t+1}\left(x_{t+1} - x_{t+2}\right) - r_t b_{t+1} + \left(b_{t+2} - b_{t+1}\right) \\ & - c_y\left(y_{t+1}/x_{t+1}\right)\left(x_{t+1} - x_{t+2}\right) - c_x\left(y_{t+1}/x_{t+1}\right) x_{t+1}] \} \end{aligned} \tag{A11}$$

Use equation (A6) $E_t \left\{ m_{t+1} \left[q_{t+1} - c_y\left(y_{t+1}/x_{t+1}\right) - c_x\left(y_{t+1}/x_{t+1}\right) \right] \right\} = q_t - c_y\left(y_t/x_t\right)$ and equation (A7), $E_t \left[m_{t+1}\left(1 + r_t\right) \right] = 1$ to simplify the right-hand side of (A11):

$$\begin{aligned} & \left[F_t - q_t + c_y\left(y_t/x_t\right) \right] x_{t+1} + \left(G_t + 1\right) b_{t+1} \\ & \quad = E_t \left\{ m_{t+1} \left[F_{t+1} - q_{t+1} + c_y\left(y_{t+1}/x_{t+1}\right) \right] x_{t+2} + \left(G_{t+1} + 1\right) b_{t+2} \right\} \end{aligned} \tag{A12}$$

This equation is of the form $Z_t = E_t \left[m_{t+1}\left(Z_{t+1}\right) \right]$ which has as the only non-bubble solution that $Z_t = 0$. Thus, we confirm the guessed solution and find that $F_t = q_t - c_y\left(y_t/x_t\right)$ and $G_t = -1$ for all t. Hence, we obtain equation (7) in the text

$$p_t = \left[q_t - c_y\left(y_t/x_t\right) \right] x_{t+1} - b_{t+1}. \tag{A13}$$

The stock return, $1 + r^S_{t+1} = (p_{t+1} + d_{t+1})/p_t$, is obtained from (A2) and (A13), using $c(y_t, x_t) = c_y(y_t/x_t)\, y_t + c_x(y_t/x_t)\, x_t$:

$$1 + r^S_{t+1} = \frac{q_{t+1} - c_y(y_{t+1}/x_{t+1}) - c_x(y_{t+1}/x_{t+1}) - (1 + r_t)(b_{t+1}/x_{t+1})}{q_t - c_y(y_t/x_t) - (b_{t+1}/x_{t+1})}. \tag{A14}$$

Using (A13) and (A14) the excess return then equals, as given in equation (8) in the text:

$$r^S_{t+1} - r_t = \frac{\{[(q_{t+1} - q_t)/q_t] - r_t\} + [(1 + r_t)\, c_y(y_t/x_t) - c_y(y_{t+1}/x_{t+1}) - c_x(y_{t+1}/x_{t+1})]/q_t}{p_t/q_t x_{t+1}} \tag{A15}$$

Given the investment return from (A6): $1 + r^I_{t+1} = \frac{q_{t+1} - c_y(y_{t+1}, x_{t+1}) - c_x(y_{t+1}, x_{t+1})}{q_t - c_y(y_t, x_t)}$ (equal to the return on assets here), it is easy to confirm that

$$r^I_{t+1} = (1 - \lambda_t)\, r^S_{t+1} + \lambda_t r_t, \lambda_t = b_{t+1}/(p_t + b_{t+1}). \tag{A16}$$

Appendix B

Chapter 2 Appendix

B1 Appendix A

The hypotheses development is guided by the investment-based model for commodity producers proposed by Kogan et al. (2009) and developed by Yang (2013). Following Yang (2013), I assume the commodity production economy to be competitive, with each commodity having many identical producers, allowing their complex behavior to be reduced to a single representative producer problem. In this commodity production economy, there are N commodities, and the representative producer of commodity j is assumed to have a Cobb-Douglas production function:

$$Q^S_{j,t} = AK^{\alpha}_{j,t} \tag{A.1}$$

where $Q^S_{j,t}$ is the production of commodity j at time t, (j $= 1, 2 \ldots$ N), A is the total factor productivity, and $K_{j,t}$ is the capital at time t for producing commodity j. Production is assumed to be deterministic so A = 1 and $\alpha = 1$, indicating constant returns to scale. Following Fisher (2006), the capital of the producer accumulates as

$$\mathrm{K_{j,t+1}} = (1 - \delta)\mathrm{K_{j,t}} + (\mathrm{Y_t} - C_t)\,\mathrm{I_{j,t}} \tag{A.2}$$

where δ is the rate of physical depreciation, and $I_{j,t}$ is the amount of investment. Yang (2013) introduces an investment-specific technology as Y_t, which impacts the investment of all commodity producers ($\mathrm{j} = 1, 2 \ldots$ N). A positive investment shock increases the efficiency of investment, leading to higher future capital $K_{j,t+1}$ for the same level of investment $I_{j,t}$. Higher efficiency in investment corresponds to lower investment costs or lower price of investment goods (Yang 2013). Building on Yang's model, I introduce a carbon pricing shock (C_t) that negatively impacts the investment of all commodity producers. The carbon pricing shock, reflecting any unexpected increase in carbon prices, reduces investment efficiency by increasing production costs through greenhouse gas emission allowances or other charges related to climate change mitigation. Specifically, the carbon pricing shock substantially increases production costs for commodities with high greenhouse gas emissions, such as precious and industrial metals. Commodities with lower emissions may also be affected by increased transportation or utility costs due to carbon pricing. The introduced carbon pricing shock affects investment channels, reflecting uncertainty about future economic activity paths aimed at mitigating climate change. Overall, the capital accumulation function (eq.A.2) indicates that commodity producers' investment efficiency is influenced by both the conventional investment shock (Y_t) and carbon pricing shock (C_t), albeit in opposite directions.

On the supply side of the commodity production economy, a representative producer of commodity j maximizes their firm's valuation by choosing the amount of investment $I_{j,t}$ at each time t , subject to the constraint that investment is irreversible.

$$V_{j,0} = \max_{\mathrm{I_{j,t}}} \mathrm{E_0} \left[\sum_{\mathrm{t=0}}^{\infty} \left(\prod_{\mathrm{u=0}}^{\mathrm{t}} \mathrm{M_u} \right) \left(\mathrm{P_{j,t} Q^s_{j,t}} - \mathrm{I_{j,t}} \right) \right] \tag{A.3}$$

$$\text{s.t. } \mathrm{I_{j,t}} \geq 0$$

where M_{t+1} denotes the stochastic discount factor (SDF) from time t to time $t+1$, with $M_0 = 1$. Following Kogan Livdan, and Yaron (2009) and Yang (2013), the demand for commodity j is assumed to be exogenous, driven by two types of exogenous demand shocks: aggregate demand shock X_t and idiosyncratic demand shock $Z_{j,t}$. The inverse demand function for commodity j is defined as

$$P_{j,t} = \left(\frac{X_t Z_{j,t}}{Q^D_{j,t}} \right)^{\eta} \tag{A.4}$$

where $Q^D_{j,t}$ is the demand quantity for commodity j at time t , and η measures the price elasticity of demand. With the supply and demand functions defined as above, the commodity production economy experiences four types of exogenous shocks: aggregate investment shock Y_t, carbon pricing shock C_t, aggregate demand shock X_t, and idiosyncratic demand shock $Z_{j,t}$.

The log aggregate investment shock, $\log Y_t$, is assumed to follow the AR(1) process as in Justiniano, Primiceri, and Tambalotti (2010),

$$\log Y_{t+1} = \rho_y \log Y_t + \sigma_y u_{t+1} \tag{A.5}$$

The log carbon pricing shock, $\log C_t$, is assumed to follow the AR(1) process since it is considered a type of "negative" investment shock:

$$\log C_{t+1} = \rho_c \log C_t + \sigma_c \xi_{t+1} \tag{A.6}$$

Further, the log aggregate demand shock, $\log X_t$, is assumed to follow a random walk process, while the log idiosyncratic demand shock, $\log Z_{j,t}$, is assumed to follow an AR(1) process as follows:

$$\log X_{\mathrm{t+1}} = \log \mathrm{X_t} + \mathrm{g_x} + \sigma_{\mathrm{x}} \mathrm{e_{t+1}} \tag{A.7}$$

$$\log \mathrm{Z_{j,t+1}} = \rho_{\mathrm{z}} \log Z_{\mathrm{j,t}} + (1 - \rho_{\mathrm{z}}) \overline{\mathrm{Z}} + \sigma_{\mathrm{z}} \epsilon_{\mathrm{j,t+1}} \tag{A.8}$$

Following Berk et al. (1999), Zhang 2005,and Yang 2013, the SDF in this model is parameterized directly as

$$M_{t+1} = \frac{1}{\phi_t} \exp\left[-r_f - \gamma_x \left(\log X_{t+1} - \log X_t\right) - \gamma_y \left(\log Y_{t+1} - \log Y_t\right) \right. \tag{A.9}$$
$$\left. -\gamma_c \left(\log C_{t+1} - \log C_t\right)\right]$$

Where ϕ_t is a compensator so that the no-arbitrage condition $E_t\left[M_{t+1}\right] = e^{-r_f}$ is satisfied over all states of the economy. The risk-free rate r_f is assumed to be constant.

Spot price

The market-clearing condition for each commodity j is that the supply quantity equals to the demand quantity.

$$Q_{j,t}^{S} = Q_{j,t}^{D} \tag{A.10}$$

From the market-clearing condition and the production technology $\mathrm{Q}_{\mathrm{j,t}}^{\mathrm{s}} = \mathrm{K}_{\mathrm{j,t}}^{\alpha}$, it is evident that the quantity of commodity j is determined by the producers' capital $\left(Q_{j,t}^{S} = Q_{j,t}^{D} = K_{j,t}^{\alpha}\right)$. Taking the difference of the log inverse demand function over time:

$$\begin{aligned} &\log P_{j,t+1} - \log P_{j,t} \\ &\qquad = \eta\left(\log X_{t+1} - \log X_{t}\right) + \eta\left(\log Z_{j,t+1} - \log Z_{j,t}\right) \\ &\qquad - \eta \log\left(\frac{K_{j,t+1}}{K_{jt}}\right) \end{aligned} \tag{A.11}$$

Assuming the optimal investment rate to be $i^{*} = \frac{I_{j,t}^{*}}{K_{j,t}}$, the dynamics of capital can be written as:

$$\frac{K_{j,t+1}}{K_{j,t}} = 1 - \delta + \left(Y_{t} - C_{t}\right) i_{j,t}^{*} \tag{A.12}$$

Plugging the dynamics of capital and the exogenous shocks into the price first difference equation, the log spot price dynamics can be written as:

$$\log P_{j,t+1} - \log P_{j,t} = g_{j,t}^{p} + \eta\sigma_{x}e_{t+1} + \eta\sigma_{z}\epsilon_{j,t+1} \tag{A.13}$$

Where $\mathrm{g}_{\mathrm{j,t}}^{\mathrm{p}} = -\eta \log\left(1 - \delta + \mathrm{Y}_{\mathrm{t}}\mathrm{i}_{\mathrm{j,t}}^{*} - \mathrm{C}_{\mathrm{t}}\mathrm{i}_{\mathrm{j,t}}^{*}\right) + \eta \mathrm{g}_{\mathrm{x}} + \eta\left(1 - \rho_{\mathrm{z}}\right)\left(\overline{\mathrm{z}} - \log \mathrm{Z}_{\mathrm{j,t}}\right)$, and η is the price elasticity of demand, $\mathrm{i}_{\mathrm{j,t}}^{*}$ is the optimal investment, g_{x} is the growth rate of the aggregate demand shock, ρ_{z} is the persistence of the idiosyncratic demand shock, $\bar{Z}$ is the long-run mean of idiosyncratic demand shock, and σ is the conditional volatility of the corresponding exogenous shocks.

Futures price

A commodity futures contract is a claim written on the commodity that is sold on the spot market at prevailing spot price. The payoff of longing a futures contract of commodity j with maturity T is $P_{j,T} - F_{j,t,T}$ at maturity, where $P_{j,T}$ is the spot price of commodity j at time T and $F_{j,t,T}$ is the price of the future contract on commodity j with maturity T . The initial value of the futures contract is zero when both counterparties enter into the agreement. Therefore, at the initial time t, the following equation is satisfied:

$$0 = E_t \left[M_T \left(P_{j,T} - F_{j,t,T} \right) \right] = E_t \left[\left(\prod_{u=t+1}^{T} M_u \right) \left(P_{j,T} - F_{j,t,T} \right) \right] \tag{A.14}$$

Under the risk-neutral measure, the futures price $F_{j,t,T}$ can be computed as the expected spot price of commodity j at time T:

$$P_{j,t} = \frac{F_{j,t,T}}{e^{r(T-t)}} = \frac{E^Q \left[F_{j,t+1,T} \right]}{e^{r(T-t)}} = E_t \left[M_{t+1} \frac{F_{j,t+1,T}}{e^{r(T-(t+1))}} \right] = \frac{E_t \left[M_{t+1} F_{j,t+1,T} \right]}{e^{r(T-t-1)}} \tag{A.15}$$

Following eq.A.15, the futures price can be written as:

$$F_{j,t,T} = E^Q \left[F_{j,t+1,T} \right] = e^{r_f E_t \left[M_{t+1} F_{j,t+1,T} \right]} \tag{A.16}$$

The boundary condition is $F_{j,T,T} = P_{j,T}$, which is the no-arbitrage condition for futures contracts. Since no analytical expression exists for the above expectation (eq. A.16), I use it to show that the spot price dynamics determine the risk exposure for the futures return as the futures price converges to its spot price when the futures contracts approach the maturities ($F_{j,T,T} = P_{j,T}$). Therefore, the identified four risks in the spot price dynamics also drive the futures returns: aggregate investment shock (Y_t), carbon pricing shock (C_t), aggregate demand shock (X_t), and idiosyncratic demand shock ($Z_{j,t}$). These risks are systematic except the idiosyncratic demand shock ($Z_{j,t}$), so the risk exposure to these systematic risks (Y_t, C_t, X_t) determines the risk premiums of futures contracts. Moreover, spot price dynamics (eq. A.13) determine the risk exposure of futures contracts as well since the futures contract is a claim written on the spot price. Based on spot price dynamics (eq. A.13), the risk loading on carbon pricing shock (C_t) is approximately $\eta i^*_{j,t}$, which is positive and dependent on the optimal investment rate $\left(i^*_{j,t}\right)$ (eq. A.17).

$$\beta^C_{j,t+1,T} \approx \eta i^*_{j,t} > 0 \tag{A.17}$$

$\beta^C_{j,t+1,T}$ is the risk exposure of a futures contract of commodity j with maturity T at time t + 1 to the carbon pricing risk, and η is the price elasticity of demand.

The approximate risk loading in eq.A. 17 implies that investors demand higher returns on futures with carbon pricing shock. Because the capital of producers determines the future supply and hence the prices of commodities in the future, carbon pricing shock will raise the investment cost of commodity production, which implies a decrease in future capital installed given the level of investment, and hence decreases the future supply and appreciates the commodity futures prices. The futures prices of high investment commodities are more sensitive to investment shocks and carbon pricing shocks than low investment ones for the same reason. Since the loading on the carbon pricing risk is positive, I expect the risk premium associated with the carbon emission risk to be positive, ensuring that riskier commodity assets yield higher returns.

B2 Appendix B

The commodity sorting characteristics, Basis, Momentum, Basis-Momentum, Value, and Volatility are defined as follows:

a. Basis: Defined as $B_{i,t} = \frac{\log F_{i,t,1} - \log F_{i,t,2}}{D_2 - D_1}$, where $F_{i,t,1}$ is the futures price of commodity i for the nearest contract at time t, and D_1 is the number of days to maturity for this future contract. $F_{i,t,2}$ is the futures price on time t for the second nearest contracts and D_2 is the number of days to maturity of the second nearest contract at time t.

b. Momentum: The cumulative excess futures return from the prior 12 months.

c. Basis-Momentum: Following (Boons and Prado (2019)), it is the difference between the momentum in the first- and second-nearest contracts.

d. Volatility: Defined as the coefficient of variation (CV), the variance per absolute mean of the futures returns during the prior 36 months: $CV_t^j = \left(\frac{\sigma_{j,t}^2}{|\mu_{j,t}|}\right)$.

e. Value: The average futures prices of the first nearby futures contracts 4.5 to 5.5 years ago over of the nearest future price at time t, Value $_t^i = \ln\left(\frac{F_{i,t-60}^{T_1}}{F_{i,t}^{T_1}}\right)$, where $F_{i,t-60}^{T_1}$ is the average price of the nearest contract from 54 months to 66 months ago.

Appendix C

Chapter 3 Appendix

C1 Appendix A:Algorithms of machine learning models

C1.1 Random forest regression

1. Bootstrapping: Randomly select N samples from the dataset with replacement to create multiple subsets for training individual trees.

2. Feature Selection: Randomly select a subset of features (denoted by m) from the total M features.

3. Decision Trees: Build a decision tree using the selected subset of data and features. Split the nodes based on the best split criteria until a leaf node is reached.

4. Ensemble Learning: Repeat steps 1-3 for n times to create n decision trees.

4. Prediction: Predict the target variable by averaging the predictions of all K trees for regression. For each input data point, the output is the average of the

individual tree predictions.

5: Aggregation: The predictions are averaged across all decision trees.

$$\hat{y} = \frac{1}{n} \sum_{i=1}^{n} y_i \tag{C.1}$$

where $\hat{y}$ represents the predicted value; y_i represents the individual predictions from each decision tree.

C1.2 Random forest classification

1. Bootstrapping: Randomly select N samples from the dataset with replacement to create multiple subsets for training individual trees.

2. Feature Selection: Randomly select a subset of features (denoted by m) from the total M features.

3. Decision Trees: Build a decision tree using the selected subset of data and features. Split the nodes based on the best split criteria until a leaf node is reached.

4. Ensemble Learning: Repeat steps 1-3 for n times to create n decision trees.

4. Prediction: Predict the target variable by averaging the predictions of all K trees for regression. For each input data point, the output is the average of the individual tree predictions.

5: Aggregation: The mode (most frequent class label) is taken as the predicted class.

$$\hat{y} = \text{mode}(y_1, y_2, ..., y_n) \tag{C.2}$$

where $\hat{y}$ represents the class label; y_i represents the individual predictions from each decision tree.

C1.3 Gradient boosting regression

1. Initialize Predictions: Initialize the prediction with the mean of the target variable for regression tasks.

$$F_0(x) = \text{mean}(y) \tag{C.3}$$

2. Calculate Residuals: Calculate the residuals with input x by subtracting the current predictions from the actual target values.

$$r_i = y_i - F_{m-1}(x_i) \tag{C.4}$$

3. Fit a Weak Learner[1] (Decision Tree): Fit a weak learner (decision tree) to the residuals to capture the errors made by the current ensemble.

$$h_m(x) = \text{argmin}_h \sum_{i=1}^{N} (h(x_i) - r_i)^2 \tag{C.5}$$

where $h_m(x)$ are weak learners in the context of boosting.

[1]For gradient boosting regression model, we use squared error as the loss function to be optimized.

4. Update the Prediction with Learning Rate (η): Update the predictions by adding a scaled version of the predictions from the weak learner to the current predictions.

$$F_m(x) = F_{m-1}(x) + \eta h_m(x) \tag{C.6}$$

5. Repeat: Repeat steps 2-4 for a specified number of iterations to build multiple weak learners and improve the overall prediction.

C1.4 Gradient boosting classification

1. Initialize Class Probabilities: Initialize the model by setting the initial predicted values $P_0(C_i|x)$ for sample x.

$$P_0(C_i|x) = \log(\frac{p(x)}{1-p(x)}) \tag{C.7}$$

2. Calculate Pseudo-Residuals: Calculate pseudo-residuals, which are gradients of the loss function with respect to the current predictions.

$$r_i^{(m)} = -\left[\frac{\partial L(y_i, P_{m-1}(x_i))}{\partial P_{m-1}(x_i)}\right]_{P_{m-1}(x_i)=P_{m-1}(C_1|x_i)} \tag{C.8}$$

where is $L(y_i, P_{m-1}(x_i))$ the log loss (deviance) for binary classification.

3. Fit a Weak Learner[2] (Decision Tree): Fit a weak learner (decision tree) to the pseudo-residuals to capture the gradients made by the current ensemble.

[2]For gradient boosting classification model, we use binomial deviance (negative log-likelihood) as the loss function to be optimized.

$$h_m(x) = \text{argmin}_h \sum_{i=1}^{N} \left(h(x_i) - r_i^{(m)} \right)^2 \tag{C.9}$$

4. Update the Prediction with Learning Rate (η): Update class probabilities by adding a scaled version of the weak learner's predictions to the current probabilities.

$$P_m(C_1|x) = P_{m-1}(C_1|x) + \eta h_m(x) \tag{C.10}$$

5. Repeat: Repeat steps 2-4 for a specified number of iterations to build multiple weak learners and improve the overall prediction.