UNIVERSITY OF OXFORD

Saïd Business School

(Generative) AI in Financial Economics

Focus: Asset Pricing

Shumiao Ouyang

Saïd Business School

Mo, H., & Ouyang, S. (2025). (Generative) AI in Financial Economics [Working paper]. University of Oxford.

impact *from within*

Overview

Table of Contents

Traditional AI/ML vs. Generative AI/LLMs in Finance

	Traditional AI/ML	**Generative AI / LLMs**
Goal	Predictive modeling, classification, pattern detection	Content generation, summarization, reasoning
Techniques	Regression, trees, SVMs, neural nets, clustering	Transformers (e.g., GPT), GANs, VAEs
Data	Mostly structured (e.g., numerical finance data)	Primarily unstructured text/code; increasingly multimodal
Strength	Statistical learning from data patterns	Contextual understanding, natural language generation
Use Cases	Credit scoring, fraud detection, risk modeling, trading signals	Robo-advising, sentiment/narrative analysis, report/code generation
Risks	Overfitting, bias, lack of interpretability	Hallucinations, bias, privacy leaks, explainability, misuse

LLMs as Tools

- Predict stock returns (e.g., Chen et al., 2022)—via sentiment identification (Garcia et al., 2023; Chang et al., 2023), understand investor behavior (Chen et al., 2024d), measure market uncertainty (Audrino et al., 2024), help trading (Chen et al., 2024c)
- AI financial analyst (Zhou et al., 2024; Dong, 2024)
- Investment companies' reliance on generative AI (Sheng et al., 2024), analysts' AI usage (Christ et al., 2024)
- Writing academic papers (Novy-Marx and Velikov, 2025) and facilitate academic research (Korinek, 2023)
- Shocks to workers (e.g., Eloundou et al., 2024; Brynjolfsson et al., 2025)
- Analyzing unstructured information (e.g., Cong et al., 2024), e.g., 10-k filings (Shaffer and Wang, 2024; Serafeim, 2024), analyst reports (Lv, 2024; Li et al., 2024; Bastianello et al., 2024), SEC filings (Krockenberger et al., 2024)

LLMs as Economic Agents

- Social experiments as a *homo silicus* (Horton, 2023; Bini et al., 2025)
- Ethics and risk preferences of LLMs (Ouyang et al., 2024)
- Rational budgetary decisions (Chen et al., 2023)
- LLM-based pricing agents (Fish et al., 2024)
- Engage in social interactions (Manning et al., 2024)
- LLMs rely on associative memory to make decisions (Zheng, 2025)
- Professional forecasters (Hansen et al., 2024)

New Data

- Text-based industry classifications (e.g., Hoberg and Phillips, 2016)
- Extracting manager expectations (e.g., Jha et al., 2024b)
- Analyzing conference calls (Jha et al., 2024a)
- Global business networks (Breitung and Müller, 2025)
- Production networks (Fetzer et al., 2024)
- Competitor networks (Hoberg et al., 2024)
- Fed speak (Hansen and Kazinnik, 2023)
- Evaluating innovations (Chen et al., 2019)
- Text algorithms in economics (Ash and Hansen, 2023)

Table of Contents

AI Investment and Firm Performance

Firm value, growth, and product innovation:

- AI as a tool to spur growth: Product Innovation & Process Efficiency
- ↑ market valuation (Ahmadi et al., 2023; Eisfeldt et al., 2023; Babina et al., 2024; Bertomeu et al., 2025; Rock, 2019)
- ↑ growth in sales and employment(Babina et al., 2024)
- ↑ product innovation (Cockburn et al., 2018; Babina et al., 2024)
- ↑ product quality (Fedyk et al., 2022)
- ↓ probability to exit and be merged for large firms (Lu et al., 2024)

Measure of AI investments: job posting and resume data (e.g., Babina et al., 2024)

AI as an External Force

Organizational transformation

- Shifts in workforce skills and flattening hierarchies (Babina et al., 2023a)
- Autonomous vs. non-autonomous AI (Ide and Talamas, forthcoming)
- AI manager (Campello et al., 2023)
- Path dependency (Schubert, 2025)

Risk

- Higher systematic risk (Babina et al., 2023b)
- AI-adoption can backfire—Agency problem (Chen and Han, 2024)
- "AI washing" (Barrios et al., 2024)

Financial markets communications

- Fraud detection (Hobson et al., 2012)
- Startup pitches videos (Hu and Ma, 2024), corporate executive presentations (Cao et al., 2024a), Earnings Conference Q&A (Bai et al., 2023)
- How to talk when a machine is listening? — Sentiment management in disclosures (Cao et al., 2023)

Table of Contents

How AI Influences Households

Financial inclusion

- ML leads to disparity in rates (Fuster et al., 2022)
- Cashless payment adoption increases credit access (e.g., Ouyang, 2021)

Financial services

- Interest rate liberalization through FinTech (Buchak et al., 2021)
- Robo-advising provides some benefits (Rossi and Utkus, 2024; D'Acunto et al., 2019; D'Acunto and Rossi, 2023; Chak et al., 2022), but human experts continue to add unique value (Greig et al., 2024)
- AI advising on soft information and hard decisions (Huang and Ouyang, 2025)

Decision making

- Algorithm aversion on robo-advising (Greig et al., 2024)
- Data privacy and data sharing (Tang, 2019; Chen et al., 2021; Bergemann and Bonatti, 2024)
- Data protection (e.g., Matos and Adjerid, 2022)

Table of Contents

AI as a Shock to Labor Market

- Shifts in labor demand (Jiang et al., 2025b; Acemoglu et al., 2022; Lyonnet and Stern, 2022; Gofman and Jin, 2024) and occupational exposure (Webb, 2019; Jiang et al., 2025a; Hampole et al., 2025)
- Productivity impact (Seamans and Raj, 2018; Alderucci et al., 2020; Eloundou et al., 2024)—randomized experiments (Brynjolfsson et al., 2025; Kanazawa et al., 2022; Noy and Zhang, 2023; Peng et al., 2023).
- Human-AI interactions (Agrawal et al., 2019b), and for high skilled workers (Grennan and Michaely, 2020)
- Risks of AI (Acemoglu, 2021; Acemoglu and Restrepo, 2018), implementation lags (Brynjolfsson et al., 2019), and policy implications (Agrawal et al., 2019a; Furman and Seamans, 2019)
- Accelerate discovery rates in complex knowledge spaces (Agrawal et al., 2018)
- Knowledge production (Abis and Veldkamp, 2024)

Table of Contents

Risk and Challenges in the AI Era

- Algorithmic bias of ChatGPT (Fedyk et al., 2024)
- Look-ahead bias (Glasserman and Lin, 2023) and the "Garbage in, Garbage out" critique of LLM (Bender et al., 2021)
- Regulating algorithmic decisions (e.g., Clark and Hadfield, 2019; Blattner et al., 2021)
- Biases in AI- and ML-generated variables (Battaglia et al., 2024), substantial gap between average accuracy and self-reported confidence (Yoo, 2024; Chen et al., 2024a)
- AI collusion (e.g., Johnson and Sokol, 2020; Dou et al., 2024)
- Setting AI standards (Canayaz and Wang, 2024)
- The general data protection regulation (GDPR) (e.g., Matos and Adjerid, 2022; Johnson et al., 2023; Goldberg et al., 2024)
- Managers' perceptions on ethical issues related to AI (Cuéllar et al., 2024)

Future Directions and Open Questions

- **Causality over Correlation**
  How can we isolate the *causal impact* of AI on firms, markets, and households?
- **Interpretability & Economic Insight**
  Can we “open the black box” to uncover underlying *economic mechanisms*?
- **LLMs as Economic Agents**
  What are the potentials & limits of using LLMs in *simulated economic experiments*?
- **Governance and Regulation**
  How do we build regulatory frameworks to address bias, collusion, and privacy?
- **Welfare & Long-Term Structure**
  What are AI’s *structural effects* on inequality, competition, and market stability?
- **Human-AI Complementarity**
  Which financial tasks should remain human-led? How do we build *hybrid systems* that foster trust and inclusion?

Other Related Review Articles

- Effect of AI on the wider economy (Furman and Seamans, 2019)
- Generative AI as a research topic in finance and as a technology shock to methods for financial research (Eisfeldt and Schubert, 2025)
- Natural language processing (NLP) tools used in financial economics research (Hoberg and Manela, 2025)
- AI as a tool to analyze alternative data (Cao et al., 2024b)
- LLM-based multi-agents (Guo et al., 2024)
- ML in portfolio decisions (Guidolin et al., 2024)

Table of Contents

Why Machine Learning in Asset Pricing?

Traditional Challenges:

- High-dimensional predictor space (100+ firm characteristics)
- Low signal-to-noise ratios in return prediction
- OLS overfits when p (predictors) large relative to n (observations)
- The "factor zoo" problem: 316 documented anomalies (Harvey et al., 2016)

What ML Brings:

- Flexible functional forms capture nonlinearities
- Automated variable selection handles high dimensions
- Regularization prevents overfitting
- Out-of-sample validation ensures robustness

The Asset Pricing Problem

Goal: Predict individual stock excess returns

$$r_{i,t+1} = \mathbb{E}_t[r_{i,t+1}] + \varepsilon_{i,t+1}$$

Conditional Expectation Model:

$$\mathbb{E}_t[r_{i,t+1}] = f(z_{i,t}; \theta)$$

where $z_{i,t}$ = vector of characteristics (size, value, momentum, profitability, etc.)

Approaches:

- **Traditional:** $f(z_{i,t}; \theta) = z_{i,t}'\beta$ (linear, restrictive)
- **Machine Learning:** $f(\cdot)$ highly nonlinear, captures interactions

Regularization Methods: Ridge and LASSO

Ridge Regression (L2):

$$\hat{\beta}^{ridge} = \arg\min_{\beta} \left\{ \sum_{i=1}^{n} (r_i - z_i'\beta)^2 + \lambda \sum_{j=1}^{p} \beta_j^2 \right\}$$

- Shrinks all coefficients toward zero
- Handles multicollinearity
- No variable selection

LASSO (L1):

$$\hat{\beta}^{LASSO} = \arg\min_{\beta} \left\{ \sum_{i=1}^{n} (r_i - z_i'\beta)^2 + \lambda \sum_{j=1}^{p} |\beta_j| \right\}$$

- Sets many coefficients exactly to zero (sparsity)
- Performs automatic variable selection
- High-frequency prediction (Chinco et al., 2019), factor testing (Feng et al., 2020)

Elastic Net

Combines L1 and L2 penalties:

$$\hat{\beta}^{EN} = \arg\min_{\beta} \left\{ \sum_{i=1}^{n} (r_i - z_i'\beta)^2 + \lambda \left[\alpha \|\beta\|_1 + (1-\alpha)\|\beta\|_2^2 \right] \right\}$$

Key Properties:

- Performs variable selection (via L1)
- Handles correlated predictors (via L2)
- Encourages grouping effect
- Two tuning parameters: λ (penalty strength) and α (L1/L2 mix)

Applications:

- Market return prediction (Dong et al., 2022)
- Mutual fund selection with positive alpha (DeMiguel et al., 2023)

Principal Component Analysis (PCA)

Idea: Extract latent factors from high-dimensional characteristics

Method: Find orthogonal directions maximizing variance

$$PC_j = Zw_j, \quad j = 1, \ldots, k$$

where w_j are eigenvectors of covariance matrix $Z'Z$

Properties:

- Unsupervised (doesn't use returns)
- Orthogonal components by construction
- May capture variance but miss return predictability

Asset Pricing Extensions:

- **IPCA** (Kelly et al., 2019): Characteristics instrument time-varying factor loadings
- **RP-PCA** (Lettau and Pelger, 2020): Identifies factors explaining both covariance and returns
- **Group LASSO** (Freyberger et al., 2020): Nonparametric selection, finds 11-14 key characteristics

Partial Least Squares (PLS)

Key Difference from PCA: Supervised dimension reduction

- PCA: Maximizes variance in predictors (unsupervised)
- PLS: Maximizes covariance with returns (supervised)

Algorithm: Sequentially extract components

$$w_j = \arg\max_{\|w\|=1} \mathrm{Cov}(Z^{(j)} w, r^{(j)})$$

Advantage: Explicitly targets predictive power for returns

Performance: Outperforms PCA and Fama-MacBeth regression for return prediction (Light et al., 2017)

Decision Trees and Random Forests

Decision Trees:

- Recursive partitioning: Split predictor space into regions
- Prediction: Average return within each region
- Automatically detect interactions (e.g., value effect stronger for small stocks)
- Problem: High variance, unstable

Random Forests: Ensemble of many trees

- Bootstrap sampling + random feature selection
- Average predictions across trees
- Dramatically reduces variance
- Among the top performers: $R^2_{OOS} \approx 0.33\%$ monthly (Gu et al., 2020)

Gradient Boosting

Key Difference: Sequential learning (vs. parallel in Random Forests)

Algorithm:

1. Initialize: $\hat{f}^{(0)}(z) = \bar{r}$
2. For $m = 1$ to M:
   - Compute residuals: $u^{(m)} = r - \hat{f}^{(m-1)}(z)$
   - Fit shallow tree to residuals
   - Update: $\hat{f}^{(m)} = \hat{f}^{(m-1)} + \nu \cdot h^{(m)}$

Regularization:

- Learning rate ν (shrinkage): typical 0.01–0.1
- Shallow trees (depth 3–6): weak learners
- Subsampling: stochastic gradient boosting

Variants: XGBoost, LightGBM, CatBoost

Neural Networks: Feedforward Architecture

Basic Structure: Composition of nonlinear transformations

$$\text{Hidden Layer: } h^{(1)} = g\left(W^{(1)}z + b^{(1)}\right)$$

$$\text{Output: } \hat{r} = W^{(out)}h^{(L)} + b^{(out)}$$

Activation Functions:

- **ReLU:** $g(x) = \max(0, x)$ — most popular
- Creates piecewise-linear, highly flexible functions

Regularization:

- **Dropout:** Randomly drop neurons (sets to zero) during training
- **L2 penalty:** Penalize large weights
- **Early stopping:** Stop when validation error increases

Neural Networks in Asset Pricing

Direct Return Prediction (Gu et al., 2020):

- Architecture: 1-5 hidden layers, 2-32 neurons each
- Out-of-sample $R^2 \approx 0.40\%$ monthly (NN3, top performer)
- Captures nonlinear effects and interactions

SDF Estimation (Chen et al., 2024b):

- **Three-network system:**
  - LSTM (Long Short-Term Memory): Extracts hidden states from macro time series
  - Feedforward NN: Estimates SDF weights ω from firm characteristics + macro states
  - GAN (Generative Adversarial Network): Adversarially constructs optimal test assets (maximizes pricing errors)
- **Performance (annual, out-of-sample):**
  - Sharpe ratio: 2.6 GAN (vs. 1.5 FFN, 0.8 FF5)
  - Explained variation of individual stock returns: 8% (2$\times$ benchmarks)
  - Cross-sectional R^2: 23%
- **Key innovation:** No-arbitrage condition as criterion function; adversarial approach based on Hansen & Jagannathan (1997) minimax objective

Advanced Architectures: Autoencoders

Standard Autoencoder: Encoder (compress) + Decoder (reconstruct)

Conditional Autoencoder for Asset Pricing (Gu et al., 2021):

$$r_{i,t+1} = \beta_i(z_{i,t})' f_{t+1} + \varepsilon_{i,t+1}$$

where $\beta_i(z_{i,t}) = \text{NeuralNet}(z_{i,t})$ (nonlinear in characteristics)

Key Results (Out-of-Sample):

- **Managed portfolios:** Total $R^2 = 92\%$ (IPCA, 3 factors) vs. 70% (FF 3-factor)
- **Individual stocks:** Total $R^2 = 14\%$ (CA1, 6 factors) vs. 3% (FF 3-factor)
- **Predictive R^2:** 0.58% (CA2) vs. negative for FF models
- **Sharpe ratio:** 2.63 (CA2 equal-weighted) vs. -0.40 (FF)

Economic Insight: Most return predictability from characteristics works through time-varying nonlinear betas, not alpha

Transformers and Attention Mechanisms

Motivation: Traditional models treat assets independently. But returns are interdependent (co-movement, contagion, spillovers)

Self-Attention: Each asset "attends to" other relevant assets

$$\text{Attention score: } a_{ij} = \frac{\exp(Q_i' K_j / \sqrt{d})}{\sum_k \exp(Q_i' K_k / \sqrt{d})}$$

Transformer-Based SDF (Kelly et al., 2025):

- Cross-asset information sharing via attention
- 30% lower out-of-sample pricing errors than NNs
- Attention weights reveal economic linkages

The Complexity Paradox: Millions of parameters, yet outperforms simpler models out-of-sample

The Factor Zoo Problem

The Challenge: 316 published return predictors (Harvey et al., 2016)

- Multiple testing bias (false discoveries)
- Publication bias (file drawer problem)
- Data mining (p-hacking)

ML Solutions:

1. Robust Testing (Feng et al., 2020):

- Double-selection LASSO + Fama-MacBeth
- Most new factors redundant after controlling for existing factors

2. Sparse SDF (Kozak et al., 2020):

- Characteristics-sparse SDFs fail, but sparsity works in PC space (6-10 PCs)

3. Nonlinear Factor Models (Gu et al., 2021):

- Characteristics predict returns via time-varying nonlinear factor loadings, not alpha

4. Asset Pricing Trees (Bryzgalova et al., 2025):

- Tree-based algorithm to create optimal test assets
- Captures high-dimensional interactions, better SDF spanning

Overfitting and Out-of-Sample Validation

Bias-Variance Trade-off:

$$\text{Total Error} = \text{Bias}^2 + \text{Variance} + \text{Irreducible Error}$$

- Simple models: Low variance, high bias (underfitting)
- Complex models: Low bias, high variance (overfitting)

Cross-Validation for Time Series:

- **Never** use future data to predict past (look-ahead bias)
- **Expanding window:** Training set grows over time
- **Validation:** Tune hyperparameters on validation set
- **Test:** Final evaluation on held-out test set

Regularization Techniques:

- L1/L2 penalties, Dropout, Early stopping
- Ensemble methods (bagging, boosting)

Interpretability vs. Performance

The Spectrum:

Method	Interpretability	Performance	Application
OLS	High	Low	Research
LASSO	High	Medium	Variable selection
EBM	High	High	Best of both worlds
Random Forest	Low	High	Prediction
Deep NN	Very Low	High	Max accuracy

Interpretability Techniques:

- Feature importance (tree-based models)
- Partial Dependence Plots (PDPs)
- SHAP values (game-theoretic approach)
- Explainable Boosting Machines (EBMs)

Implementation Challenges

1. **Data Quality:**
   - Missing values, outliers, survivorship bias
   - Solution: Winsorization, tree-based methods (handle missingness)
2. **Transaction Costs:**
   - Paper profits $\neq$ real profits
   - Must incorporate: Bid-ask spread, price impact, commissions
3. **Model Decay:**
   - Structural breaks, strategy crowding
   - Solution: Online learning, continuous monitoring
4. **Risk Management:**
   - Unintended factor exposures, concentration risk
   - Solution: Factor-neutral constraints, position limits
5. **Regulatory Compliance:**
   - Explainability requirements
   - Solution: Use interpretable models (EBM) or tools (SHAP)

Generative AI in Asset Pricing

LLM Applications:

- Predicting stock returns via sentiment (Lopez-Lira and Tang, 2023; Chen et al., 2022)
- Textual factors (Cong et al., 2024)
- Writing anomalies papers (Novy-Marx and Velikov, 2025)

Performance:

- Mixed results: Some find good performance (Chen et al., 2022; Kelly et al., 2025), others find miscalibration (Chen et al., 2024a)
- Context-dependent: Works better for sentiment-driven moves

Challenges:

- Hallucinations and errors
- Look-ahead bias (Glasserman and Lin, 2023)
- Model opacity and alignment issues

Future Directions in ML for Asset Pricing

1. **Causal Inference:**
   - Move beyond prediction to causal understanding
   - Causal forests, double machine learning
2. **Alternative Data:**
   - Multimodal learning: Text + images + networks
   - Satellite imagery, social media, geospatial data
3. **Improving Interpretability:**
   - Better tools for understanding complex models
   - Hybrid models: Theory + ML flexibility
4. **Real-Time Learning:**
   - Online learning, continuous adaptation
   - Handle regime changes and structural breaks
5. **Regulatory Frameworks:**
   - Balance innovation with risk management
   - Audit mechanisms for algorithmic decisions

Beyond Equity Markets

Fixed Income (Bianchi et al., 2021):

- Extreme trees and deep NNs outperform linear models for Treasury bonds
- Group-ensembled NNs leverage economic priors (macro categories)
- Nonlinearities *within* groups drive improved prediction

International Markets (Leippold et al., 2022):

- Neural networks excel in Chinese equity markets
- Return predictability stronger in China than US
- ML profitable even through 2015 crash and COVID

Challenging Weak-Form EMH (Murray et al., 2024):

- CNN-LSTM on 12-month returns: 1.08% monthly long-short return (Sharpe 0.78)
- Robust in large caps and recent period
- Driven by nonlinear interactions, distinct from momentum/reversal
- Technical patterns contain alpha unexplained by risk, violating weak-form EMH

Paper 1: Novy-Marx & Velikov (2024)

"AI-Powered (Finance) Scholarship"

Research Question:

- Can LLMs automate academic research production from hypothesis to full paper?
- What are implications for academic integrity and peer review?

Key Contributions:

1. **Complete automation pipeline**: Data mining $\rightarrow$ testing $\rightarrow$ paper writing
2. **Industrial scale**: 380 complete academic papers generated
3. **Multiple theoretical frameworks**: Same results, different "stories"
4. **Cautionary tale**: Demonstrates potential for HARKing industrialization

Novy-Marx & Velikov: Methodology

Step 1: Signal Mining

- 30,000+ potential predictors from accounting data (COMPUSTAT)
- Ratio and difference-type signals: X/Y, $\Delta X/lag(Y)$

Step 2: Rigorous Filtering ("Assaying Anomalies" protocol)

- 30,000 candidates $\rightarrow$ 95 signals passing all criteria
- Require: significant across decile/quintile sorts, VW/EW portfolios, FF6 alphas
- Benchmark against 200+ documented anomalies

Step 3: AI Paper Generation (Claude Opus 4.1)

- Signal naming: e.g., "Liquidity Leverage Intensity" (ACO/SEQ)
- Four distinct theoretical frameworks per signal:
    1. Unrestricted (general economic mechanisms)
    2. Behavioral (slow diffusion of information)
    3. Production-based asset pricing
    4. Consumption-based asset pricing
- Complete papers: intro, data, results, conclusion, references

Novy-Marx & Velikov: Key Findings

Quality of Data-Mined Signals:

- t-statistics distribution matches published anomalies
- Equal-weighted: 87% overlap with published literature
- Value-weighted: Shorter right tail but similar modes
- **Implication**: Peer review may not distinguish mining from discovery

Quality of AI-Generated Content:

- Readability (Flesch-Kincaid): 16–18 years (college graduate level)
- Remarkably consistent across theoretical frameworks
- Convincing prose with (occasionally hallucinated) citations
- 30 pages each: intro, methods, 8 figures, 5 tables

Efficiency:

- Data mining + validation: $\sim$1 day
- 380 complete papers generated: 12 hours
- Cost: Minimal computational expense

Novy-Marx & Velikov: Implications

Existential Threats:

- **Peer review**: Could overwhelm journals (380 papers = very high review costs)
- **Citation gaming**: Automated papers cite strategically, inflating metrics
- **Academic standards**: Distinction between discovery and fabrication blurs

The Good:

- Democratizes research production
- Accelerates hypothesis testing
- May speed market efficiency

Recommendations:

1. **Full accountability**: Authors responsible for all content (including AI-generated)
2. **Enhanced validation**: Citation verification, theoretical consistency checks
3. **Out-of-sample emphasis**: Focus on practical significance and novel predictions

Key Insight: Newton observed an apple, then theorized. Post-hoc theorizing is scientific practice—but industrial-scale automation crosses a line.

Paper 2: Gu, Kelly & Xiu (2020)

"Empirical Asset Pricing via Machine Learning" (*Review of Financial Studies*)

Data & Sample:

- **Universe**: 30,000+ stocks (NYSE/AMEX/NASDAQ), 1957–2016
- **Predictors**: 94 firm characteristics $\times$ (8 macro variables + 1) + 74 industry dummies = 920 signals
- **Sample split**: 18 years training (1957–1974), 12 years validation (1975–1986), 30 years out-of-sample testing (1987–2016)

Methods Compared:

- **Linear**: OLS, Elastic Net
- **Dimension reduction**: PCR, PLS
- **Trees**: Random Forest (bagging), GBRT (boosting)
- **Neural Networks**: NN1–NN5 (1 to 5 hidden layers, ReLU + dropout)

Gu et al.: Detailed Results

Out-of-Sample R^2 (Monthly Stock-Level):

Method	All Stocks	Top 1,000	Bottom 1,000
OLS-3	0.16%	0.31%	0.17%
Elastic Net	0.11%	0.25%	0.20%
PCR	0.26%	0.06%	0.34%
PLS	0.27%	-0.14%	0.42%
Random Forest	0.33%	0.63%	0.35%
GBRT	0.34%	0.52%	0.32%
NN1	0.33%	0.49%	0.38%
NN3	**0.40%**	**0.70%**	0.45%
NN5	0.36%	0.64%	0.42%

Key Patterns:

- Nonlinear methods dominate: NN3 achieves 0.40% vs linear 0.11–0.27%
- Large stocks: ML advantage much stronger (NN3: 0.70% vs OLS-3: 0.31%)
- Optimal depth: NN3 peaks; NN4-5 show diminishing returns (overfitting)

Gu et al.: Economic Significance

Portfolio Strategies:

Market Timing (S&P 500):

- Buy-and-hold Sharpe ratio: 0.51
- NN forecast Sharpe ratio: 0.77 (+51% improvement)

Long-Short Decile Spread (Value-Weighted):

- OLS-3 strategy: SR = 0.61, Mean return = 0.94%/month
- Neural Network (NN3): SR = 1.35, Mean return = 2.12%/month
- SR improvement: +121% (more than doubling performance)
- Turnover: 58% (OLS-3), 124% (NN3) monthly

Where Does ML Excel?

- Large-cap stocks: ML advantage strongest (NN3: 0.70% vs OLS: 0.31%)
- Portfolio-level predictions stronger than individual stocks
- Captures complex interactions missed by linear models

Gu et al.: Variable Importance & Nonlinearities

Top Predictor Categories (Consensus Across Methods):

1. **Price trends**: 12-month momentum, short-term reversal (1-month), industry momentum, momentum change, max return, long-term reversal
2. **Liquidity**: Market cap, dollar volume, turnover, bid-ask spread, Amihud illiquidity, zero trading days
3. **Risk/Volatility**: Total volatility, idiosyncratic volatility, market beta, beta squared
4. **Valuation & Fundamentals**: Earnings-to-price, sales-to-price, asset growth, earnings increases

Key Nonlinear Effects & Interactions (NN3):

- **Size**: Smaller stocks earn 2.4% more (median $\rightarrow$ 20th percentile)
- **Volatility**: Higher vol stocks earn 3.0% less (median $\rightarrow$ 80th percentile)
- **Size** $\times$ **Momentum**: Momentum works best for large stocks
- **Size** $\times$ **Reversal**: Linear for small stocks, concave for large stocks
- Linear models miss these patterns (prefer zero coefficients)

Gu et al.: Theoretical Interpretation & Limitations

What ML Achieves:

- Measures conditional expected returns: $\mathbb{E}_t[r_{i,t+1}|\mathbf{z}_{i,t}]$
- Leverages regularization (penalization, validation, ensembles) to avoid overfit
- Captures nonlinear interactions among 100+ predictors
- Combines all sources: systematic risk + idiosyncratic risk + potential mispricing

Key Findings:

- **Trees & NNs dominate**: Allow complex predictor interactions
- **Shallow beats deep**: NN3 optimal; NN4-5 overfit (low signal-to-noise in finance)
- **Best for large/liquid stocks**: Portfolio predictions stronger than individual stocks

Limitations:

- **Complexity**: Models are "black boxes"—inspectable but opaque
- **Data intensive**: Requires long time series, many stocks, quality data
- **Interpretation**: Difficult to extract simple economic mechanisms

Paper 3: Kelly et al. (2025)

"Artificial Intelligence Asset Pricing Models"

Research Question:

- Can transformer architectures (from ChatGPT/LLMs) improve asset pricing by enabling cross-asset information sharing?

Key Innovation: The AIPM (AI Pricing Model)

- Embeds transformer in the stochastic discount factor (SDF)
- **Cross-asset attention mechanism**: Asset i's return prediction uses information from *all* assets
- Inspired by how words need context in language ("bank" = financial or river?)

Key Insight:

Traditional: Return(i) = f(Characteristics(i))

AIPM: Return(i) = f(All Assets' Characteristics)

Kelly et al.: Model Architecture

SDF Representation: $m_{t+1} = 1 - w(X_t)' R_{t+1}$ where $w(X_t)$: conditional portfolio weights; X_t: $N \times D$ matrix of characteristics

Linear Portfolio Transformer: $w_t = A_t X_t \lambda = (X_t W X_t') X_t \lambda$

Attention Matrix: $A_t = X_t W X_t'$

- Measures similarity between assets
- Dynamic, conditional on characteristics
- Enables information sharing

Nonlinear Portfolio Transformer (K blocks):

1. **Multi-head attention**: $A(Y) = \sum_{h=1}^{H} \sigma(Y W_h Y') Y V_h$ (softmax for selectivity)
2. **Feed-forward network**: $F(Y) = \max[0, Y W_1 + b_1] W_2 + b_2$ (nonlinear transformations)
3. **Residual connections**: Skip connections for stability

Scaling: 1–10 blocks, up to 1 million parameters

Kelly et al.: Empirical Setup & Results

Data:

- Period: 1963–2022 (US stocks)
- Characteristics: 132 from Jensen et al. (2023)
- Training: 60-month rolling windows

Out-of-Sample Sharpe Ratios (1968–2022):

Model	Sharpe Ratio	vs. FF6
FF6	1.05	Baseline
HXZ (q-factors)	1.80	+71%
BSV (132 linear factors)	3.60	+243%
DKKM (shallow NN)	3.91	+272%
Linear Attention	3.89	+270%
MLP (deep NN)	4.31	+310%
Transformer (10 blocks)	**4.57**	**+335%**

Key Finding: Attention adds 0.3–0.7 to Sharpe ratio over best NN

Kelly et al.: Cross-Asset Information & Complexity

Principal Portfolios Analysis:

- **Symmetric component** (own-asset effects): SR = 3.23
- **Anti-symmetric component** (pure cross-prediction): SR = 3.10
- Correlation between components: only 32%!
- **Interpretation**: Cross-asset prediction captures industry spillovers, supply chains, correlated fundamentals

The Virtue of Complexity:

- Linear transformer (129 heads): SR = 3.89
- Nonlinear transformer (10 blocks, $\sim$1M params): SR = 4.57
- **No plateau yet!** Performance still increasing with complexity

Size Effects:

- **Mega-cap stocks only**: Linear SR = 1.05–1.18; Transformer SR = 1.84 (1.6–1.8$\times$)
- Attention crucial even for liquid stocks

Kelly et al.: Theoretical Insights & Implications

UNIVERSITY OF OXFORD Saïd Business School

Why Does Attention Work?

1. **Noise reduction**: Individual characteristics are noisy proxies of expected returns; cross-sectional information sharing refines them
2. **Cross-sectional dependence**: Sufficient dependence among true expected returns enables information borrowing across assets
3. **Context-aware prediction**: Like words in LLMs, asset returns understood through surrounding asset context

Connection to Factor Timing:

- Linear attention model is a high-dimensional factor timing representation
- Optimizes timing of all factors jointly (vs. one-at-a-time in prior work)

Practical Implications:

- **OOS performance**: Sharpe ratio increases from 3.8 (1 block) to 4.6 (10 blocks)
- **Virtue of complexity**: Performance still improving at $\sim$1M parameters
- **Information sharing**: Cross-asset attention dominates own-asset models

Paper 4: Bell, Kakhbod, Lettau & Nazemi (2024)

"Glass Box Machine Learning and Corporate Bond Returns"

Research Question:

- Can we achieve SOTA ML performance while maintaining full interpretability?
- What drives corporate bond returns? (nonlinear relationships, interactions)

The Dilemma:

- **Black box** (RF, XGBoost, NN): High accuracy, zero interpretability
- **Glass box** (OLS, LASSO): Interpretable, low accuracy

Solution: Explainable Boosting Machine (EBM)

$$r_{l,t+1} = \beta + \sum_{i=1}^{N} f_i(x_{i,l,t}) + \sum_{i>j} f_{ij}(x_{i,l,t}, x_{j,l,t}) + e_{l,t+1}$$

- Additive structure: Each f_i and f_{ij} visualizable
- Learned via gradient boosting on decision trees
- Bootstrap for standard errors

Bell et al.: Data & Performance

Dataset:

- **Period**: July 2002–August 2020
- **Universe**: 1,207 firms, 106,265 firm-months
- **Predictors (81 total)**: Firm-level (41) + Market-level (40)
- **Sample Split**: Train (7/2002–6/2010), Valid (7/2010–6/2011), Test (7/2011–8/2020)

Out-of-Sample R^2 (Jul 2011–Aug 2020):

Model	Training R^2	Testing R^2
OLS	31.2%	-2.3%
LASSO	17.5%	8.0%
Random Forest	46.9%	13.3%
XGBoost	42.5%	12.0%
EBM	24.8%	**12.1%**

Key Insight: EBM achieves 100% of XGBoost accuracy with full interpretability.

Bell et al.: Most Important Predictors

Top 5 Variables (by Mean Absolute Score):

1. **ΔUNCf** (Financial Uncertainty Change): Most important
2. **ΔUNC** (Macro Uncertainty Change): Jurado et al. (2015) measure
3. **TERM** (Term Structure Factor): Long-short Treasury spread
4. **LTREVB** (Long-term Reversal): Bond-specific factor
5. **spread** (Bond Yield Spread): Over Treasury curve

Key Pattern: Changes in uncertainty matter more than levels.

- Correlation(ΔUNCf$_t$, r_{t+1}) = -0.50 (t-stat = -8.93)
- Correlation(UNCf$_t$, r_{t+1}) = -0.06 (insignificant)

Interpretation: Markets react to *changes* in uncertainty, not levels. Large increases/decreases in financial uncertainty predict significant bond return movements.

Bell et al.: Nonlinear Shape Functions

1. **Financial Uncertainty (ΔUNCf):**
   - Monotonic negative, **nonlinear** effect on bond returns
   - Large increases: -66 bps; Large decreases: $+64$ bps
   - Rising uncertainty lowers returns, falling uncertainty boosts returns
2. **Term Structure (TERM):**
   - Highly **asymmetric**: Large positive TERM $\to$ +100 bps effect
   - **No corresponding negative effect** for low TERM values
   - Right tail only—captures flight-to-quality and duration effects
3. **Macro Uncertainty (ΔUNC):**
   - **Step function**: Negligible for most values
   - Large negative effects only at extreme increases (right tail)
   - Asymmetric: Market reacts more to rising than falling uncertainty
   - Triggered during 2008 financial crisis and 2020 COVID

Bell et al.: Interaction Effects

Most Important: ΔUNCf $\times$ spread

Pattern Discovered:

- High-spread bonds **super-sensitive** to uncertainty changes
- Uncertainty falls: Low-spread bonds gain little, high-spread bonds gain much larger
- **Asymmetric**: Falling uncertainty effect $>$ Rising uncertainty effect
- Represents **outsized price recovery** for risky bonds as uncertainty diminishes

Economic Story:

- **Risk-on/risk-off dynamics**: High-spread bonds sold first in stress, bought first in recovery
- Credit risk amplification and flight-to-quality reversal

Also Important: ΔUNCf $\times$ sprxrtg (spread $\times$ rating)

- Low-grade high-spread bonds benefit most from falling uncertainty
- Similar asymmetric pattern with disproportionate gains

Bell et al.: Portfolio Performance & Implications

Long-Short Portfolio (Based on EBM Predictions):

- Mean return: 8.04% annually (0.67% monthly)
- Sharpe ratio: 0.38 monthly (1.32 annualized)
- Monotonic pattern across deciles

EBM-pos Strategy (Market Timing):

- If $\mathbb{E}[r] < 0$: Invest in risk-free rate instead
- Avoids downturns (2008 crisis, 2020 COVID)

Glass Box Advantages:

1. **Exact shapes**: See nonlinearities, asymmetries, thresholds
2. **Interactions**: Which variables interact, how, and when
3. **Heterogeneity**: Different firms respond differently
4. **Regulatory compliance**: Fully explainable
5. **Strategy design**: Target effects (e.g., high-yield bonds when ΔUNCf<0)

Key Takeaway: EBM matches XGBoost performance with full interpretability

References I

Abis, Simona and Laura Veldkamp (2024). "The changing economics of knowledge production". In: *The Review of Financial Studies* 37.1, pp. 89–118.

Acemoglu, Daron (2021). *Harms of AI*. Tech. rep. National Bureau of Economic Research.

Acemoglu, Daron, David Autor, Jonathon Hazell, and Pascual Restrepo (2022). "Artificial intelligence and jobs: Evidence from online vacancies". In: *Journal of Labor Economics* 40.S1, S293–S340.

Acemoglu, Daron and Pascual Restrepo (2018). "Artificial intelligence, automation, and work". In: *The economics of artificial intelligence: An agenda*. University of Chicago Press, pp. 197–236.

Agrawal, Ajay, Joshua Gans, and Avi Goldfarb (2019a). "Economic policy for artificial intelligence". In: *Innovation policy and the economy* 19.1, pp. 139–159.

Agrawal, Ajay, Joshua S Gans, and Avi Goldfarb (2019b). "Artificial intelligence: the ambiguous labor market impact of automating prediction". In: *Journal of Economic Perspectives* 33.2, pp. 31–50.

References II

- Agrawal, Ajay, John McHale, and Alexander Oettl (2018). "Finding needles in haystacks: Artificial intelligence and recombinant growth". In: *The economics of artificial intelligence: An agenda*. University of Chicago Press, pp. 149–174.
- Ahmadi, A, AMBRUS KECSKÉS, RONI MICHAELY, and PA NGUYEN (2023). "Producing AI Innovation and Its Value Implications". In.
- Alderucci, Dean, Lee Branstetter, Eduard Hovy, Andrew Runge, and Nikolas Zolas (2020). "Quantifying the impact of AI on productivity and labor demand: Evidence from US census microdata". In: *Allied social science associations—ASSA 2020 annual meeting*.
- Ash, Elliott and Stephen Hansen (2023). "Text algorithms in economics". In: *Annual Review of Economics* 15.1, pp. 659–688.
- Audrino, Francesco, Jessica Gentner, and Simon Stalder (2024). "Quantifying uncertainty: a new era of measurement through large language models". In.

References III

- Babina, Tania, Anastassia Fedyk, Alex He, and James Hodson (2024). "Artificial intelligence, firm growth, and product innovation". In: *Journal of Financial Economics* 151, p. 103745.
- Babina, Tania, Anastassia Fedyk, Alex X He, and James Hodson (2023a). *Firm investments in artificial intelligence technologies and changes in workforce composition.* Vol. 31325. National Bureau of Economic Research.
- Babina, Tania, Anastassia Fedyk, Alex Xi He, and James Hodson (2023b). "Artificial intelligence and firms' systematic risk". In: *Available at SSRN.*
- Bai, John Jianqiu, Nicole M Boyson, Yi Cao, Miao Liu, and Chi Wan (2023). "Executives vs. chatbots: Unmasking insights through human-AI differences in earnings conference Q&A". In: *Northeastern U. D'Amore-McKim School of Business Research Paper* 4480056.

References IV

Barrios, John Manuel, John L. Campbell, Ryan G. Johnson, and Y. Christine Liu (2024). *Signals or Smoke? The Determinants and Informativeness of Corporate Artificial Intelligence (AI) Disclosures*. Available at SSRN: https://ssrn.com/abstract=5133107.

Bastianello, Francesca, Paul H Décaire, and Marius Guenzel (2024). "Mental Models and Financial Forecasts". In: *Marius, Mental Models and Financial Forecasts (October 30, 2024)*.

Battaglia, Laura, Timothy Christensen, Stephen Hansen, and Szymon Sacher (2024). "Inference for Regression with Variables Generated by AI or Machine Learning". In.

Bender, Emily M., Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell (2021). "On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?" In: *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. New York, NY, USA: Association for Computing Machinery, pp. 610–623.

References V

- Bergemann, Dirk and Alessandro Bonatti (2024). “Data, competition, and digital platforms”. In: *American Economic Review* 114.8, pp. 2553–2595.
- Bertomeu, Jeremy, Yupeng Lin, Yibin Liu, and Zhenghui Ni (2025). “The Impact of Generative AI on Information Processing: Evidence from the Ban of ChatGPT in Italy”. In: *Journal of Accounting and Economics*, p. 101782.
- Bianchi, Daniele, Matthias Büchner, and Andrea Tamoni (2021). “Bond risk premiums with machine learning”. In: *The Review of Financial Studies* 34.2, pp. 1046–1089.
- Bini, Pietro, Lin William Cong, Xing Huang, and Lawrence J Jin (2025). “Behavioral Economics of AI: LLM Biases and Corrections”. In: *Available at SSRN 5213130.*
- Blattner, Laura, Scott Nelson, and Jann Spiess (2021). “Unpacking the black box: Regulating algorithmic decisions”. In: *arXiv preprint arXiv:2110.03443.*
- Breitung, Christian and Sebastian Müller (2025). “Global Business Networks”. In: *Journal of Financial Economics* 166, p. 104007.

References VI

- Brynjolfsson, Erik, Danielle Li, and Lindsey Raymond (2025). "Generative AI at work". In: *The Quarterly Journal of Economics*, qjae044.
- Brynjolfsson, Erik, Daniel Rock, and Chad Syverson (2019). "Artificial intelligence and the modern productivity paradox". In: *The economics of artificial intelligence: An agenda* 23.2019, pp. 23–57.
- Bryzgalova, Svetlana, Markus Pelger, and Jason Zhu (2025). "Forest through the trees: Building cross-sections of stock returns". In: *The Journal of Finance* 80.5, pp. 2447–2506.
- Buchak, Greg, Jiayin Hu, and Shang-Jin Wei (2021). *FinTech as a financial liberator.* Tech. rep. National Bureau of Economic Research.
- Campello, Murillo, Lin William Cong, and Luofeng Zhou (2023). "AlphaManager: A Data-Driven-Robust-Control Approach to Corporate Finance". In: *Available at SSRN 4590323.*
- Canayaz, Mehmet and Zhe Wang (2024). "Crafting an AI Compass: The Influence of Global AI Standards on Firms". In: *Available at SSRN 4481608.*

References VII

- Cao, Sean, Yichen Cheng, Meng Wang, Yusen Xia, and Baozhong Yang (2024a). "Visual information and ai divide: Evidence from corporate executive presentations". In: *SSRN Electronic Journal*.
- Cao, Sean, Wei Jiang, Baozhong Yang, and Alan L Zhang (2023). "How to talk when a machine is listening: Corporate disclosure in the age of AI". In: *The Review of Financial Studies* 36.9, pp. 3603–3642.
- Cao, Sean Shun, Wei Jiang, Lijun (Gillian) Lei, and Qing (Clara) Zhou (2024b). "Applied AI for finance and accounting: Alternative data and opportunities". In: *Pacific-Basin Finance Journal* 84, p. 102307.
- Chak, Ida, Karen Croxson, Francesco D'Acunto, Jonathan Reuter, Alberto G Rossi, and Jonathan M Shaw (2022). *Improving household debt management with robo-advice*. Tech. rep. National Bureau of Economic Research.
- Chang, Anne, Xi Dong, Xiumin Martin, and Changyun Zhou (2023). "AI democratization, return predictability, and trading inequality". In: *Available at SSRN 4543999*.

References VIII

- Chen, Hui, Antoine Didisheim, and Luciano Somoza (2024a). "Out of the (Black) Box: AI as Conditional Probability". In: *Available at SSRN.*
- Chen, Joanne and Brandon Yueyang Han (2024). "Corporate AI Backfire and Potential Remedies". In: *Available at SSRN.*
- Chen, Long, Yadong Huang, Shumiao Ouyang, and Wei Xiong (2021). *The data privacy paradox and digital demand.* Tech. rep. National Bureau of Economic Research.
- Chen, Luyang, Markus Pelger, and Jason Zhu (2024b). "Deep learning in asset pricing". In: *Management Science* 70.2, pp. 714–750.
- Chen, Mark A, Qinxi Wu, and Baozhong Yang (2019). "How valuable is FinTech innovation?" In: *The Review of financial studies* 32.5, pp. 2062–2106.
- Chen, Shuaiyu, T Clifton Green, Huseyin Gulen, and Dexin Zhou (2024c). "What Does ChatGPT Make of Historical Stock Returns? Extrapolation and Miscalibration in LLM Stock Return Forecasts". In: *arXiv preprint arXiv:2409.11540.*

References IX

- Chen, Shuaiyu, Lin Peng, and Dexin Zhou (2024d). "Wisdom or Whims? Decoding Investor Trading Strategies with Large Language Models". In: *Decoding Investor Trading Strategies with Large Language Models (June 19, 2024).*
- Chen, Yifei, Bryan T Kelly, and Dacheng Xiu (2022). "Expected returns and large language models". In: *Available at SSRN 4416687.*
- Chen, Yiting, Tracy Xiao Liu, You Shan, and Songfa Zhong (2023). "The emergence of economic rationality of GPT". In: *Proceedings of the National Academy of Sciences* 120.51, e2316205120.
- Chinco, Alex, Adam D Clark-Joseph, and Mao Ye (2019). "Sparse signals in the cross-section of returns". In: *The Journal of Finance* 74.1, pp. 449–492.
- Christ, Margaret H, MJ Kim, and Michael A Yip (2024). "Survey Evidence on the Determinants and Consequences of Artificial Intelligence Use in Sell-Side Equity Research". In: *Available at SSRN 5071836.*
- Clark, Jack and Gillian K Hadfield (2019). "Regulatory markets for AI safety". In: *arXiv preprint arXiv:2001.00078.*

References X

- Cockburn, Iain M, Rebecca Henderson, Scott Stern, et al. (2018). *The impact of artificial intelligence on innovation.* Vol. 24449. National bureau of economic research Cambridge, MA, USA.
- Cong, Lin William, Tengyuan Liang, Xiao Zhang, and Wu Zhu (2024). *Textual factors: A scalable, interpretable, and data-driven approach to analyzing unstructured information.* Tech. rep. National Bureau of Economic Research.
- Cuéllar, Mariano-Florentino, Benjamin Larsen, Yong Suk Lee, and Michael Webb (2024). "Does information about AI regulation change manager evaluation of ethical concerns and intent to adopt AI?" In: *The Journal of Law, Economics, and Organization* 40.1, pp. 34–75.
- D'Acunto, Francesco and Alberto G Rossi (2023). "Robo-advice: Transforming households into rational economic agents". In: *Annual Review of Financial Economics* 15.1, pp. 543–563.

References XI

- DeMiguel, Victor, Javier Gil-Bazo, Francisco J Nogales, and André AP Santos (2023). "Machine learning and fund characteristics help to select mutual funds with positive alpha". In: *Journal of Financial Economics* 150.3, p. 103737.
- Dong, G Nathan (2024). "Can AI Replace Stock Analysts? Evidence from Deep Learning Financial Statements". In: *Evidence from Deep Learning Financial Statements (March 31, 2024)*.
- Dong, Xi, Yan Li, David E Rapach, and Guofu Zhou (2022). "Anomalies and the expected market return". In: *The Journal of Finance* 77.1, pp. 639–681.
- Dou, Winston Wei, Itay Goldstein, and Yan Ji (2024). "Ai-powered trading, algorithmic collusion, and price efficiency". In: *Jacobs Levy Equity Management Center for Quantitative Financial Research Paper*.
- D'Acunto, Francesco, Nagpurnanand Prabhala, and Alberto G Rossi (2019). "The promises and pitfalls of robo-advising". In: *The Review of Financial Studies* 32.5, pp. 1983–2020.

References XII

- Eisfeldt, Andrea L and Gregor Schubert (2025). "Generative AI and Finance". In: *Annual Review of Financial Economics* 17.
- Eisfeldt, Andrea L, Gregor Schubert, and Miao Ben Zhang (2023). *Generative AI and firm values*. Tech. rep. National Bureau of Economic Research.
- Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock (2024). "GPTs are GPTs: Labor market impact potential of LLMs". In: *Science* 384.6702, pp. 1306–1308.
- Fedyk, Anastassia, James Hodson, Natalya Khimich, and Tatiana Fedyk (2022). "Is artificial intelligence improving the audit process?" In: *Review of Accounting Studies* 27.3, pp. 938–985.
- Fedyk, Anastassia, Ali Kakhbod, Peiyao Li, and Ulrike Malmendier (2024). "Chatgpt and perception biases in investments: An experimental study". In: *Available at SSRN 4787249*.
- Feng, Guanhao, Stefano Giglio, and Dacheng Xiu (2020). "Taming the factor zoo: A test of new factors". In: *The Journal of Finance* 75.3, pp. 1327–1370.

References XIII

- Fetzer, Thiemo, Peter John Lambert, Bennet Feld, and Prashant Garg (2024). "AI-generated production networks: Measurement and applications to global trade". In.
- Fish, Sara, Yannai A Gonczarowski, and Ran I Shorrer (2024). "Algorithmic collusion by large language models". In: *arXiv preprint arXiv:2404.00806* 7.
- Freyberger, Joachim, Andreas Neuhierl, and Michael Weber (2020). "Dissecting characteristics nonparametrically". In: *The Review of Financial Studies* 33.5, pp. 2326–2377.
- Furman, Jason and Robert Seamans (2019). "AI and the Economy". In: *Innovation policy and the economy* 19.1, pp. 161–191.
- Fuster, Andreas, Paul Goldsmith-Pinkham, Tarun Ramadorai, and Ansgar Walther (2022). "Predictably unequal? The effects of machine learning on credit markets". In: *The Journal of Finance* 77.1, pp. 5–47.
- Garcia, Diego, Xiaowen Hu, and Maximilian Rohrer (2023). "The colour of finance words". In: *Journal of Financial Economics* 147.3, pp. 525–549.

- Glasserman, Paul and Caden Lin (2023). "Assessing look-ahead bias in stock return predictions generated by gpt sentiment analysis". In: *arXiv preprint arXiv:2309.17322.*
- Gofman, Michael and Zhao Jin (2024). "Artificial intelligence, education, and entrepreneurship". In: *The Journal of Finance* 79.1, pp. 631–667.
- Goldberg, Samuel G, Garrett A Johnson, and Scott K Shriver (2024). "Regulating privacy online: An economic evaluation of the GDPR". In: *American Economic Journal: Economic Policy* 16.1, pp. 325–358.
- Greig, Fiona, Tarun Ramadorai, Alberto G Rossi, Stephen P Utkus, and Ansgar Walther (2024). "Human Financial Advice in the Age of Automation". In: *Available at SSRN 4301514.*
- Grennan, Jillian and Roni Michaely (2020). "Artificial intelligence and high-skilled work: Evidence from analysts". In: *Swiss Finance Institute Research Paper* 20-84.
- Gu, Shihao, Bryan Kelly, and Dacheng Xiu (2020). "Empirical asset pricing via machine learning". In: *The Review of Financial Studies* 33.5, pp. 2223–2273.

References XV

Gu, Shihao, Bryan Kelly, and Dacheng Xiu (2021). "Autoencoder asset pricing models". In: *Journal of Econometrics* 222.1, pp. 429–450.

Guidolin, Massimo, Giulia Panzeri, and Manuela Pedio (2024). *Machine Learning in Portfolio Decisions*. Research Paper 233. Available at SSRN: https://ssrn.com/abstract=4988124 or http://dx.doi.org/10.2139/ssrn.4988124. BAFFI CAREFIN Centre.

Guo, Taicheng, Xiuying Chen, Yaqi Wang, Ruidi Chang, Shichao Pei, Nitesh V Chawla, Olaf Wiest, and Xiangliang Zhang (2024). "Large language model based multi-agents: A survey of progress and challenges". In: *arXiv preprint arXiv:2402.01680*.

Hampole, Menaka, Dimitris Papanikolaou, Lawrence DW Schmidt, and Bryan Seegmiller (2025). *Artificial intelligence and the labor market*. Tech. rep. National Bureau of Economic Research.

- Hansen, Anne Lundgaard, John J Horton, Sophia Kazinnik, Daniela Puzzello, and Ali Zarifhonarvar (2024). "Simulating the Survey of Professional Forecasters". In: *Available at SSRN.*
- Hansen, Anne Lundgaard and Sophia Kazinnik (2023). "Can chatgpt decipher fedspeak". In: *Available at SSRN.*
- Harvey, Campbell R, Yan Liu, and Heqing Zhu (2016). "… and the cross-section of expected returns". In: *The Review of Financial Studies* 29.1, pp. 5–68.
- Hoberg, Gerard, Craig Knoblock, Gordon Phillips, Jay Pujara, Zhiqiang Qiu, and Louiqa Raschid (2024). *Using representation learning and web text to identify competitor networks.* Tech. rep. Working Paper.
- Hoberg, Gerard and Asaf Manela (2025). *The Natural Language of Finance.* Working Paper. Available at SSRN: `https://ssrn.com/abstract=5119322` or `http://dx.doi.org/10.2139/ssrn.5119322`. USC Marshall School of Business Research Paper Sponsored by iORB.

References XVII

- Hoberg, Gerard and Gordon Phillips (2016). "Text-based network industries and endogenous product differentiation". In: *Journal of political economy* 124.5, pp. 1423–1465.
- Hobson, Jessen L, William J Mayew, and Mohan Venkatachalam (2012). "Analyzing speech to detect financial misreporting". In: *Journal of Accounting Research* 50.2, pp. 349–392.
- Horton, John J (2023). *Large language models as simulated economic agents: What can we learn from homo silicus?* Tech. rep. National Bureau of Economic Research.
- Hu, Allen and Song Ma (2024). *Persuading investors: A video-based study.* Tech. rep.
- Huang, Jing and Shumiao Ouyang (2025). "Soft Information, Hard Decisions: AI Advising". In: *Working Paper.*
- Ide, Enrique and Eduard Talamas (forthcoming). "Artificial intelligence in the knowledge economy". In: *Journal of Political Economy.*
- Jha, Manish, Jialin Qian, Michael Weber, and Baozhong Yang (2024a). *ChatGPT and corporate policies.* Tech. rep. National Bureau of Economic Research.

References XVIII

- Jha, Manish, Jialin Qian, Michael Weber, and Baozhong Yang (2024b). "Harnessing Generative AI for Economic Insights". In: *arXiv preprint arXiv:2410.03897*.
- Jiang, Wei, Junyoung Park, Rachel Jiqiu Xiao, and Shen Zhang (2025a). *AI and the Extended Workday: Productivity, Contracting Efficiency, and Distribution of Rents*. Tech. rep. National Bureau of Economic Research.
- Jiang, Wei, Yuehua Tang, Rachel J Xiao, and Vincent Yao (2025b). "Surviving the FinTech disruption". In: *Journal of Financial Economics* 171, p. 104071.
- Johnson, Garrett A, Scott K Shriver, and Samuel G Goldberg (2023). "Privacy and market concentration: intended and unintended consequences of the GDPR". In: *Management Science* 69.10, pp. 5695–5721.
- Johnson, Justin and D Daniel Sokol (2020). "Understanding AI collusion and compliance". In: *Cambridge Handbook of Compliance,(D. Daniel Sokol & Benjamin van Rooij, editors),(Forthcoming)*.
- Jurado, Kyle, Sydney C Ludvigson, and Serena Ng (2015). "Measuring uncertainty". In: *American Economic Review* 105.3, pp. 1177–1216.

References XIX

- Kanazawa, Kyogo, Daiji Kawaguchi, Hitoshi Shigeoka, and Yasutora Watanabe (2022). *AI, skill, and productivity: The case of taxi drivers.* Tech. rep. National Bureau of Economic Research.
- Kelly, Bryan T, Boris Kuznetsov, Semyon Malamud, and Teng Andrea Xu (2025). *Artificial Intelligence Asset Pricing Models.* Tech. rep. National Bureau of Economic Research.
- Kelly, Bryan T, Seth Pruitt, and Yinan Su (2019). "Characteristics are covariances: A unified model of risk and return". In: *Journal of Financial Economics* 134.3, pp. 501–524.
- Korinek, Anton (2023). "Generative AI for economic research: Use cases and implications for economists". In: *Journal of Economic Literature* 61.4, pp. 1281–1317.
- Kozak, Serhiy, Stefan Nagel, and Shrihari Santosh (2020). "Shrinking the cross-section". In: *Journal of Financial Economics* 135.2, pp. 271–292.

References XX

- Krockenberger, Vanessa, Anthony Saunders, Sascha Steffen, and Paulina Verhoff (2024). "Covenantai-new insights into covenant violations". In: *Available at SSRN 4640653.*
- Leippold, Markus, Qian Wang, and Wenyu Zhou (2022). "Machine learning in the Chinese stock market". In: *Journal of Financial Economics* 145.2, pp. 64–82.
- Lettau, Martin and Markus Pelger (2020). "Factors that fit the time series and cross-section of stock returns". In: *The Review of Financial Studies* 33.5, pp. 2274–2325.
- Li, Kai, Feng Mai, Rui Shen, Chelsea Yang, and Tengfei Zhang (2024). "Dissecting corporate culture using generative AI–Insights from analyst reports". In: *Available at SSRN 4558295.*
- Light, Nathaniel, Denys Maslov, and Oleg Rytchkov (2017). "Aggregation of information about the cross section of stock returns: A latent variable approach". In: *The Review of Financial Studies* 30.4, pp. 1339–1381.

References XXI

- Lopez-Lira, Alejandro and Yuehua Tang (2023). “Can chatgpt forecast stock price movements? return predictability and large language models”. In: *arXiv preprint arXiv:2304.07619*.
- Lu, Yao, Gordon M Phillips, and Jia Yang (2024). *The impact of cloud computing and ai on industry dynamics and concentration.* Tech. rep. National Bureau of Economic Research.
- Lv, Linying (2024). “The Value of Information from Sell-side Analysts”. In: *arXiv preprint arXiv:2411.13813*.
- Lyonnet, Victor and Léa H Stern (2022). “Venture capital (mis) allocation in the age of AI”. In: *Proceedings of the EUROFIDAI-ESSEC Paris December Finance Meeting*.
- Manning, Benjamin S., Kehang Zhu, and John J. Horton (2024). *Automated Social Science: Language Models as Scientist and Subjects.*
- Matos, Miguel Godinho de and Idris Adjerid (2022). “Consumer consent and firm targeting after GDPR: The case of a large telecom provider”. In: *Management Science* 68.5, pp. 3330–3378.

- Murray, Scott, Yusen Xia, and Houping Xiao (2024). "Charting by machines". In: *Journal of Financial Economics* 153, p. 103791.
- Novy-Marx, Robert and Mihail Z Velikov (2025). *AI-powered (finance) scholarship*. Tech. rep. National Bureau of Economic Research.
- Noy, Shakked and Whitney Zhang (2023). "Experimental evidence on the productivity effects of generative artificial intelligence". In: *Science* 381.6654, pp. 187–192.
- Ouyang, Shumiao (2021). "Cashless payment and financial inclusion". In: *Available at SSRN 3948925*.
- Ouyang, Shumiao, Hayong Yun, and Xingjian Zheng (2024). "AI as Decision-Maker: Ethics and Risk Preferences of LLMs". In: *Available at SSRN 4851711*.
- Peng, Sida, Eirini Kalliamvakou, Peter Cihon, and Mert Demirer (2023). "The impact of ai on developer productivity: Evidence from github copilot". In: *arXiv preprint arXiv:2302.06590*.

References XXIII

- Rock, Daniel (2019). "Engineering value: The returns to technological talent and investments in artificial intelligence". In: *Available at SSRN 3427412*.
- Rossi, Alberto G and Stephen Utkus (2024). "The diversification and welfare effects of robo-advising". In: *Journal of Financial Economics* 157, p. 103869.
- Schubert, Gregor (2025). "Organizational Technology Ladders: Remote Work and Generative AI Adoption". In: *Available at SSRN*.
- Seamans, Robert and Manav Raj (2018). *AI, labor, productivity and the need for firm-level data*. Tech. rep. National Bureau of Economic Research.
- Serafeim, George (2024). "Climate Solutions, Transition Risk, and Stock Returns". In.
- Shaffer, Matthew and Charles CY Wang (2024). "Scaling Core Earnings Measurement with Large Language Models". In: *Available at SSRN*.
- Sheng, Jinfei, Zheng Sun, Baozhong Yang, and Alan L Zhang (2024). "Generative AI and asset management". In: *Available at SSRN 4786575*.
- Tang, Huan (2019). "The value of privacy: Evidence from online borrowers". In: *Available at SSRN*.

References XXIV

- Webb, Michael (2019). "The impact of artificial intelligence on the labor market". In: *Available at SSRN 3482150*.
- Yoo, Minji (2024). *How Much Should We Trust Large Language Model-Based Measures for Accounting and Finance Research?* https://ssrn.com/abstract=4983334. The Wharton School Research Paper.
- Zheng, Xiang (2025). "How can innovation screening be improved? A machine learning analysis with economic consequences for firm performance". In: *Journal of Financial and Quantitative Analysis*, pp. 1–57.
- Zhou, Tianyu, Pinqiao Wang, Yilin Wu, and Hongyang Yang (2024). "FinRobot: AI Agent for Equity Research and Valuation with Large Language Models". In: *arXiv preprint arXiv:2411.08804*.