vba Journaal

CFA Society Netherlands

AI EN CHATGPT IN DE NEDERLANDSE BELEGGINGSSECTOR

JAARGANG 39 | NUMMER 155 | WINTER 2023

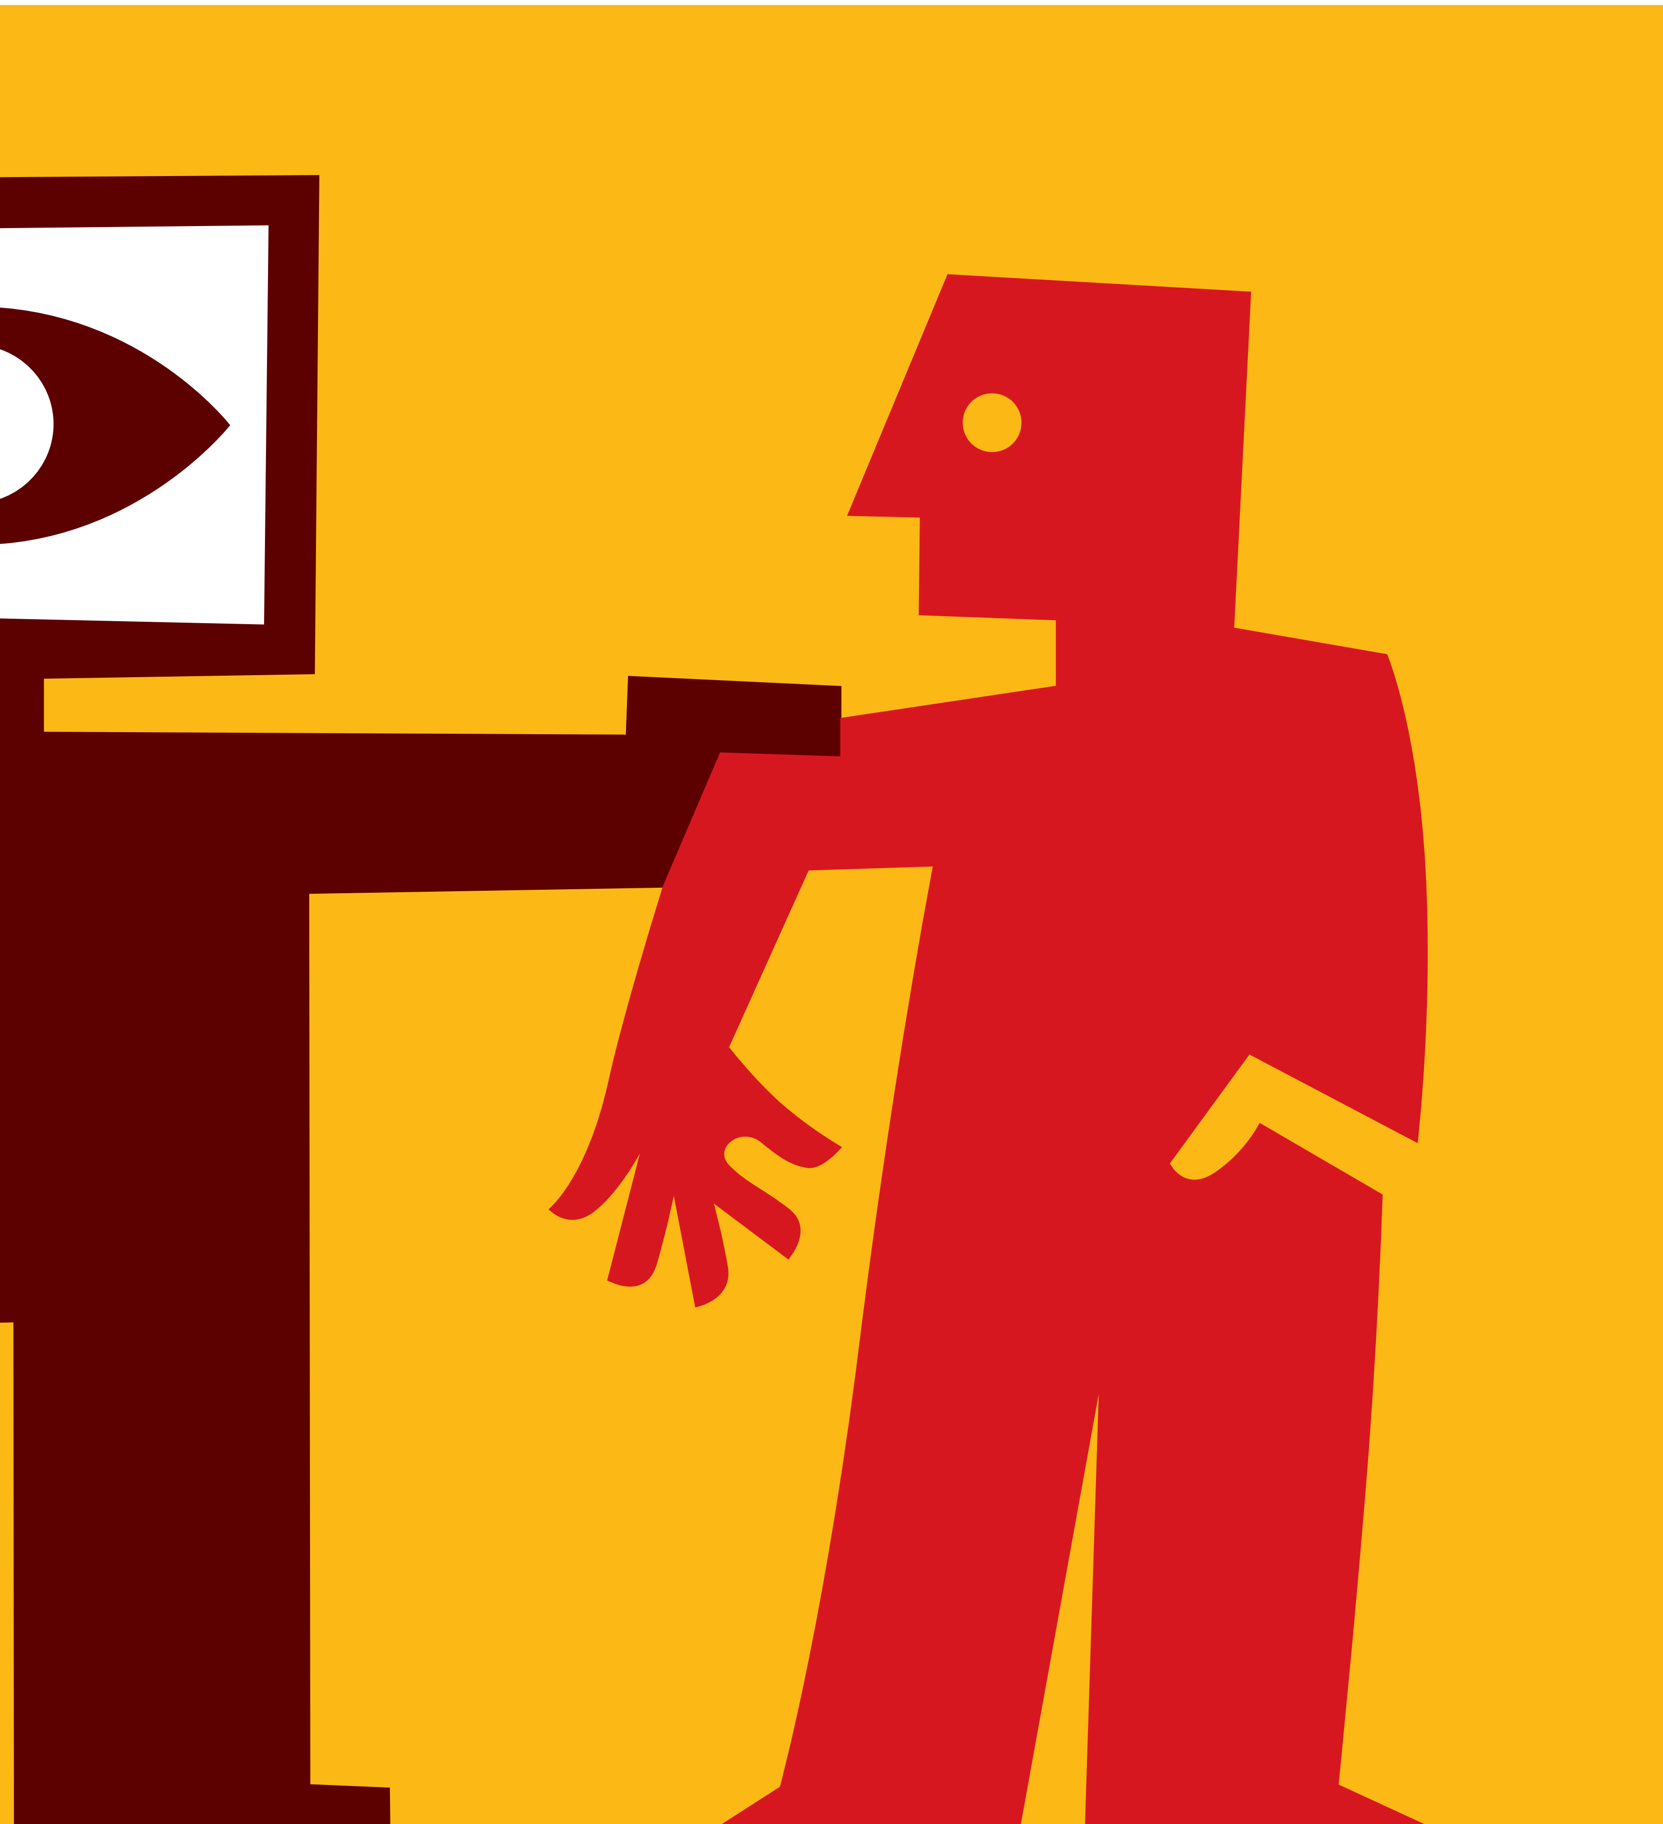

Inhoud

CALL FOR PAPERS

A new era for Fiduciary Management

Sinds de eerste uitbesteding in 2002 van het vermogensbeheer door de verzekeraar VGZ aan Goldman Sachs is fiduciair management inmiddels bij Nederlandse pensioenfondsen en verzekeraars de norm. Ook buiten Nederland wint fiduciair management aan populariteit. Het aantal aanbieders en het aantal varianten van fiduciaire modellen zijn wereldwijd gegroeid.

De redactie signaleert dat het fiduciaire model in beweging is. De Nederlandse sector heeft voldoende ervaring om de sterke en zwakke punten van dit model en haar implementatie te kunnen evalueren en evolueren. Zo lezen wij dat de toezichthouder een grotere verantwoordelijkheid ziet voor de pensioenfondsen om betrokken te zijn bij belangrijke uitvoeringsbeslissingen. Tevens zien wij dat de Nederlands implementatie van fiduciair management verschilt van de implementatie in Anglo-Saksische landen. Tenslotte zien wij zien divergentie in fiduciaire proposities.

De redactie nodigt auteurs uit om de ontwikkelingen binnen fiduciair management te ontrafelen. Om auteurs te inspireren heeft de redactie een aantal mogelijke onderwerpen bedacht. Om te beginnen is de redactie op zoek naar een auteur om van deze editie ook een naslagwerk te maken door middel van een overzichtsartikel over wat anno 2024 fiduciair management is in Nederland. Welke dienstverlening hoort hier bij en welke niet: bijvoorbeeld balansmanagement, managerselectie? Vervolgens zijn er de volgende onderwerpen waarvan de redactie denkt dat interessante artikelen over kunnen worden geschreven voor ons lezerspubliek.

- Welke ontwikkelingen heeft fiduciair management in Nederland de afgelopen 20 jaar doorgemaakt? Wat verklaart specifiek het Nederlandse fiduciaire model? En waarom wordt de implementatie veelal door Nederlandstaligen gedaan?
- Is een fiduciair een manager of een adviseur?
- Wat is de toekomst van het (traditionele) Nederlandse fiduciaire model?
- Hoe en in welke mate wijkt de selectie van een fiduciair manager af van de selectie van een vermogensbeheerder? In welke mate speelt het track record een rol?
- Verandert het perspectief van de toezichthouder op het fiduciaire model?
- Welke invloed heeft het MVB beleid op fiduciair management? Welke verschillen bestaan er tussen de Europa en de Verenigde Staten?
- Hoe verandert de rol en het takenpakket van de fiduciair manager bij verschillende bestuurs- en governancemodellen? Bestaat er een optimum? Heeft de oprichting van bestuursbureaus de fiduciaire taken overgenomen? Hoe beïnvloedt de nieuwe wet- en regelgeving, zoals de IORP II richtlijn vanuit de Europese Unie om het toezicht op pensioenuitvoerders te harmoniseren, het fiduciaire model?
- Wat is fiduciair management in Wet toekomst pensioenen (Wtp) en onder Pan-European Personal Pension Plan (PEPP)?
- Hoe verhoudt het Nederlandse fiduciaire model zich tot de Anglo-Saksische aanpak? Ofwel een internationale vergelijking van de implementatie van het fiduciaire model. We denken hierbij onder meer aan het zogenaamde outsourced CIO concept.
- Robo/AI advies als Fiduciair 2.0
- Wat is de toegevoegde waarde van fiduciair management? Kwalitatief? Kwantitatief?
- Wat zijn de verschillen en overeenkomsten, en waarom, in fiduciair modellen tussen pensioenfondsen en verzekeraars en banken?

Via deze call for papers roepen wij geïnteresseerde auteurs op om uiterlijk 15 februari 2024 contact op te nemen met de redactie (irma.willemsen@cfasociety.nl). Het moet daarbij gaan om oorspronkelijk werk dat nog niet eerder is gepubliceerd. Het artikel kan zowel in het Engels als Nederlands worden aangeboden.

Kunstmatige intelligentie en ChatGPT in de Nederlandse beleggingssector: recente ontwikkelingen en impact op de praktijk

De komst van ChatGPT heeft de wereld in slechts twaalf maanden veranderd. Dit taalmodel voert levensechte gesprekken, schrijft gedichten en programmeert. Kunstmatige intelligentie kan de economie een boost geven, stelden Vleugels en Van Wijnen onlangs in Het Financieele Dagblad van 20 november. Tegelijk berichtten Jongsma en Van der Leij op 7 december in dezelfde krant dat de econoom Anna Salomons waarschuwt voor ontwrichting van de arbeidsmarkt. In tegenstelling tot eerdere automatiseringen lijkt het nu vooral hoogopgeleide banen te raken.

De redactie heeft ook Kunstmatige intelligentie en grote taalmodellen zoals ChatGPT ingezet bij het maken van dit winternummer. Zo heeft deze technologie geholpen bij het schrijven van de call for papers en heeft het de uitwerking van het interview vergemakkelijkt. In alle gevallen keek de redactie nauwlettend mee en waar nodig grepen we in zodat we de kwaliteit ten alle tijden konden blijven garanderen. Deze technologie stroomlijnt niet alleen het maken van uw winternummer, maar ook verandert het het investeringslandschap in bredere zin.

In deze wintereditie van het VBA Journaal verkennen we de recente ontwikkelingen en de impact van deze technologieën op de Nederlandse investeringspraktijk. Dit nummer bevat een reeks bijdragen die zowel de mogelijkheden alsook de potentiële risico's en uitdagingen onderzoeken. Anna Dijkman schrijft al dat de komst van nieuwe technologieën zoals AI vaak zowel utopische verwachtingen als dystopische angstbeelden oproept.

Bij een aantal artikelen komen de utopische verwachtingen goed naar voren. Christian Hull bespreekt AI-toepassingen op investeringsbeslissingen, handel, risicobeheer en operaties. De bijdrage van Tjeerd van Capelle verkent hoe investeringsfirma's grote taalmodellen kunnen inzetten voor sentimentanalyse, ESG-classificatie, voorspellingen en thematisch onderzoek. Het artikel van Mike Chen, Iman Honarvar en Harald Lohre van Robeco geeft een overzicht van de adoptie van AI en belangrijke toepassingen zoals het voorspellen van rendementen, het clusteren van bedrijven en duurzaamheidsanalyse. Het artikel van Valentijn van Nieuwenhuijzen benadrukt hoe *natural language processing* inzichten kan bieden voor duurzaam beleggen. Het interview met Simona Paravani-Mellinghoff van BlackRock geeft een overzicht van de geschiedenis van AI en toepassingen zoals het analyseren van alternatieve data, het automatiseren van taken en het genereren van ideeën. Anna Dijkman is positief gestemd over de kansen die AI biedt, maar benadrukt dat de precieze impact nog onzeker is.

De dystopische angstbeelden komen wat duidelijker naar voren in de column van ChatGPT, aangestuurd door Loranne van Lieshout van de commissie risicomanagement. Die column waarschuwt voor potentiële nadelen van AI zoals marktvolatiliteit, flash crashes en ethische kwesties, maar waarbij de kanttekening geplaatst wordt door de commissie risicomanagement dat aan de juistheid soms ernstig getwijfeld kan worden. De column van Mark Geene onderzoekt hoe AI nieuwe vooroordelen gerelateerd aan gezichtsaantrekkelijkheid en vocale toon introduceert in investeringsbeslissingen. Het waarschuwt analisten om die vooroordelen te verminderen door onder meer het vermijden van persoonlijke ontmoetingen, te trainen om vocale aanwijzingen te detecteren, en niet te veel waarde hechten aan uiterlijk.

Maar er is nog meer! Marten Laudi laat zien dat financiële instellingen duurzame investeringen kunnen verhogen door informatie te verstrekken aan investeerders met duurzaamheidsvoorkeuren. Laudi wijst er tevens op dat prijsdiscriminatie de aantrekkelijkheid van dergelijke investeringen op lange termijn bedreigt.
Tjitsger Hulshoff beschrijft "Demographics Unravelled" door Amlan Roy dat laat zien dat demografie verder reikt dan bevolkingspiramides en geboortecijfers. Het beïnvloedt consumentengedrag, economische analyses en activaprijzen.

In dit nieuwe jaar 2024, staan we aan de vooravond van wat een boeiende periode belooft te worden. Lees verder voor een blik op de nieuwste ontwikkelingen op het gebied van AI in het Nederlandse vermogensbeheer!

Wij wensen u een goed 2024 en veel leesplezier toe!

Roy Hoevenaars
Tjitsger Hulshoff
Sander Nooij

Bridging the Gap: Dutch Institutional Investors & Dutch Venture Capital

By Berg De Bleecker and Som Toohey, Private Equity & Venture Capital Committee

The Netherlands has emerged as a dynamic hub for startups, becoming renowned for innovative companies and entrepreneurship. However, a closer examination reveals a challenge in scaling these startups into larger enterprises, an aspect where the nation lags behind the European average according to a 2022 study by McKinsey & Company. This perspective leads to the question: Can the availability of local Venture Capital (VC) funding and its success help scaling young companies in the region?

To engage with this topic deeper, the Private Equity & Venture Capital Committee of CFA Society The Netherlands decided to arrange an event in November titled *"Bridging the Gap between Dutch Institutional Investors and Dutch Venture Capital"* at the Symphony Offices in Amsterdam to bring together key stakeholders to discuss existing challenges and explore potential solutions.

Setting the Scene

After the event Chair, Gabriele Todesca from the European Investment Fund (EIF), opened up proceedings, Felix Zwart of Nederlandse Vereniging voor Participatiemaatschappijen (NVP) presented data revealing that only 6% of the EUR 12.4bn raised by Dutch VC funds over 2007-2022 came from pension funds, of which most were international players outside of the Netherlands. Other European countries, especially in the Nordic region, scored higher on this metric. Encouragingly, 2023 data indicates a positive shift, suggesting increased involvement of Dutch pension funds in VC.

Public Investors' Perspectives

Roger Havenith, Deputy Chief Executive at EIF, highlighted in his speech risks in Europe's lag in growing VC ventures beyond a certain size, emphasizing the need for public investors to attract private capital and foster a sustained track record of returns. The call to action resonated: all investors were urged to incorporate and generate positive impact with their activities.

In their interview by the event chair, Rinke Zonneveld (CEO of Invest-NL) and Wendy de Jong (Chair of ROM-Nederland and CEO of OostNL) provided valuable insights, emphasizing Invest NL and Regional Development Agencies' (RDAs) roles in catalysing investments into early-stage and growing innovative companies. Their collaborative efforts include building the entire VC ecosystem and catalysing diverse fundraising options for startups, to complement the investments of RDAs and Invest NL.

All the three top public investors shared details on their activities. On the day before the event and very much fitting in with the theme, EIF and Invest-NL announced Dutch Future Fund II, their new joint programme for emerging VC funds with a focus on The Netherlands. Meanwhile, ROM-Nederland's activities were described, including working closely with Invest-NL, as the collaboration of all nine RDAs. RDAs combined investments currently amount to c. EUR 2.3bn and two RDAs rank among the top-three venture capital investors in The Netherlands according to Dealroom.

Constantijn van Orange and Cabriele Todesca

Rinke Zonneveld, Gabriele Todesca and Wendy de Jong

Interview with Constantijn van Oranje (Special Envoy Techleap.nl)

The important role of venture capital in addressing societal challenges and sustainable transitions was emphasized during an engaging interview with Constantijn van Oranje. To capitalise on the Netherlands' innovation strengths a balanced approach between public and private funds was needed to help build innovations these into scalable businesses – emphasising the need to reduce friction in startup fundraising. Successfully encouraging the Dutch pension fund industry to commit more funds to venture capital might move the needle and help drive The Netherlands towards the top of the rankings in terms of scaling start-ups into the large companies of the future.

Pension fund VC investor perspective

Signs are emerging that VC is gaining traction on the agendas of pension fund managers, aligning with a shared interest in making a positive impact. Victor van Will, Principal Product Strategy & Investment Solutions at MN Services, shared insights into their progressive approach to including venture capital in their portfolio. This involved revising policies throughout the investment process to allow for the characteristics of VC funds. MN had worked hard on embedding VC into the portfolio and is sharing lessons learned with other investors. This and generating a track record of the approach in terms or returns and exits would stimulate more pension funds to participate in the sector.

Panel Discussion Highlights

The panel discussion, led by Juul Vaandrager, Director Venture Capital at NVP, featured four prominent Dutch VC fund managers from SET Ventures, Rubio Impact Ventures, Gilde Healthcare and Deep Tech XL. They made a strong case for investments in VC funds, reminding the audience of the wealth of quality funds available in The Netherlands and growing track records. The diverse panel touched upon challenges faced by first-time fund managers, conversations around fees, and valuable operational due diligence. The audience was also treated to several examples of successful engagement of the funds growing their portfolio companies.

Wrap-up

Judging by the event's success, as indicated by the high turnout, positive atmosphere, audience interaction, and subsequent social media coverage, it is evident that the conversation surrounding bridging the gap between Dutch institutional investors and Dutch venture capital is gaining momentum. The hope is that future discussions will focus more on flourishing partnerships in the space, reflecting a narrowing of the gaps between stakeholders. The Private Equity & Venture Capital Committee remains optimistic about the collaborative efforts shaping the future of the Dutch VC landscape.

AGENDA

16 januari 2024
ALV en nieuwjaarsreceptie

26 januari 2024
Alternatives binnen het pensioenstelsel
PL 2 CREDITS

9 februari 2024
The Sustainability Outlook
PL 2 CREDITS

12 februari 2024
WIM workshop Negotiations Skills
PL 1 CREDIT

1 maart 2024
Finals CFA Research Challenge Benelux

LET'S MEET!

CFA Society Netherlands
Diversity, Equity and Inclusion (DEI) Committee

Did you know that "75% of organizations with frontline decision-making teams reflecting a diverse and inclusive culture will exceed their financial targets"?[1]

We are thrilled to announce that the CFA Society Netherlands Women in Investment Management Committee (formerly "WIM" Committee) is going to expand its reach and will become the Diversity, Equity and Inclusion (DEI) Committee starting 1 January 2024. In line with the values and objectives of CFA Society Netherlands and CFA Institute,[2] the new Committee will encourage inclusion and diversity at every level in the investment industry by including different groups of professionals in the Netherlands. Jovita Razauskaite, CFA is appointed as the Chairperson of this Committee and will be supported by the core team members Agnese Pavlovska, CFA, Marilu Ortega, CFA, Cody Martian, and Rasha Zakkak, CFA.

The key focus of the Committee will be to firstly establish a platform for everyone where different groups could be proactively invited to participate in, engage and learn from each other and external experts. Next to the events that are content-oriented, the DEI Committee will put efforts to create a platform for the soft skills development, such as, but not limited to, personal branding, negotiation and persuasion skills, leadership, communication, power of culture and others. Such skills are the basis of the DEI factors leading for an individual's success within and outside the workplace as well as the performance of the organisation itself.

Secondly, we will work on establishing a network to foster the retention and career development of our community members in the financial services industry. We have noted that DEI leadership is moving from human resources (HR) to business leaders[3] and thus, through our networking events, mentorship programs and workshops we hope to equip the industry participants with the right skills and opportunities to make a positive change and encourage DEI culture in investment management. We plan to do so also in close collaboration with other CFA Society NL Committees for a greater positive ripple effect.

Finally, our DEI platform would promote diversity and inclusion in the investment industry at universities and via partnerships with groups of similar affinities in order to bridge the gap between the academic world of diversity and inclusion and what we see in our industry.[4] After all, current students are poised to be our future leaders.

Jovita Razauskaite

Marilu Ortega

Rasha Zakkak

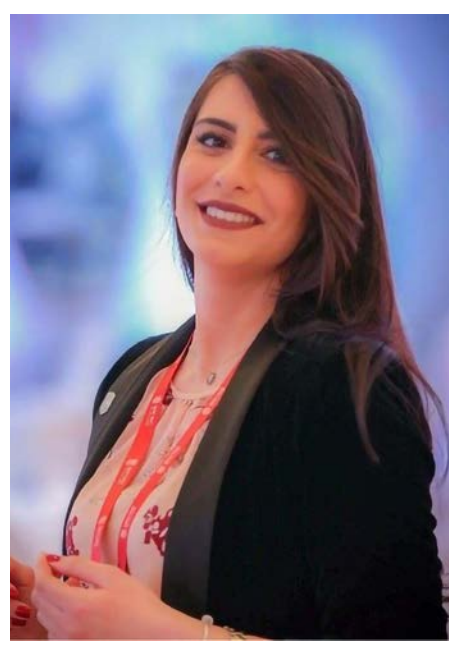

Rebranding the WIM Committee to DEI Committee is driven by the changing market environment and growing importance of different diversity groups as well as their added value to the profession and the society at large. According to Garner's research, "75% of organizations with frontline decision-making teams reflecting a diverse and inclusive culture will exceed their financial targets. And gender-diverse and inclusive teams outperformed gender-homogeneous, less inclusive teams by 50%, on average."[5] Clearly, a broader aspect of DEI positively contributes to the businesses.[6]

Additionally, DEI is on the agenda for many companies. 76% of European organisations have publicly declared their commitment to DEI.[7] However, according to the same PwC research, only a small percentage of DEI programs reach the highest level of maturity.

Therefore, our DEI Committee aims to actively contribute to the positive changes in this field through our activities and community engagement.

So we are entering 2024 with a lot of excitement! Next year we are planning to diversify our own Committee further and organize a number of insightful events. So keep a close eye on our activities via the CFA Society Netherlands Newsletter and Linkedin page and we look forward to meeting you soon!

Did this article trigger your interest about DEI? Or do you already have some ideas on how we could reach our goals? We are happy to hear from you!

Jovita Razauskaite, CFA (Chairperson). Committee members: Agnese Pavlovska, CFA, Marilu Ortega, CFA, Cody Martian, and Rasha Zakkak, CFA.

Notes

1 https://www.gartner.com/smarterwithgartner/diversity-and-inclusion-build-high-performance-teams
2 https://rpc.cfainstitute.org/-/media/documents/code/dei/dei-code-overview.pdf
3 https://rpc.cfainstitute.org/-/media/documents/article/industry-research/accelerating-change.pdf
4 https://rpc.cfainstitute.org/en/topics/diversity-equity-and-inclusion/types
5 https://www.gartner.com/smarterwithgartner/diversity-and-inclusion-build-high-performance-teams
6 https://www.mckinsey.com/capabilities/people-and-organizational-performance/our-insights/delivering-through-diversity
7 https://www.pwc.com/gx/en/services/people-organisation/global-diversity-and-inclusion-survey/european-report.pdf

Cody Martian

Agnese Pavlovska

Impact Investing in Publieke Aandelen Markten

CFA SOCIETY NETHERLANDS EVENT, 16 NOVEMBER 2023, ROSARIUM, AMSTERDAM

Door Hilde Veelaerts en Ridzert van der Zee, Commissie Verantwoord Beleggen

Inleiding

Impact beleggen kende over de afgelopen jaren een sterke groei. Meer en meer groeit de overtuiging dat de financiële sector naast de overheid en publieke sector een belangrijke rol heeft in het verwezenlijken van de transitie naar een meer duurzame wereld. De term "impactbeleggen" werd in 2007 bedacht door The Rockefeller Foundation. Zij wilden hiermee een naam geven aan investeringen die gedaan worden met de bedoeling om zowel financieel rendement als sociale en/of ecologische impact te genereren. Inmiddels bedraagt de AuM van de markt van impact beleggingen wereldwijd volgens het Global Impact Investing Netwerk (GIIN) al 1,164 biljoen USD en volgens de International Finance Corporation (IFC), die een bredere definitie hanteert, al 2,3 biljoen USD. Volgens beide organisaties is impact beleggen een strategie die in verschillende activaklassen kan worden toegepast. Recent markt onderzoek van de GIIN laat zien dat impact beleggen in publieke aandelen de snelst groeiende activaklasse is.

Als duurzaamheidscommissie wilden wij graag een event organiseren waarbij we meer inzicht wilden geven over hoe de markt voor impactbeleggingen in Nederland is samengesteld en belangrijke organisaties in de context van impact-beleggen, als de GIIN en NAB, voorstellen. Verder wilden we de vragen beantwoorden als 'hoe beleg je met impact in publieke aandelenmarkten en welke overwegingen kan je maken over activaklassen heen?'

De sprekers op het event waren: Laure Wessemius-Chibrac (NAB), Wouter Koelewijn (GIIN), Neal Hegeman (ASR), Jos Gijsbers (ASR), Hilde Veelaert (Cardano). Het event werd gemodereerd door Ridzert van der Zee (Triple A – Risk Finance).

Introductie GIIN en NAB

Global Impact Investing Network: De GIIN, opgericht in 2009, is een wereldwijd netwerk, dat impact beleggen wil doen groeien en barrieres verlagen zodat meer investeerders kapitaal kunnen alloceren aan oplossingen voor 's werelds grootste uitdagingen.
De GIIN bereikt dit doel door ondermeer: educatie, onderzoek, kennisuitwisseling en het bouwen van instrumenten en middelen die helpen om de ontwikkeling van de sector te versnellen. De GIIN heeft wereldwijd meer dan 480 organisaties als lid, in meer dan 60 landen. De GIIN lanceerde ook IRIS+, een systeem voor het meten en managen van impact dat vrij beschikbaar[1] is.

Netherlands Advisory Board on Impact investing: De NAB is een onafhankelijke non-profitorganisatie die tot doel heeft de groei van de markt van impact beleggen te versnellen en de effectiviteit te verbeteren. Dit doen ze door het bewustzijn en het vertrouwen in impactvolle investeringen te vergroten en belemmering voor impact beleggen aan te pakken. De NAB is onderdeel van een wereldwijd netwerk van nationale advies raden die gegroepeerd zijn onder het Global Steering Group for impact investing (GSG).

1 Bronnen: https://iris.thegiin.org

Definitie van impact beleggen

Volgens de definitie van de GIIN zijn impactbeleggingen, "beleggingen die gedaan worden met de intentie om positieve, meetbare sociale en ecologische impact te genereren naast financieel rendement". Impact beleggingen situeren zich op het "Spectrum of Capital" naast traditionele beleggingen (finance-only), verantwoord, duurzaam en filantropie (impact-only). Afhankelijk van hoe de afweging wordt gemaakt tussen impact, risico en financieel rendement kan een impact strategie een marktconform rendement nastreven of een lager rendement accepteren ten voordele van hogere positieve impact uitkomsten.

Impactbeleggingen in Nederland

> *"Allocaties naar impactbeleggingen moeten groeien naar minstens 10%"*

De NAB heeft een onderzoek gedaan naar de status van de Nederlandse Impact Investing Sector (NAB: Are we punching below our weight – March 2022).[1] Daarbij inventariseert de NAB de omvang van de Nederlandse impact markt en doet ze uitvraag naar eventuele belemmeringen die de groei van de sector tegenhouden. Volgens de meting zou de sector in Nederland 150 tot 180 miljard EUR AuM in impact beleggingen hebben, wat overeenkomst met 4 tot 6% van de totale assets beheerd in Nederland. Hiermee maken impact beleggingen in Nederland al een grotere proportie van de assets uit dan wereldwijd waarvoor het percentage 1 tot 2% bedraagt. Wereldwijd zou volgens de NAB een allocatie van 5 tot 7% nodig zijn om de SDG's in 2030 te halen. Qua type belegger zien we de grootste allocaties naar impact beleggingen terug bij pensioenfondsen gevolgd door fonds- en assetmanagers. De gekozen activaklasse verschilt per type deelnemer aan het onderzoek. 75% van de institutionele beleggers melden dat ze in obligaties of andere vastrentende waarden beleggen, terwijl 60% van de niet-institutionele beleggers eerder kiezen voor private equity beleggingen.

Qua SDG's geeft het onderzoek aan dat sommige SDG's meer nagestreefd worden dan andere. SDG 13 (Climate action), SDG 8 (Decent work and economic growth) en SDG 7 (Affordable and clean energy) worden het meest frequent nagestreefd. Belangrijkste belemmeringen die vastgesteld worden zijn gerelateerd aan Nederlandse en Europese wetgeving. Volgens beleggers worden bijvoorbeeld hogere kapitaalvereisten geëist voor illiquide beleggingen of beleggingen in opkomende landen. Een andere belangrijke belemmering is het gebrek aan standaardisatie, definities en data. Ook vereist impact beleggen een langere beleggingshorizon en hogere tolerantie voor minder liquide beleggingen. De NAB draagt bij aan het verlagen van deze belemmeringen door bijvoorbeeld het reageren op consultaties of het publiceren van research papers (het meeste recent over perceptie van risico en kapitaalvereisten onder Solvency II[2]).

Figuur 1
Spectrum of capital

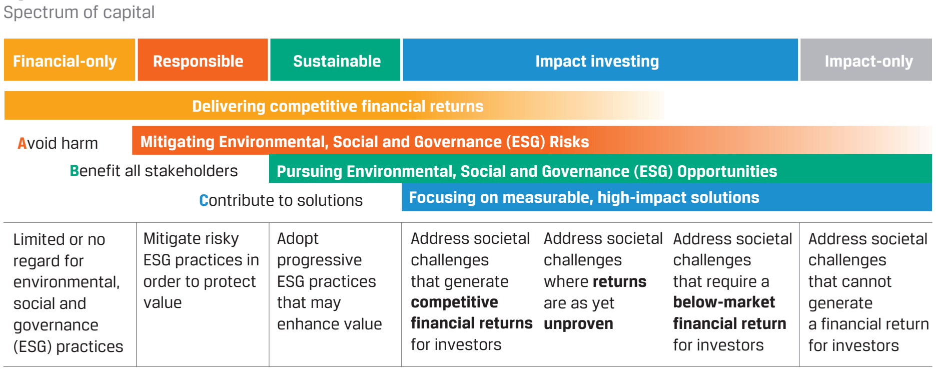

Bron: Bridges Capital, IMP

Figuur 2
Impact allocation

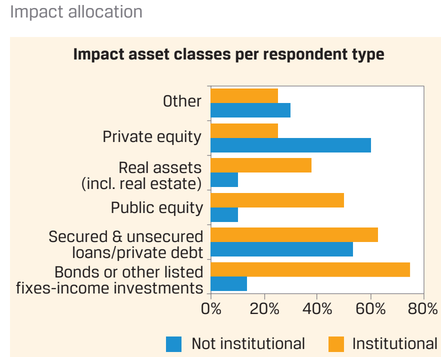

Bron: NAB Report – Are we punching below our weight?

Impactbeleggen in Publieke Aandelen Markten

> *"Having impact is more than having exposure."*

Inspelend op de groei van impact beleggen in publieke aandelen lanceerde de GIIN in 2023 de "Guidance for Pursuing Impact in Listed Equities".[3] Volgens de GIIN heeft een impact belegging vier belangrijke kenmerken, die ook gelden voor publieke aandelenmarkten:

Intentionaliteit:
Een impact belegger wenst intentioneel bij te dragen aan meetbare sociale of milieu doelen.

Contributie:
Als impact belegger is het belangrijk om duidelijk en inzichtelijk te maken hoe je als belegger bijdraagt aan het maken van positieve impact. Voor publieke aandelen is het belangrijkste middel om de positieve impact van beleggingen te vergroten: voting en engagement.

Gebruik van impact data of bewijsmateriaal in de opzet van het beleggingsproces:
De integratie van impact data of impact-overwegingen in het beleggingsproces zorgt ervoor dat de uiteindelijke beleggingskeuzes bijdragen aan de sociale of ecologische doelen die de strategie nastreeft.

Management van impact:
Een impactbelegger monitort en rapporteert de impactresultaten van beleggingen met de intentie om deze impact resultaten verder te verbeteren.

De twee belangrijkste concepten voor een impactfonds of -strategie in publieke aandelenmarkten zijn de Theory of Change en de contributie van de belegger zelf.

CONCEPT #1: THEORY OF CHANGE

Een belangrijk concept voor impact beleggers in publieke aandelen is de *Theory of Change*. Een Theory of Change maakt duidelijk welk probleem de belegger wil helpen oplossen en beschrijft vervolgens op welke wijze de belegger zal bijdragen aan verder vergroten van de positieve impact die de belegging genereert. Een Theory of Change is meestal opgebouwd uit vijf stappen.

Figuur 3
Theory of Change: Climate example

Linking activities with desired impact – additionality of the investee

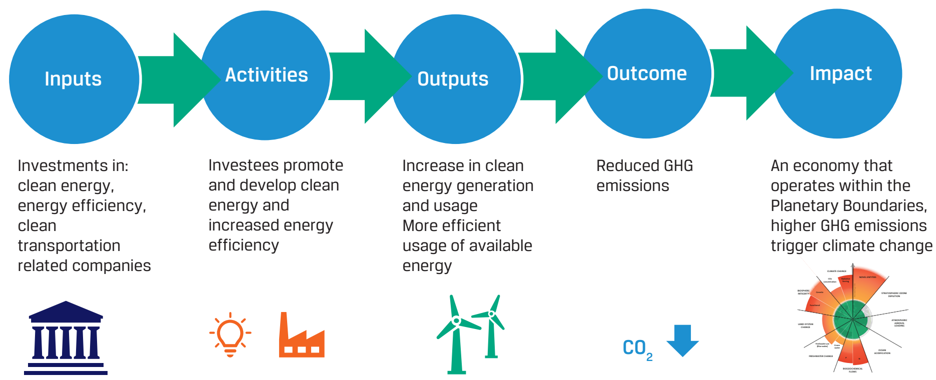

Bron: Cardano

De inputs zijn de middelen die gebruikt worden om de activiteiten van de gekozen belegging te bereiken. De activiteiten leiden tot een bepaald resultaat of "output", dit zijn vaak producten of diensten. De outputs leiden tot een bepaalde uitkomst of "outcome". De "outcome" kan een bepaald niveau van welzijn zijn of een status van het milieu. De "impact" is de verbetering in het niveau van welzijn of de status van het milieu.

CONCEPT #2: CONTRIBUTIE VAN DE BELEGGER

Contributie van de belegger is een belangrijk concept voor impactbeleggen in publieke aandelenmarkten. Met alleen het aankopen van aandelen op (secundaire) aandelenmarkten creëert een belegger nog geen impact in de reële economie. Het beleggen in beursgenoteerde aandelen van bedrijven die exposure hebben op of "aligned" zijn met impactdoelen is dus niet genoeg om te kunnen spreken van een impactbelegging. Aantoonbare contributie van de manager is een belangrijke vereiste om te kunnen spreken van 'impact' vanuit publieke aandelen. Het is met name de contributie van de belegger zelf die van een beleggingsstrategie in publieke markten een impact strategie maakt. De belangrijkste middelen om deze contributie te leveren voor een belegger in publieke aandelen zijn stemmen en engagement. Traditioneel kan engagement voor duurzame beleggingsfondsen een breed scala aan ESG-onderwerpen bestrijken. Voor impact fondsen zijn de onderwerpen voor engagement direct gekoppeld aan de impactdoelstellingen die de beleggingsstrategie nastreeft.

Het streven naar de gestelde impactdoelen moet voor een impactstrategie geïntegreerd zijn door het hele beleggingsproces heen: van het stellen van de doelen en Theory of Change tot aan de portfolio constructie, de selectie van aandelen, voting, engagement, rapportage en monitoring van impact data.

Investment process

> *"Impact beleggen is een ambacht"*

Selectie
Cardano geeft een concreet voorbeeld van hoe de selectie van aandelen voor hun impact strategie wordt uitgevoerd. Het samenstellen van het impact universum van aandelen die in aanmerking komen voor de impact portefeuille vindt plaats in twee stappen. De eerste stap is kwantitatief, op basis van data en positieve screening wordt het universum van 20.000 beursgenoteerde aandelen gereduceerd tot 300 aandelen. In het screeningsproces worden aandelen geselecteerd die bijdragen aan de impact doelen van de strategie, maar wordt er ook systematisch gescreened op het vermijden van eventuele negatieve impact. In de tweede kwalitatieve stap beoordeeld Cardano de impact die de bedrijven maken volgens de "Five Dimensions of Impact". Deze vijf dimensies gebruikt Cardano om de omvang en additionaliteit van de impact van de betrokken bedrijven beter te begrijpen en te beoordelen.

Figuur 4
Qualitative Review

Five Dimensions of Impact:
- What outcomes is the company contributing to and how important are these to stakeholders?
- Who are the impacted stakeholders and how underserved were they prior to the company's effect?
- How much is the number of stakeholders that experienced the outcome, what degree of change did they experience and for how long?
- Contribution assessment of whether an enterprise's and/or investor's efforts resulted in outcomes that were likely better than what would have occurred otherwise
- Risk assessment of the likelihood that the impact will be different from what was expected

Intentionality
- What is the priority level assigned by the company to create impact?

Negative check
- Are there any adverse impacts that are or might be generated by the company?

Bron: Cardano

Figuur 5
Transition risk in government bonds

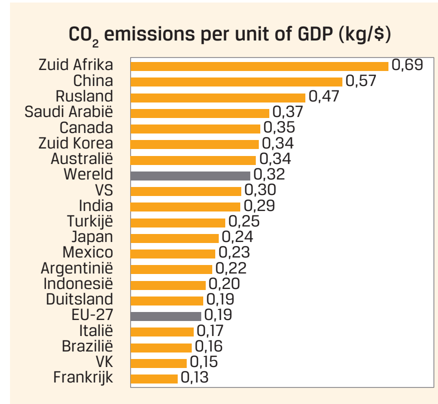

Bron: Our World in Data

Meten en rapporteren
Het meten en rapporteren van positieve en negatieve sociale en milieu impact van beleggingen is een belangrijk kenmerk van impact beleggen. Het helpt bij het nemen van beslissingen maar verhoogt ook de transparantie naar eindbeleggers. De beschikbaarheid van data is hierbij een uitdaging. Daarbij is het ook belangrijk dat de data die geselecteerd wordt gerelateerd is aan de "outcomes" die je wilt bereiken met de impact belegging. Cardano geeft een concreet voorbeeld van een bedrijf dat positieve impact maakt in de tuincultuur door het aanreiken van oplossingen voor minder gebruik van water. De concrete "outcome" KPI's die gemonitored worden zijn gerelateerd aan water efficiency en werden vastgesteld tijdens de kwalitatieve analyse in het selectie proces. In de engagement gesprekken wordt het bedrijf aangemoedigd om breder doelen te stellen, niet alleen op de water "footprint" maar ook de water "handprint", alsook uit te breiden naar gebieden met meer water schaarste. Het monitoren van vermindering van waterstress in de betrokken regio's of de uitbreiding naar nieuwe regio's zijn in dit geval voorbeelden van "impact" KPI's.

OVERWEGINGEN BIJ HET ALLOCEREN NAAR IMPACT OVER ACTIVAKLASSEN HEEN

ASR alloceert in haar portefeuille naar impact in verschillende activaklassen: van publieke aandelen, overheids- en bedrijfsobligaties, private leningen en aandelen tot hypotheken en vastgoed. Elke activaklasse heeft haar eigen criteria, definities, processen en governance.

Impact is voor een duurzame belegger een extra dimensie om rekening mee te houden, naast bestaande dimensies als rendement en risico, maar ook bijvoorbeeld fiscaliteit of wettelijke kapitaalvereisten. Als belegger komt het er steeds op aan om de optimale portefeuille te maken rekening houdend met de verschillende dimensies en restricties, waar impact doelstellingen onderdeel van zijn.

Bij het samenstellen van de portefeuille onderzoekt ASR bijvoorbeeld wat het effect is van klimaatverandering op het verwachte rendement en risico. Hiervoor gebruikt ASR een aantal klimaat scenario's. Een belangrijke conclusie hierbij is dat alle scenario's in meer of mindere mate een negatief effect hebben op het verwachte rendement en de solvabiliteitsratio, waarmee ook kansen voor impact investeringen worden geïdentificeerd. Een concreet voorbeeld van hoe klimaatverandering het verwachte rendement op overheidsobligaties kan beïnvloeden is gerelateerd aan de ratio CO2 emissies/GDP. Als er een wereldwijde CO2 taks wordt ingevoerd dan zullen landen die relatief veel CO2 uitstoten per GDP relatief harder geraakt worden dan landen die minder uitstoten.

Conclusie

Op basis van een peiling onder de aanwezigen bij het event, was de meerderheid er van overtuigd dat het nastreven van impact heel goed mogelijk is via beleggen in publieke markten.

Het classificeren als een impact belegging in publieke markten vergt wel een bepaalde discipline van de belegger: de intentie om positieve impact te maken, de actieve contributie van de belegger zelf (bijvoorbeeld door engagement) en het meten en rapporteren van daadwerkelijke impact. Per activaklasse zal de uitwerking van impact beleggen verschillen, afhankelijk van beschikbare tools en data (voor de contributie per manager of het screenen van bedrijven), karakteristieken (zoals beleggingshorizon en risico/rendement), beleggingscriteria (waaronder de keuze van beleggingsinstrumenten en wijze van beleggen). Impact kan dus geïntegreerd worden in de samenstelling van een beleggingsportefeuille over activa-klassen heen. Impact is een extra dimensie die in de gehele beleggingscyclus kan worden geïntegreerd.

Noten
1 NAB Team Site – NAB report – Are We Punching Below Our Weight.pdf – All Documents (sharepoint.com)
2 https://www.nabimpactinvesting.nl/post/capital-requirements-for-institutional-investors
3 https://thegiin.org/research/publication/listed-equities-working-group

Buluitreiking

Op 24 november jl. studeerde de 20ste lichting cursisten af aan de Postgraduate VU-VBA opleiding Investment Management. CFA Society Netherlands en de Vrije Universiteit feliciteren de volgende personen van harte met het succesvol afronden van de opleiding.

Jurgen Muijs van de Moer en Kevin Sallé staan niet op de foto.

Simon Zijlstra, Bas Neelis, Sander Rijpma en Jakar van Blaaderen staan wel op de foto maar horen niet bij de laatste afgestudeerden. Zij hadden hun diploma al.

Nieuwe RBA'ers

Naam	Bedrijf
Toon Admiraal	State Street Bank
Kees de Koning	VI Company
Jeremy van der Laan	AF Advisors
Jurgen Muijs van de Moer	Blue Sky Group
Kevin Sallé	PGB Pensioendiensten
Christiaan Wevers	PGGM
Cheng Lian Zhou	Rabobank

Finding our feet in a world of AI

Christian Hull

If you were in AI in Financial Services in 2023 it was a wild ride. Whereas AI was a tough sell in previous years; executives didn't see the value, and readily cited explainability, or untested outcomes as a reason not to pursue the technology. Machine learning was quietly being used by quants and hedge funds, but mainstream financial services were asleep at the wheel. Then, OpenAI launched ChatGPT, and the world literally changed overnight, suddenly budgets were open, and expectations were sky high!

Even now at the end of 2023 many executives are unclear on what Artificial Intelligence (AI) is all about, so in this article I will try to pull back the veil and explain what the value is, which types of AI are delivering value, and how to build an AI capability at your firm.

WHAT IS AI?

If only to dispel misbeliefs, it is always best to start with some basic concepts. So, what is AI? At its simplest AI is a black box with data as an input and probability as an output. When you dig a little deeper into the world of data science you soon discover new terms like machine learning, deep learning, supervised learning, classification, regression, etc. It can be daunting, but if you are not building actual models yourself it is safe to lump all the complexity into a single category called AI. It is, however, important to know the different types of outcomes an AI could derive from your data. The Turing Institute published a paper[1] that examines the techniques and use cases in greater technical detail that I would recommend if you wish to dive deeper into the subject.

- Predictive Analytics – Simply put this is the ability to predict outcomes, think of portfolio valuation trends, settled trade positions, cash balances, asset prices, etc. Doing this manually is often near impossible, but AI models can happily process a wide range of factors needed to calculate a future outcome. This is hugely important in financial services, and I'd argue the lion's share of AI models built so far are predictive.
- Anomaly Detection – "One of these things is not like the other". This is useful for fraud detection, anti-money laundering, or simple data quality checking.
- Simulations – AIs can be built to run complex scenario models like the famous Monte Carlo simulation. Futurists also know that this is one of the main abilities of quantum computing, and so we'll see more of this emerging in the coming years.
- Automation – Most people think this is the primary use case for AI, but I'd argue that most of the automation use cases have already been executed using existing technologies, and now we have a long tail of small tasks that don't justify the expense of AI. However, if you have a process that is expensive and was not possible to automate with existing technology, AI could be the innovation you have been waiting for. Just ensure the business case is sound.
- Large Language Models – This is the category that contains ChatGPT, and other foundation models. The novelty of these generative AI models is that they are a new technology that allows complex language-based tasks that were previously impossible to achieve economically using natural language processing.

Christian Hull
Ex Head of Innovation, EMEA, BNY Mellon

WHAT WILL AI DO FOR ME?

A recent survey[2] conducted by Oliver Wyman on UK Finance members showed that 91% of firms had deployed Predictive AI versus only 22% who had deployed any generative AI solutions in any way. This is typical of the trend seen across all geographies. Generative AI models, typically large language models that generate new content based upon the data they have been trained on, are attracting all of the hype, but the real work is being done by more traditional narrow-focus AI models. In my experience, predictive analytic models are immediately valuable to financial services firms; many others can add real value but are harder to embed into the current state process. The simple ability to predict an outcome before it happens can easily be added to existing process flows to enhance the value of the data.

It has been amazing to see how Generative AI has galvanized the executive suite into action. I see clear evidence of this in the amount of hiring that firms have been doing. The Evident AI Index[3] noted that there has been a 10% increase in AI talent in banks in the period May to September 2023, clearly showing that a significant investment is underway. Also, UK Government research[4] cites Finance and Insurance as the industry with job roles most exposed to AI vs all sectors. In summary, all the evidence is indicating that AI is a huge disruptor in the finance industry.

ARTIFICIAL INTELLIGENCE IS A HUGELY DISRUPTING TECHNOLOGY THAT FINANCIAL SERVICES FIRMS NEED TO QUICKLY ADOPT

So, everybody is doing it, but what are they actually using AI for? A better question would be "*what should I use AI for?*". Again, as a business leader, you don't need to understand how AI works, it's more important to focus on outcomes that can be addressed with AI. Use case examples range across front, middle, and back office:[5]

(I) Credit tasks like assessing credit risk, predicting default, or assigning a loan approval probability are all being worked on by banks already.

(II) Sales and Marketing teams can be augmented with AI to help identify new clients, or product opportunities, and to uncover insights about your existing clients' purchase behaviours that will allow you to optimize your products and services.

(III) Enhancing the customer experience with AI is becoming easier and easier, and I expect banks to start using generative AI-based chatbots for client interaction in 2024.

(IV) Trading is a pure data industry, so lends itself perfectly to AI. Almost every step can be enhanced with AI. You can understand the market better and predict price movements;[6] algo-traders have been doing this for years, and the technology is getting better and better. Specialist chatbots can eradicate trade support by capturing trader discussion and converting it into formatted instructions for the traders to agree. Error-checking deals, issuing SSIs, etc. The list goes on and on.

(V) Automation and operation efficiency are emerging as areas where generative AI will be able to help a lot, but for now you'll need to focus on use cases where the pay off is enough to cover the cost of developing a bespoke model. That said, the toolkit is expanding daily with low code, RPA, co-pilot support, etc., already making a big difference. A recent paper presented at ICAIF 23 showed very promising progress going beyond the prescriptive behaviour of RPA to a more spontaneous approach currently only possible with human operatives.[7]

(VI) So much of risk management is finding a needle in the proverbial haystack, and human minds are not well optimized to this task. However, AI never gets bored or distracted, so can offer much better outcomes, if well managed. Systems can detect anomalies in KYC data, or fraudulent behaviours, detect sanction busting schemes, etc. This technology is already relatively mature in cyber risk, where practitioners have always needed to trawl vast amounts of data quickly. Companies like Chainalaysis have been using AI to examine blockchain data and identify criminality, and their techniques are just as relevant to the traditional finance space.

(VII) There are many self-serving use cases where AI can be deployed to assist in the preparation of data for use in model building, for example by cleaning data. Synthetic data is emerging as a powerful tool that allows you to create a synthetic set of data that has the same characteristics as your production data but is anonymised and cannot be traced back to identifiable clients. This data can be used to test external vendors' capabilities, ensure that no production data is used in your development regions, or to allow academic collaborations that don't need strict NDAs.

(VIII) Although the return on investment may not be as high, many firms are building AI-based services to serve their internal use cases as this is a safer way to build experience with the technology. Generative AI fronted chatbots that allow semantic search of the firms' proprietary data, policies, procedures, etc., is a popular use case; as are more hidden use cases aimed at insider threat or cyber risk behaviours.

A general trend is that AI use cases assist a process, not replace it. The art of successful AI deployment is being able to make a mental leap from a business process that may be sub-optimal, and to then realise that AI could be deployed to enhance it. When you consider the return on your investment, consider if a 10% improvement in the outcome would make a meaningful difference. 10% optimization should be easily achievable for most AI-augmented processes, and many will do significantly better than this. Over time and with greater experience you can move to full automation, but starting with process assistance is the right way to start.

There is not one particular domain of finance that has an advantage over another when it comes to AI. There are advantages to reduce risk, identify alpha, increase efficiency, reduce cost, increase revenue, etc. All things that improve the way financial services works, ultimately providing a better outcome for the end consumer.

THE IMPACT WILL BE FELT IN ALL ASPECTS OF OUR BUSINESSES, NOT JUST IN THE QUANT TEAMS WHERE IT HAS BEEN USED FOR YEARS

Dream big but accept that AI is not a magic bullet. Some use cases are not solvable. Your data may not be large enough to train a model, or may not be clean enough – there could be irreparable bias issues that will produce unwanted outcomes. Some outcomes, like predicting security prices, are beyond AI's ability to predict, as there are just too many variables. Consider, however, if you really need to predict a spot price, or would it still be valuable to predict the direction of price movement and maybe even the magnitude. In many cases, understanding the nature of the output probability is valuable in itself.

HOW DO I INTRODUCE AI TO MY FIRM?

Now you know what AI is, and understand that it is strategically important, let's discuss what it takes to build an AI capability in your firm.

There are different operating models for AI that can be adopted, depending on your firm's strategy and goals. In financial services most firms don't have widespread AI experience yet, so it is advised to build a centralised multidisciplinary capability to lead your AI effort. The cost of these teams should be considered too, where larger firms may be able to build the full capability in house, other firms may decide to outsource aspects of the function to third parties. It doesn't matter if your operating model is in-house, outsourced or hybrid, the same components need to be considered.

- Strategy – First, strong executive sponsorship is vital; as building an AI capability is a significant challenge, and is likely to face many headwinds before it is successful, the sponsor should be able to provide air cover to the team while they focus on the capability.
- Data Scientists – These are the specialists with the technical knowledge to build AI models from scratch. These specialists are highly trained and are therefore also highly sought after. Ideally you would have some data scientists in-house, but there are many firms offering data science as a service. 'Try before you buy' may be a good strategy to get your AI capability started.
- Risk, Legal and Compliance – AI presents new challenges that need to be understood so that you can build within your agreed risk appetite. Model Risk specialists have existed for years generally looking at the work of quant teams or central planning teams, but their experience is vital to AI success. Likewise, data risk people are needed to ensure that the data that feeds your models is being used appropriately, and that the output does not have unacceptable biases. Ethics is another significant risk where AI brings specific concerns. Depending upon your budget you will decide if you hire new staff, or upskill existing staff, but the costs of ignoring AI risks will be very significant as regulators are paying very close attention to how AI is being used by firms.
- Technologists – Your existing IT team have a large part to play to ensure that hungry AI models have access to the data they need and have the compute available to run. They will decide if an 'on prem' or cloud architecture is right for your needs; both work, and there are distinct advantages to both too. One oft-ignored role is that of data engineer, as these people are expert at wrangling the data into a format that a data scientist can gain value from. Many data scientists report that a significant part of any project is taken preparing the data. As you scale it is prudent to consider assigning this task to lower paid data engineers rather than letting expensive data scientists do it for you.
- Business Experts – The last set of people you need are business experts who understand what AI technology is capable of. These experts then engage with your business leaders to understand their challenges and scout for opportunities where AI can add value. Finding use cases that are valuable to your firm is the most important task that any AI capability must achieve, without this you run the real risk of wasting time doing AI theatre.

THERE ARE NEW RISKS TO CONSIDER, BUT THE ROUTE TO ADOPTION IS NOT INSURMOUNTABLE

A popular operating model is a 'hub and spoke' operating model, where you build expertise centrally and define standards so that teams in your business can build on that to solve problems. This gives you consistency, risk management, scalable tooling, but also business flexibility to focus on the challenges most important to each team. Obviously, this is not the only operating model that you could consider. It could be more appropriate to embed AI capability in the product lines, or to focus on one aspect of your business that will generate the most gains by leveraging AI.

The US National Institute of Standards and Technology (NIST)[8] does a good job of describing what good AI risk management looks like. Standards are coalescing around the world: the EU AI Act will form a set of crucial guidelines in Europe, and other regions are following a similar path. NIST looks at risks through three lenses: risks to people, organisation and to ecosystem. AI risk goes deeper to consider these categories in terms of data risks including privacy and bias, legal risks covering who owns the data and what contractual protection you have for using it, model risks covering if the model does what you think it does, cyber risks, cost risks, strategic risks, ethics risks, etc. AI risk is a complex space, and it is well worth upskilling your risk management function so they can help you achieve the strategic outcomes you desire safely. For a deep dive into AI Risk, an OECD report[9] published in 2021 covers the space very well.

'Garbage in, garbage out' is a phrase often repeated when discussing an AI data pipeline. The good news is that the finance industry has been doing a lot of work over the last few years improving the quality of its data and AI is the payoff. Having good data governance practices converges with the data science talent and techniques to suddenly make the impossible possible. This is why AI technology that has been around since the 1960s feels so new. Cloud computing is also an accelerating factor, as it allows you to manage your data at scale and connect to data science solutions more easily. Well-managed data could be 'on prem' just as easily but there are advantages to cloud computing, for example scaling and cost. However, it's easy enough to build an AI model that uses 'on prem' data as it is to use cloud-stored data.

THE SKILLS AND KNOWLEDGE OF EXISTING EXECUTIVES IS CRITICAL AS THEY WILL IDENTIFY THE MOST IMPACTFUL USE CASES TO WORK ON

Undoubtedly your firm already has experienced technologists who can leverage your data and deliver AI solutions, but it is likely you'll need to do some work to adjust your risk governance and run some education both in technology and within the business teams. One big mistake that happens too often is that leaders delegate AI to teams who are not equipped to implement it. Technology tends not to have deep understanding of the business goals or close enough relationships. Risk and compliance people are often too cautious with a technology they don't understand. Data people naturally obsess over data structures and governance, but don't tend to have the magic touch needed to build an AI capability.

Building a better data architecture and growing deeper understanding of data and AI in your risk teams will lead to better cyber risk too. There is a naïve view that AI will 'get away from us', yet I believe that the very well-regulated financial services industry is well positioned to adopt this technology, because we already have the muscle strength of good governance. Robust risk practices and strong oversight are traits that are aligned with AI success. When we pair that strength with a desire for meaningful business outcomes it becomes a powerful driver for success.

It is hard to define what it takes to build this AI capability; however, the leadership needs to be aligned very closely to the business. To identify good AI use cases, you need people who can discuss business challenges articulately with your business leaders and clients, but who also have a deep enough understanding of AI to identify when a challenge should leverage AI to build a solution. Typically, at this point the AI capability engagement team will bring their data scientists to the discussion to dive deeper and validate the proposed solution. People who can stand in both data science and the business are the ones you should prioritise identifying as these are the people who will make or break the AI capability at your firm.

Hopefully you now understand what AI is, its value to your business and how you may go about creating an AI capability at your firm. If you already have a capability, you can use this article to assess if your firm is on track. AI is likely the largest single disrupter to financial services in a long while, both revolutionising how we run our firms, but also how our clients interact with our products and services. AI-enabled firms will naturally focus on the quality of their data, which will lead to further efficiencies, and they will develop more agile change practices that when paired together with AI will make them very hard to compete with. This is therefore a critical time to invest for the future health of your firm.

Notes

1 Maple et al, The AI Revolution: Opportunities and Challenges for the Finance Sector, The Alan Turing Institute, 2023
2 The impact of AI in financial services: opportunities, risks and policy considerations | Policy and Guidance | UK Finance), 2023
3 Evident AI Index, Banks Key Finding Report, November 2023
4 GOV.UK Impact of AI on UK jobs and training (publishing.service.gov.uk), November 2023
5 Evident AI Outcomes Report, October 2023
6 Gu, S., Kelly, B., Xiu, D. Empirical asset pricing via machine learning. The Review of Financial Studies, 33(5), 2233-2273, 2020. DOI: 10.1093/rfs/hhaa009
7 ICAIF '23: Proceedings of the Fourth ACM International Conference on AI in Finance November 2023 Pages 73–81
8 AI Risk Management Framework | NIST, January 2023
9 OECD (2021), Artificial Intelligence, Machine Learning and Big Data in Finance: Opportunities, Challenges, and Implications for Policy Makers, https://www.oecd.org/finance/artificial-intelligence-machine-learningbig-data-in-finance.htm.

AI for investment analysts: Triangle of Silicon Valley, Hollywood & Wall Street

Mark Geene, Senior Investment Consultant, PGGM

In 2023 the advances in AI caused extensive strikes in Hollywood hitting movie studios and streaming platforms. Writers and actors were scared AI could and would replace their plots, voices and faces. AI has the ability to revolutionize and further digitalize movie experiences. The finance industry has been similarly revolutionized by the speed and access of digital technology amplified by recent developments in AI that can replace as well as augment investment analysts. Similar as in Hollywood and (un?)surprisingly, several of these advances in AI concern voice and facial techniques. These techniques identified and underscored biases investors have in relation to voice and physical appearances of people.[1] Awareness of these techniques and the results of related studies during manager selection and fundamental company analysis is required to, symbolically, not end up like the villain in a James Bond movie.

Recently several academic studies have revealed, again not surprisingly and in line with previous research on 'good looks', that also investors are influenced by the facial attractiveness of individuals. Several papers using AI to score facial attractiveness indicate that unattractive PM's outperformed attractive PM's by over 2% annually. However, papers also show that PM's with the looks of 'Brad Pitt' or 'Angelina Jolie' attracted significant higher inflows, especially when their photos were accessible to investors. Authors of these papers indicate that potential explanations are predominantly related to biases of end investors including the so called 'beauty effect'.
Related studies show good looking analysts produced more accurate earnings forecasts than less attractive analysts. In addition, their stock recommendations were more profitable. These analysts gained more media exposure, had better connections to institutional investors and received more internal support from their employers. Again, evidence of the beauty effect.
This raises interesting puzzles and moral considerations when performing a due diligence on a manager. Should one abolish face-to-face with PM's eventhough these interviews are currently considered a crucial part of each due diligence? Should one hire due diligence analysts, that by using AI-techniques, score high on attractiveness and trustworthiness as they can 'extract' more and better information during the conversations? Is it ethical to use photos and actually rank PM's on the basis of their looks? And because also PM's pick up the results of these studies and if one persists on having meetings face-to-face meetings: should one try to pull on the supposedly ugly looking PM's hair or skull to test if this is a real face and not a mask from Mission Impossible?
Parallel research signaled the presence of the beauty effect in the corporate world. Using AI-techniques authors show that attractive bank CEO's received higher pay, eventhough they were prone to lead to lower shareholder returns. Investment analysts should pay attention to this 'beauty trap', but even more importantly to the tone of voice of CEO's as the number of face-to-face meetings with CEO's reduced since Reg FD in 2000. Many asset managers already use machine learning to go over transcripts of earnings calls and therefore its value will erode over time. Likewise, CEO's have been trained to use certain words and sentences to counter or 'game' the algo. Therefore, recently several asset managers started researching how words are actually spoken. These techniques identify for instance hesitation and subsequently reveal the emotions and the true sentiment by assessing their tone of voice. What are the next steps in this rat race between CEO's and investment analysts? More voice training with actors by CEO's like common for politicians or will the CEO-analyst encounters become more boring and emotionless? Will CEO's use the Steven Hawkings Voice Generator during an earnings call, while not transforming into a voice (and AI on the loose) like Hal9000 in A Space Odyssey?

How should investors take the findings of AI research on the role facial attractiveness and voice into account to debias their investment processes? Like for all the behavioral biases and heuristics an investor first and foremost has to recognize the existence of these and be humble enough to acknowledge that even he or she, like all financial actors, will be influenced by it.
The above findings indicate that you are able to at least build some 'defenses' in your investment process against it and even use it to your own benefit. However, it remains a rat race between PM's and CEO's trying to convey an image and perfect story, smile and face towards prospective and current investors and analysts. In between (or their techniques and tools) are the servers of OpenAI, Alphabet and Amazon as well as moral considerations. For sure, Wall Street pros have to travel to Silicon Valley to pick up the latest AI techniques, but an extensive visit to Hollywood is necessary to not be influenced by a mirage and siren song.

Note

1 Including papers from the Journal of Accounting Research, Journal of Behavioral and Experimental Finance, Journal of Economics and Business or (NBER) working papers of academics.

There is Alpha in Large Language Models

Tjeerd van Cappelle[1]

Artificial Intelligence and Large Language Models (LLMs) are adapted by society on an increasingly larger scale. Consequentially, they also make their way into the investment industry. This article discusses LLMs and their application in investments.

The first part of this article discusses LLMs in general and how they impact the investment industry. It describes the various kinds of LLMs, the risk involved in the use of LLMs and the types of applications they can be used for. The author argues that traditional fundamental analysts aren't wired to identify which textual information is already captured by the numbers (see also earlier research by Van Cappelle and Niesert, 2021). LLMs, on the other hand, can be trained for specific tasks and can be used at scale. This is true for many models, including those that predict future revenue growth, costs, or analysis of Environmental, Social, and Governance (ESG)-related topics. At the end of the first part, the article elaborates on how investment firms will need to adapt to use LLMs or work with data derived by LLMs.

The second part of this article focuses on quantitative strategies. The author uses forecasts that are exclusively available by LLMs to conduct an empirical analysis. The analysis compares the LLM-derived forecasts with more traditional forecasts. Additionally, the author scrutinizes the LLM-derived forecasts by formal tests that involve established factor models. The analysis shows that these new forecasts are complementary to traditional quantitative factors and robust to implementation choices. The article concludes that quantitative investors will have to embrace the use of LLMs to stay relevant and add new sources of robust alpha.

THE RISE OF LARGE LANGUAGE MODELS

Automated processing of text (Natural Language Processing) has existed for many years. However, NLP models didn't enter the mainstream media until the launch of ChatGPT in November 2022. Since this time, ChatGPT has impressed society at large both by its breadth of applications and its seemingly understanding of questions. As a result, 2023 is seen by many as the year of Large Language Models (LLMs), and reservations about using LLMs have largely disappeared. To many, they are seen as a silver bullet that can fix anything.

The current success and wide acceptance of LLMs has been in the making for quite a few years; in 2018, two models emerged that are the basis of today's most well-known LLMs:

- The Generative Pretrained Transformer (GPT) (Radford et al., 2018) and
- Bidirectional Encoder Representations from Transformers (BERT) (Devlin et al., 2019)

These two models represent different types of Large Language Models: the decoder model and the encoder model.

Decoder models are LLMs trained to generate text that makes a good conversation. Note that the primary learning objective, when decoder LLMs are trained, is not to provide responses that are factually correct.[2] Rather, the objective is to provide responses that are liked by the person asking the question. OpenAI's GPT-4 model is the state-of-the-art decoder model.

Tjeerd van Cappelle
Founder & Managing Director at aiLiftoff

OpenAI has surrounded the GPT-4 model with supporting software and deep learning models to mitigate the inherent risks in using a decoder model. For instance, to prevent generating racist text. As decoder LLMs generate text, they are often referred to as Generative AI.

Encoder LLMs are good at summarizing text. An example would be a classification task in which the sentiment of a certain text is classified as being positive, negative, or neutral. One of the state-of-the-art encoder models is Microsoft's DeBERTa[3] LLM. Where decoder LLMs were dubbed Generative AI, Encoder LLMs can be either referred to as Interpretive AI or Predictive AI. Interpretive AI when it summarizes the current situation. Predictive AI when it classifies future events based on text.

A special type of classification is the so-called zero-shot classification. A zero-shot classifier is an LLM that isn't trained on a pre-defined set of labels. Instead, it receives as input both text and several labels to choose from. This makes zero-shot classifiers very flexible as, in theory, they can be used at any classification task.

APPLICATIONS OF LARGE LANGUAGE MODELS IN THE INVESTMENT INDUSTRY

With this broad overview of LLMs in mind, one can think of a variety of applications in the investment industry.

Generative LLMs can be used to support the production of investment reports. Interpretive LLMs can be used to identify current trends at companies. For instance, they could be used to identify ESG related risks. Predictive LLMs can be used to forecast company fundamentals, like future earnings or future revenue growth. Since this article describes the application of LLMs to company analysis, the focus is on predictive and interpretive LLMs.

The use of LLMs in analyzing companies has three distinct advantages:

1. LLMs can be applied at scale. Where an analyst at an investment firm typically covers 30 to 50 companies, an LLM can analyze thousands of companies.
2. LLMs produce consistent and comparable analysis. Where one analyst may have a different interpretation from another analyst, or even her own judgement could vary with time, LLMs will produce the same analysis for the same text.
3. LLMs can be trained for specific tasks. Humans are very good learners of languages, but they are not good at "not reading" certain text. LLMs on the other hand can be trained to focus on a specific task and not be distracted by other information present.

Obviously, the use of LLMs is not without risk. First, while LLMs produce consistent results, their results can still be prone to biases. As pointed out in earlier research (Van Cappelle and Niesert, 2021), at any point in time, just 10 companies account for more than 50% of all company-related news articles and social media posts. This leads to huge biases towards these companies. Additionally, it is well known that neural networks, and therefore LLMs as well, can easily learn relationships that don't exist.

Besides the risks mentioned here, there are also considerations to be taken around information security and potential intellectual property infringement. Some of the most powerful LLMs can only be invoked through an Application Programming Interface. This means that the textual data is shared with the provider of the LLM. In the case of a proprietary text source there is the risk of information leakage. Considering LLMs are trained on large swaths of text, that are available on the internet, there are increasing concerns and claims that the LLMs might infringe copyrights.[4]

In summation, while LLMs are clearly powerful, specific skills are required to understand and mitigate the risks related to their implementation.

WAYS TO USE LARGE LANGUAGE MODELS AS AN INVESTMENT FIRM

Let's look in a little more detail at how investors can benefit from the application of LLMs. What are the ways to access LLMs, and for which type of analyses can LLMs be useful?

There are broadly three ways investors can use LLMs. First, investors can use LLMs that are publicly available. For instance, there are publicly available LLMs that can classify text according to financial sentiment, such as finbert-tone, a model developed by Huang et al. (2022). The second way would be by developing an LLM oneself. This would have the benefit of having complete proprietary insights. However, the effort, investment in human capital, and the risks involved are easily underestimated. The third way would be to use commercially available forecasts or insights made by LLMs, or to license a commercial LLM. The latter would be an option in case theLLM is used to process proprietary textual data.

> INTERPRETIVE LLMS AND PREDICTIVE LLMS CAN IDENTIFY AND PREDICT TRENDS AT COMPANIES

The type of analyses which could be done with LLMs are probably endless. Yet a few of the common analyses include:

- Sentiment analysis. There are numerous commercial data sets which determine sentiment surrounding companies. Apart from those commercial data sets, there are various publicly available models for sentiment analysis, amongst others yiyanghkust/finbert-tone (Huang et al., 2022) and ProsusAI/finbert (Araci, 2019)
- ESG and SDG classification. Given the need for more ESG data among investors, LLMs provide a good way to extract ESG information, and information related to SDGs from

unstructured data. Again, there are lots of commercial data sets created by applying LLMs, and there are public LLMs for ESG or SDG classification as well.
- Forecasting company fundamentals like revenue and cost. There are vendors who sell data sets that forecast revenue, costs, and other fundamentals based on company's management discussions and press releases.
- For the analysis of human capital and governance one could use commercially available specialty data sets which analyze the language surrounding appointments and departures of company management.
- Finally, for thematic or one-off research it is recommended to use a zero-shot classifier. The upside of zero-shot classifiers is that they don't need to be trained for a specific task. Which is especially useful in case there are only a few labeled samples to train on, or the task at hand is a one-off. The flipside is that zero-shot classifiers aren't trained specifically for a task. That said, currently one of the most popular zero-shot classifiers is also trained on the financial phrase bank (Laurer et al., 2022)

ARCHITECTURE OF LARGE LANGUAGE MODELS

To provide more context to what a Large Language Model and its surrounding software might consist of, let's look at the most common parts of a language model. Broadly the model consists of a:
- pre-processing module
- tokenizer
- deep-learning network
- task-specific head
- post-processing module

Figure 1 gives an overview of the architecture with an example.

Figure 1
Architecture overview with example text

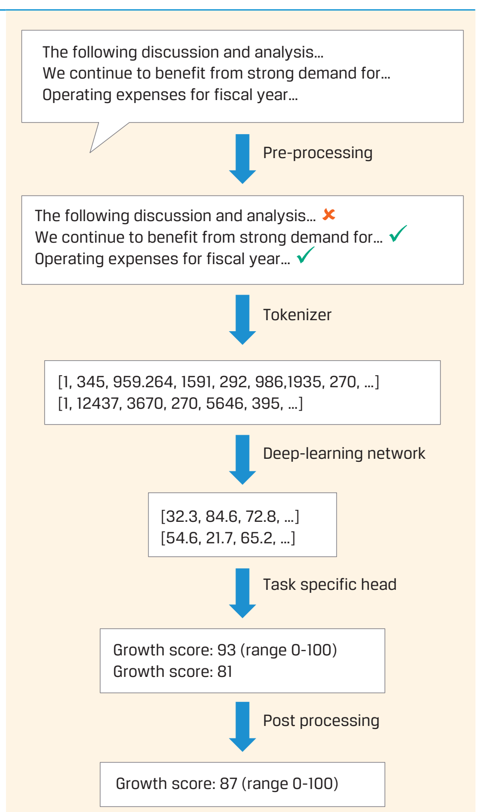

The pre-processing module is where the raw unstructured data is reformatted and possibly filtered for consumption by the LLM. So, management discussions or press releases will be converted to machine readable text. Text that runs over multiple pages will be 'glued' together; sentences or paragraphs are restored as intended. For training purposes, the training target will be calculated. Finally, there is most of the time a 'relevancy-filter' in place to take out irrelevant text. For instance, disclaimers, introductory text, or boilerplate language that is used in every report or press release of the company are not of interest.

The tokenizer is a part of the model where text is converted to numbers. This is a necessary step. When developing a model from scratch, a choice needs to be made, whether to re-use an existing tokenizer or train a tokenizer from scratch. Often an existing tokenizer is re-used, which creates the next choice: a) use a tokenizer that is specific for financial text, or b) use a tokenizer that is more generic. What works best is often an empirical question.

THE APPLICATION OF LLMS WILL HAVE AN IMPACT ON HOW INVESTMENT FIRMS OPERATE

The deep-learning network can be considered the heart of the LLM. Again, there are a few choices to be made. Will the network be trained from scratch or will a pre-trained network be re-used. Often the choice is to start off from a pre-trained network. That still leaves many questions open: go with a Bert-based model or use larger models like RoBERTa or DeBERTa? Use a pre-trained model which is trained on financial text or use a pre-trained model that is trained on more generic text?

The task specific head is the final layer that is put on top of the deep-learning network. In the case of a sentiment classifier, the task specific head would provide probabilities for the text to be positive, negative, or neutral.

The post-processing module is where the outcome of the task-specific head is processed further to an outcome. Say a large document is processed in multiple parts. Then the score for each part needs to be aggregated into a score for the overall document. Post-processing typically includes some weighting scheme or perhaps more complex functions.

For certain tasks it can be useful to have multiple models work together. In the case of ESG analysis, it could be useful to have a first model that identifies paragraphs of text that discuss a certain ESG related topic. A second model trained specifically on the ESG subject in question could further analyze the paragraphs involved in the subject.

IMPACT ON INVESTMENT FIRMS

The application of LLMs by investment firms will have an impact on how they operate. Specifically, it will have an impact on the distribution of human resources. As established earlier, there will need to be an allocation of human resources to the use of LLMs to enjoy the new possibilities that they offer.

Investment firms that decide to develop LLMs themselves, or use publicly available LLMs, will need to make a considerable investment in human capital to develop the know-how to create, test and maintain these LLMs. The requirement of special skills in developing and operating LLMs is proven by the emergence of new job titles, like "prompt engineer",[5] which didn't exist in 2021. Furthermore, investment firms that choose to buy data created with LLMs, choose to license commercial LLMs, or use public LLMs, will need to have a basic understanding of what they are using.

The shift of resources to the development and application of Large Language Models needs to be paired with a gain in efficiency. The most obvious candidate for efficiency gain is in the field of traditional corporate analysis. Investment firms will either decide to automate part of their in-house analysis, or they use the analysis done with the help of LLMs to steer their in-house analysis towards the most promising investment opportunities.

Besides efficiency gains, new possibilities will emerge that simply didn't exist before. A powerful feature of Large Language Models is that they can be trained for specific tasks without being distracted by anything that is irrelevant to the task at hand. Earlier research by Van Cappelle and Niesert (2021) has shown that LLMs are capable of forecasting future revenue growth that cannot be inferred from the income statements of companies.

In short, the impact of Large Language Models for investment firms will be a shift towards new roles in the field of AI, efficiency gains in the field of traditional analysis and access to information and forecasts that simply wasn't accessible before.

IMPACT ON QUANTITATIVE INVESTMENT STRATEGIES

The second part of this article will discuss the impact on quantitative strategies. Quantitative strategies are strategies that are fully data driven. First the data and the test setup will be explained. Next the results will be discussed.

Quantitative strategies are employed by quantitative hedge funds, mutual funds, and pension funds. Quantitative hedge funds usually employ shorter term strategies with holding periods that are expressed in days or weeks. Mutual funds and pension funds usually employ strategies with somewhat longer holding periods that are expressed in months or quarters. Therefore, the test will examine characteristics over a variety of investment horizons.

To study how quantitative strategies could benefit from LLMs, the author explores a commercially available data set. This data set is created with the use of LLMs. The test uses two forecasts that are part of this data set: Revenue Surprise forecasts and Growth Acceleration forecast. To compare the results with a more traditional dataset, the test also includes consensus revenue growth forecasts by institutional broker analysts.

The Revenue Surprise forecast aims to predict the amount of surprise in the revenue growth that will be reported in the next quarter by a company. The forecast is created as soon as companies publish their quarterly report. In the forecast a separation is made between the revenue growth that can be forecast using information from past income statements, and what extra information is contained in the management discussion and press release of a company. The extra information is captured by the Revenue Surprise forecast. It is expressed in a number ranging from 0 to 100. Values over 50 indicate a positive expectation for growth stemming from the textual data, whereas values below 50 indicate a negative expectation. The LLM, with which these forecasts are created, is trained on data from 1994 until 2007.[6]

The Growth Acceleration forecast is a more sophisticated forecast that considers the previous 10 earnings reports and predicts whether revenue growth is on an accelerating trend or not. The forecast is expressed in numbers ranging from 0 to 100. Values over 50 indicate that revenue growth is expected to accelerate. Those below 50, on the other hand, indicate that growth deceleration is expected. Like with the Revenue Surprise forecasts, the model to produce the Growth Acceleration forecasts is trained on data from 1994 until 2007.[6]

Under normal circumstances, the Revenue Surprise and Growth Acceleration forecasts are available within an hour after an earnings announcement. Yet, to be on the safe side for testing purposes, it is assumed that the forecasts are available as of the day after the earnings announcement.

The traditional data set contains forecasts for next quarter's revenues made by analysts. The data is from a commercially available database of institutional broker estimates. Every time an analyst updates the forecast for the next quarter's revenue, the data is updated in the data set. The forecasts of different analysts are averaged daily per company. The average of these forecasts, when compared to previous revenue numbers, gives what is referred to as a Consensus Growth forecast.

The stock universe on which these data sets are tested consists of the 3000 largest stocks in the US market by market capitalization. The 10 full years spanning from 2013 until 2022 were chosen to conduct the research. The start date of 2013 ensures that the test is free from any possibly forward-looking biases, as all training and parameter tuning for the LLMs happened before 2013. While as of 2013, the search for alpha from textual sources had already begun. Quantitative investors had already been applying NLP techniques to management discussions, press releases, and analyst calls.

STRATEGY EXPLORATION

For the first test, quintile portfolios are created daily for various investment horizons. It is best to describe the daily creation of portfolios in a step-by-step procedure:

1. The investible universe is determined. All stocks are selected that are among the 3000 largest stocks in market capitalization.
2. Based on the investment horizon, stocks are filtered out. Say, the investment horizon is 2 weeks, then any stock that had an earnings announcement more than 2 weeks ago is excluded.
3. The stocks are ranked according to their forecasts. Stocks with the best forecast end up in the first quintile, stocks with the worst forecast end up in the bottom quintile.
4. All the stocks in a quintile form an equally weighted portfolio.
5. From the best and the worst quintile, a long-short portfolio is created. The long/short portfolio is 50% long in the best quintile and 50% short in the worst quintile.

For each forecast, this procedure is repeated with investment horizons of 1, 2, 5, 10 and 20 weeks.

The resulting Information Ratios of the respective test/forecast combinations are shown in figure 2.

Figure 2
Information Ratio as a Function of Investment Horizon

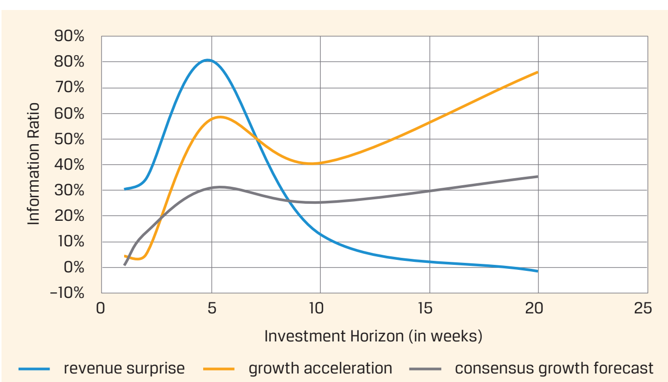

The chart shows that the Revenue Surprise forecast yields higher information ratios over shorter investment horizons than the Growth Acceleration forecast. A possible explanation is that the Growth Acceleration forecast is more sophisticated than the revenue forecast. As such, one might expect that it would take longer for the market to figure this out.

It is also noticeable that the information ratio of the Revenue Surprise forecast increases as the investment horizon increases from 1 week to 5 weeks. This doesn't necessarily mean that the information is more powerful after 5 weeks. It might also reflect the fact that with an investment horizon of 5 weeks, the quintile portfolios get better diversified. This is because more companies had their earnings announcement in the last 5 weeks than in the last week, hence each quintile portfolio will include more names. Furthermore, companies have the tendency to have their earnings announcements around the same dates (earnings season). This causes even sparser portfolios between earnings seasons.

The information ratio of the Revenue Surprise forecast fades out as the investment horizon becomes longer. This indicates that the signal is better suited for investors who pursue strategies with shorter holding periods.

REVENUE SURPRISE AND GROWTH ACCELERATION DATA UNLOCK INFORMATION THAT IS NOT AVAILABLE BY TRADITIONAL MEANS

The Information Ratio of the more traditional Consensus Growth Forecast is for any investment horizon lower than at least one of the forecasts produced by LLMs. This implies that the Revenue Surprise and Growth Acceleration data unlock information that is not available by traditional means.

Finally, it is observed that the information ratio of the Growth Acceleration forecast shows a small dip at a 10-week investment horizon, or a small bump at the 5-week horizon. The same effect can be seen to a lesser extent with the traditional Consensus Growth forecast. There is no obvious explanation for this phenomenon.

EVALUATING PERFORMANCE

Based on the initial analysis it is decided to further explore four forecast/investment horizon combinations:

1. Revenue Surprise – 1-week horizon
2. Revenue Surprise – 5-week horizon
3. Growth Acceleration – 20-week horizon
4. Consensus Growth – 20-week horizon

For the 5- and 20-week investment horizon, the creation of quintile portfolios is further enhanced to control for size. The investible universe is split into the 1,000 largest stocks and the 2,000 stocks that are smaller (like how the Russell 1000 and Russell 2000 are constructed).[7] For each group the quintile portfolios are created through the same procedure as before. The long/short portfolio is constructed by being 25% long in the best quintile of the 1,000 largest stocks, 25% long in the best quintile of the 2,000 smallest stocks, 25% short in the worst quintile of the 1,000 largest stocks and 25% short in the worst quintile of the 2000 smallest stocks.

With the 1-week investment horizon the size-control enhancement is not feasible as there would be too few stocks eligible at times to create quintile portfolios. So, for the 1-week investment horizon we keep the portfolios as before. However, a different risk measure is taken for the 1-week long/short portfolio: to mitigate the effects caused by few stocks between earnings seasons, the amount of risk is adjusted based on the breadth of the portfolio.[8]

Figure 3
Cumulative performance chart of four long-short portfolios

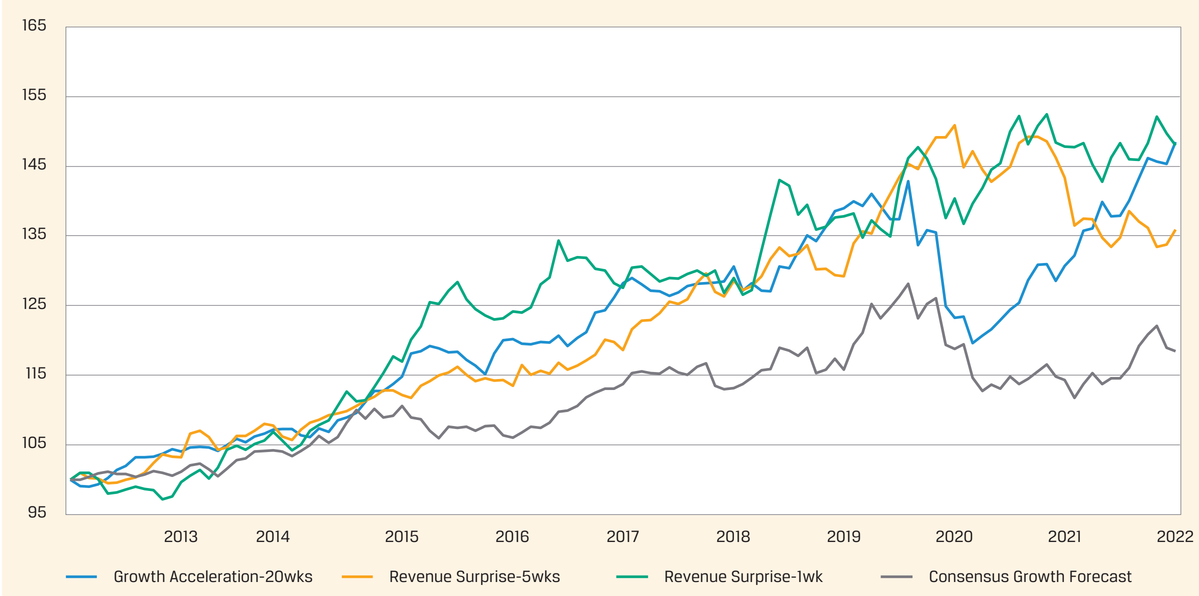

The cumulative performance of these four strategies is depicted in figure 3.

It is immediately clear that each strategy has its peaks and troughs, but visually they don't seem to happen at the same time. In other words, the three LLM strategies are not overly correlated. Numerical analysis learns that among the LLM forecasts the highest pair-wise correlation is 0.33 (between the 1-week and 5-week Revenue Surprise strategies). The highest pair-wise correlation between the Consensus Growth strategy and the LLM strategies is 0.55 (between Growth Acceleration and Consensus Growth).

The performance of the 4 strategies is summarized in table 1.

The table confirms that the LLM-based forecasts deliver higher information ratios than the more traditional growth forecasts.

Additionally, the table shows that the 1-week investment horizon Revenue Surprise portfolio benefits a lot from the risk adjustment based on the number of stocks in the portfolio. In the initial set up the 1-week investment horizon gave an information ratio of 0.3, with the risk adjustment, the information ratio increased to 0.6.

LLMS DELIVER ROBUST ALPHA

As the alpha potential of LLMs is shown, further analysis is conducted into the robustness of the alpha.

To analyze the return characteristics further, the returns are regressed on the Fama and French 5 factor model (Fama and French, 2015). The regression coefficients as well as the alpha are presented in table 2.

Table 1
Summary performance and characteristics of 4 strategies

Forecast	Revenue Surprise	Revenue Surprise	Growth Acceleration	Consensus Growth
Investment horizon	1 week	5 weeks	20 weeks	20 weeks
Annualized Return	4,4%	3,2%	4,2%	1,9%
Annualized Risk	7,4%	4,8%	5,5%	5,3%
Information Ratio	0.59	0.68	0.75	0.35

Table 2
Summary performance and characteristics of 3 strategies, regressed on the 5 Fama and French factors

Forecast	Revenue Surprise	Revenue Surprise	Growth Acceleration
Investment horizon	1 week	5 weeks	20 weeks
Market exposure	0.00	0.02 *	0.00
Size exposure	-0.05	-0.03 *	-0.05 *
Value exposure	-0.01	-0.08 *	-0.06 *
Profitabilitiy exposure	-0.15 *	-0.12 *	-0.04 *
Investing exposure	-0.09	-0.10 *	0.11 *
Alpha	5.0% *	3.4% *	3.9% *

* denotes exposures that exceed the 95% significance level

The table tells us that neither strategy is overly exposed to any factor. The strategies based on this new source of information are complementary to established quant factors. It can also be seen that the factor models are not explaining away the excess return of the long-short portfolios; again, a sign that this new source of information is additive to established factors.

All in all, it can be concluded that the forecasts based on LLMs offer alpha opportunities for active investors regardless of their preferred investment horizon.

The Fama-MacBeth (Fama and MacBeth, 1973) regression is widely accepted as a formal test to see whether exposure to a factor is rewarded. Hence, this empirical study includes such a regression as well.

The Fama-MacBeth procedure is as follows:
1. Several (at least 7, preferably more) portfolio-return series serve as input to the procedure.
2. A time series regression is done of each of the portfolio excess returns on the factor returns to establish the exposure to factors for each portfolio.
3. A cross-sectional regression is done for each period of the portfolio returns on their factor exposures from step 2. This cross-sectional regression produces the factor premia for each period.
4. The resulting time series of factor premia are averaged, and the standard errors are calculated. This gives an observed reward as well as a confidence interval.

The Fama-MacBeth regression is applied to the Revenue Surprise forecast with the 5-week investment horizon as well as to the Growth Acceleration forecast. The quintile portfolios of both the larger stocks as well as the smaller stocks are used in the regression; so, there are 10 portfolios which is sufficient for the cross-sectional regressions. The procedure is followed for the Fama and French 3 factor model (Fama and French, 1993) as well as the Fame and French 5 factor model (Fama & French, 2015). In the 3-factor model, the factors Profitability exposure and Investing exposure are not included.

Table 3 reports the reward for the respective forecasts as well as their confidence interval, as found with the Fama-MacBeth tests.

Table 3
Summary of Fama MacBeth test for the Revenue Surprise and Growth Acceleration forecasts

Factor reward	Estimated reward	90% confidence interval
Revenue Surprise – 3 factor model	3,4%	0.8% – 5.9%
Revenue Surprise – 5 factor model	3,5%	1.0% – 6.0%
Growth Acceleration – 3 factor model	3,3%	0.1% – 6.6%
Growth Acceleration – 5 factor model	4,2%	1.3% – 7.1%

The results show that the cross-sectional reward for both Revenue Surprise as well as Growth Acceleration exists indeed. They are statistically significant different from 0 (at 95% confidence), as the 90% confidence interval is strictly positive.

A final check for robustness involves a change in constructing the long-short portfolio. With the Revenue Surprise forecast, so far, two long-short portfolios have been created. One where plainly 5 quintile portfolios were created and the long-short was the difference between the best and the worst quintile. Next, a double sort on size and Revenue Surprise forecast was created, where the long-short portfolio was made up of the difference between the best large and small quintiles versus the worst large and small quintiles. Now a third long-short portfolio is added to the mix. This time it is controlled by sector.[9] We divide the Industrials, Consumer Cyclical, Health Care, and Technology

Figure 4
Cumulative performance of variations of the Revenue Surprise strategy with 5-week investment horizon

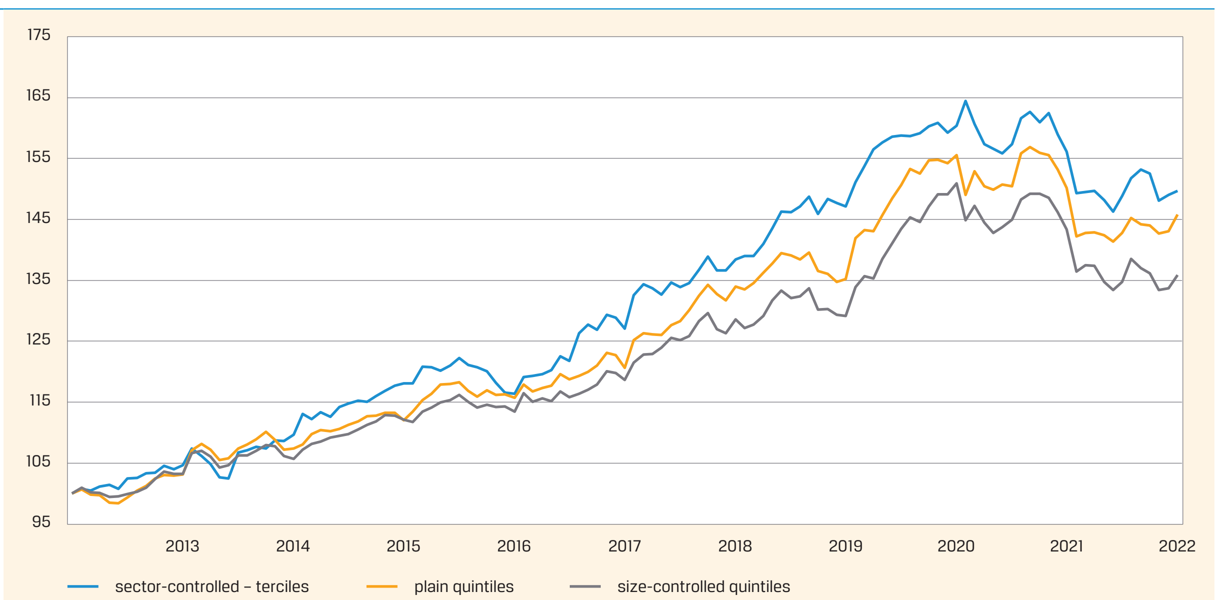

sector in terciles. The long-short portfolios reflect the difference between the best terciles per sector versus the worst sector per tercile. The cumulative return series of the three different choices for a long-short portfolio are compared as a check that the results so far are not driven by a seemingly innocent choice in constructing the portfolio.

The cumulative returns are plotted in figure 4.

The chart shows that the patterns of the strategies are quite similar. Furthermore, it is visible that the choice for size control in the factor analysis didn't inflate results. On the contrary, results without size-control or with sector-control instead of size-control are even better.

LLM-DERIVED DATA CAN OFFER COMPLEMENTARY ALPHA

Various tests were applied to the LLM-derived data. By doing all these tests it is established that the alpha of the LLM-derived forecasts is robust, and not influenced by choices in the way portfolios are constructed. Additionally, a formal Fama-MacBeth test pointed to an annualized reward of roughly 3.5% for both the Revenue Surprise as well as the Growth Acceleration forecast.

CONCLUSION

In this article the rise of Large Language Models has been discussed, and the impact on the investment industry has been analyzed. What is clear is that investors will need to adapt. Fundamental and quantitative investors alike will need to increasingly shift resources to the use of LLMs.

For fundamental investors LLMs will bring efficiency which can be used to augment or steer their traditional analysis.

It was also shown that forecasts derived using LLMs can produce alpha. This alpha is complementary to what is offered by more traditional and established data. Quantitative investors have no choice but to embrace the use of LLMs to stay relevant.

Literature

- Araci, D., 2019 FinBERT: Financial Sentiment Analysis with Pre-trained Language Models. arXiv preprint arXiv:1908.10063.
- van Cappelle, T. & Niesert, R., 2021, De cijfers voorbij: vanuit tekst omzetgroei voorspellen, VBA Journaal, nr. 147, 26-32.
- Devlin, J., Chang, M. W., Lee, K., & Toutanova, K., 2019, BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In NAACL-HLT, pages 4171–4186, Minneapolis, Minnesota.
- Fama, E., & French, K., 1993, Common Risk Factors in the Returns of Stocks and Bonds, Journal of Financial Economics, 22, 3-56.
- Fama, E. & French, K., 2015, A Five-Factor Asset Pricing Model, Journal of Financial Economics, 116, 1-22.
- Fama, E. & MacBeth, J., 1973, Risk, Return and Equilibrium: Empirical Tests, The Journal of Political Economy 81.3, 607-636.
- Huang, Allen H., Hui Wang, and Yi Yang, 2022, FinBERT: A Large Language Model for Extracting Information from Financial Text, Contemporary Accounting Research.
- Laurer, M., van Atteveldt, W., Salleras Casas, A. & Welbers, K., 2022, Less Annotating, More Classifying – Addressing the Data Scarcity Issue of Supervised Machine Learning with Deep Transfer Learning and BERT – NLI. Preprint, June. Open Science Framework. https://osf.io/74b8k.
- Radford, A., Narasimham, K., Salimans, T. & Sutskever, I., 2018, Improving Language Understanding by Generative Pre-Training. https://openai.com/research/language-unsupervised

Notes

1 The author would like to thank Alex Ward Corral, Carmen van Wuijckhuijse, Daryl Smith, Finn van Cappelle, Giulia Mantovani, Judit van der Geest, Roy Hoevenaars, Sandra Toften and three reviewers of the VBA Journaal for their valuable comments and ideas.

2 This news article provides an example that illustrates that generative AI is not trained to generate factual correct text https://www.nytimes.com/2023/06/08/nyregion/lawyer-chatgpt-sanctions.html

3 Huggingface is the place to find pre-trained LLMs, https://huggingface.co/models
Models that relate to this article include yiyanghkust/finbert-tone, MoritzLaurer/DeBERTa-v3-base-mnli-fever-anli, Microsoft/deberta-v3-large, ProsusAI/finbert and nbroad/ESG-BERT

4 Several organizations are taking steps to defend themselves against copyright infringement by ChatGPT https://edition.cnn.com/2023/08/28/media/media-companies-blocking-chatgpt-reliable-sources/index.html

5 A prompt-engineer is someone who specializes in structuring questions or instructions for a generative AI model to obtain the best results.

6 The model is refreshed every three years: the forecasts as of July 2013 are based on a refresh with data from 1994 until 2012, the forecasts as of July 2016 are based on a refresh with data from 1994 until 2015, and so on.

7 The choice to split the stocks in the 1,000 largest and the 2,000 smaller stocks is the result of balancing between the desire to have both groups represent an equal market cap (requiring fewer large stocks) and being able to construct well diversified quintile portfolios (requiring more large stocks)

8 The rule that is used to adjust risk is to multiply the portfolio positions by the square root of the number of stocks in the portfolio and divide by the average of the square root of the number of stocks in the portfolio over the entire test period.

9 The TRBC sector definitions are used. For practical reasons, 4 sectors with a large enough number of stocks were selected: Industrials, Consumer Cyclicals, Health Care and Technology.

AI: Augmenting investment insights with Simona Paravani-Mellinghoff

By Sander Nooij

We had the opportunity to engage in a discussion with Simona Paravani-Mellinghoff, Global Chief Investment Officer of Solutions at BlackRock Multi-Asset Strategies & Solutions, and Assistant Professor at the Faculty of Economics at Cambridge University where she teaches a course in AI applied to Finance. She spoke with us from BlackRock's London office, delving into the integration of AI throughout investment management and its impact on the investment landscape.[1]

Our conversation encompassed the potential of AI in generating investment opportunities and enhancing investment management tools. Simona Paravani-Mellinghoff articulated the catalysts behind the current surge in AI despite its longstanding presence in academia. She offered insights in how AI drives productivity enhancements and refines investment tools by analyzing alternative data and automating various tasks. The dialogue concluded with an exploration of the alignment between academia and practitioners, shedding light on the pivotal skills for the growth of financial professionals.

AI Landscape – History, Progress, and Public Perception

With all the media frenzy on AI, there is a lot of confusion. Public understanding is often distorted by sci-fi stereotypes or unrealistic expectations of human-like "general" AI. What is AI really and what is changing?

The origins of AI as an academic discipline trace back to pioneering work in the 1950s. As an example of early AI, she highlighted a famous Bell Labs experiment that trained a neural network to navigate a maze.[2]

Paravani-Mellinghoff broadly defines AI as systems designed to mimic human cognitive capabilities, whether winning chess matches, analyzing MRI scans or composing poetry. However, specialized, narrow AI that focuses on specific tasks are likely more viable than "general" AI, i.e. the AI that can do anything a human can.

Within AI, she distinguishes a sub-field called generative AI that creates content like text, images, audio and video, rather than analyzing existing data. Generative AI has powered recent advances like chatbots and auto-generated images. Simona emphasizes that practical applications of AI only recently became possible through a combination of effects:

1. Increases in computational power;
2. The ongoing increase in data availability. As an example, she notes estimates that just two autonomous vehicles generate an amount of data equivalent to approximately 8,000 internet users; and
3. Increase in adoption of AI tools. Which in turn follows from algorithmic advances like transformer architectures that allow

Simona Paravani-Mellinghoff is the Global CIO of Multi-Asset Strategies & Solutions (MASS) at BlackRock. She has held senior positions at HSBC and was previously a quantitative strategist at Julius Baer Asset Management. She started her career as a quantitative analyst at Orbis in 1998. Paravani-Mellinghoff is an Assistant Professor at Cambridge University where she teaches Financial Analytics and Machine Learning. She has received industry awards including Role Model and Investment Woman of the Year in 2022. Paravani-Mellinghoff sits on the board of the financial education charity MyBnk and was awarded the Italian title of Commendatore OMRI for her professional achievements and commitment to education.

modeling language context, not just individual words. This significantly improves the fluency and coherence of generated text.

Paravani-Mellinghoff notes these trends enable models to mimic a broader range of human cognitive capabilities with higher fidelity. Whereas AI previously focused narrowly on logical tasks like chess, now the generation of prose, art and media are capturing the public's attention.
Importantly, she argues growth in data and computing shows no signs of slowing. The proliferation of smartphones, digital lifestyles, IoT sensors and more will continue to generate trainable data. The availability of custom AI chips and cloud infrastructure make leveraging this data at scale even more accessible.
Likewise, Paravani-Mellinghoff suggests faster public adoption will further drive generative AI's progress and potential applications. Younger and emerging market country demographics tend to more readily embrace AI.[3] Capabilities such as speech and image generation resonate universally.

ARTIFICIAL INTELLIGENCE EMERGES AS A PRODUCTIVITY ENHANCING MEASURE

Given the sustained trajectory, Paravani-Mellinghoff expects interest and development in generative AI models to continue accelerating.[4]

AI's Investment Implications – Opportunities and Tools

When examining AI's impact on investing, Paravani-Mellinghoff divides developments into two categories:
- Investment Opportunities: At a macro level, AI is transforming sectors and companies, creating winners and losers. This has investment implications as productivity and growth shift.
- Investment Tools: Within asset management firms, AI is augmenting processes ranging from data analysis to operations.

AI is generating Investment Opportunities

When examining AI's investment opportunities, Paravani-Mellinghoff outlines two key framework pillars to assess where opportunities may emerge:
- AI is not an island – It should be evaluated within the context of a broader technology ecosystem including innovations like blockchain, robotics and 3D printing. The interplay between AI and these other technologies will shape overall economic impact.
- Adoption matters as much as innovation – The benefits depend on companies effectively incorporating AI into business processes and models, not just disruptive startups. Leaders will combine AI with other emerging tech and upskill their workforces.

Elaborating on the first pillar, Paravani-Mellinghoff notes AI's productivity impacts and investment prospects are mediated through the technology ecosystem. For instance, AI, sensors and blockchain may intersect to transform supply chains. This magnifies AI's benefit relative to assessing it in isolation.
Regarding the second pillar, she cautions that investors overly focus on AI startups as the primary beneficiaries. She suggests that established companies which flexibly integrate AI across their operations, despite legacy constraints, may also make significant gains. How they execute their strategies to absorb AI into their business matters more than innovation alone.
As an example, she cited an MIT study showing AI augmented workers improved productivity on analytical tasks by up to 40% while also enhancing quality, especially for poorer performers.[5] This suggests that competitiveness of corporations will partly be driven by the optimal use of AI.
This framework suggests assessing investment potential through a wider technology lens, while evaluating the nitty-gritty of adoption and implementation, not just high-level disruption narratives. Companies which successfully combine AI with related innovations and real-world integration will drive economic shifts and create investment opportunities.

AI is improving Investment Management Tools

Beyond broader investment opportunities, Paravani-Mellinghoff explains AI is already transforming tools and processes within asset management firms:
- Data synthesis – AI techniques help condense vast amounts of alternative data sources like satellites, credit card transactions, etc. into usable insights. This expands the scope for analysis and represents one of the most significant changes that AI is bringing forward.
- Sentiment analysis – The emergence of large language models (LLMs) have made sentiment analysis better and more accurate versus simpler existing natural language processing (NLP) models used in the past such as 'bag-of-words.' These newer models can parse textual data like news, speeches and transcripts to generate sharper signals on market sentiment shifts. Paravani-Mellinghoff explained how NLP algorithms parse text to generate sentiment dashboards, potentially driving tactical asset allocations, converting qualitative data into quantitative signals. However, she stressed these AI tools should act as "co-pilots", augmenting rather than replacing traditional financial analysis and human judgment.

AI IS NOT AN ISLAND – THE INTERPLAY WITH OTHER TECHNOLOGIES WILL SHAPE OVERALL ECONOMIC IMPACT

- New data sources – She emphasizes the tremendous potential of AI techniques to synthesize vast amounts of alternative data beyond just textual sources. The ability to condense and extract signals from the explosion of alternative data sources like satellites, credit card transactions, social media feeds, etc. represents one of the most significant transformations that AI enables. By leveraging AI to turn these new and exponentially growing alternative data sources into usable insights, the breadth of information available for investment decisions can expand dramatically.[6]

- Idea generation – By scanning large datasets, AI can identify non-obvious relationships and clusters around factors like momentum, value, quality, etc. This allows fresh approaches to multifactor and smart beta strategies. Investment teams can then explore these algorithmically generated ideas further.
- Operations – Paravani-Mellinghoff explains AI is transforming a wide range of back-office tasks including automated report writing through natural language processing, reconciling complex datasets via machine learning, flagging suspicious trading patterns using anomaly detection algorithms, extracting key details from legal documents and regulations via text analysis. This ties in to the improved quality of the work done by humans aided by AI mentioned earlier.

INVESTMENT OPPORTUNITIES DEPEND ON THE INTERPLAY BETWEEN AI AND OTHER TECHNOLOGIES AS WELL AS THE RATE OF ADOPTION

Paravani-Mellinghoff emphasized the role of synthesizing alternative data. She explains that AI techniques can help condense and extract signals from the vast and growing array of alternative data sources mentioned above. Though AI unlocks these new data sources, human critical thinking remains vital to separate signal from noise and thus AI should augment, not replace, human portfolio managers' expertise and judgement.
Looking ahead, she expects AI's application in investment processes to continue expanding.

Connecting Academic Research and Business Applications

When asked about the relationship between academia and current business practice, Simona noted that large language models like GPT-3 are still new for both academia and industry. Her position in academia allowed her to interface with these models earlier than industry, giving her some advance insight into their potential. While she had earlier access to such models, she swears not to have used them to write her published children's books.
Regarding the skills and talent needed going forward, Simona emphasized the complexity of models makes it harder to rely on standard approaches to verify outputs. She noted that as an academic, in her opinion, there are two main skillsets which are needed.
First, critical thinking, because the more complex the models, the less you can rely on standard approaches to verify outputs. The ability to determine if something makes sense will become even more crucial. Adoption of AI by businesses matters more than innovation alone
Second, the ability to work well with others, since no one person may have all the skills to critically evaluate complex model outputs on their own. Gaining others' perspectives on machine output will be important.
In response to an article suggesting interpersonal skills will become more important than traditional quant skills, Simona didn't see it as an either/or proposition. She believes critical thinking requires an understanding of probability and statistics, traditionally quant skills. But that it is also about combining different perspectives to analyze challenges.
Simona pointed to the recent 100+ page Executive Order on AI from the Biden administration,[7] highlighting its focus on training and skills development. She believes this underlines the need for education and training to focus on evaluating model outputs regardless of job function rather than just teaching coding skills.
We concluded the interview with one final question. We asked her what next breakthroughs we could expect next with her one foot in academia. Unfortunately, she wanted to keep that a secret for now.

Notes

1 Artificial intelligence – beyond the buzz, A roadmap to help assess the investment implications of artificial intelligence, https://www.blackrock.com/corporate/literature/whitepaper/bii-megaforces-november-2023.pdf
2 Claude Shannon, known as the "father of information theory", demonstrating a mechanical mouse he built that could learn to navigate a maze. By adjusting electrical circuits and relay switches, Shannon was able to train the mouse to successfully traverse the maze on its own through a form of primitive machine learning. https://www.youtube.com/watch?v=vPKkXibQXGA
3 Zhang, Baobao and Dafoe, Allan, Artificial Intelligence: American Attitudes and Trends (January 9, 2019). Available at SSRN: https://ssrn.com/abstract=3312874 or http://dx.doi.org/10.2139/ssrn.3312874
4 https://www.weforum.org/agenda/2021/07/this-is-a-visualization-of-the-history-of-innovation-cycles/
5 Study finds ChatGPT boosts worker productivity for some writing tasks, Zach Winn | MIT News Office, July 14, 2023, https://news.mit.edu/2023/study-finds-chatgpt-boosts-worker-productivity-writing-0714
6 Speaking at the *Environmental Finance* Future of ESG Data event, BlackRock's Head of Factors, Sustainable and Solutions Andrew Ang said the $237 billion data-driven investment unit of BlackRock was "underweight" Silicon Valley Bank after analyzing customer complaints data provided by the US Consumer Finance Protection Bureau.
7 FACT SHEET: President Biden Issues Executive Order on Safe, Secure, and Trustworthy Artificial Intelligence https://www.whitehouse.gov/briefing-room/statements-releases/2023/10/30/fact-sheet-president-biden-issues-executive-order-on-safe-secure-and-trustworthy-artificial-intelligence/ and https://www.whitehouse.gov/briefing-room/presidential-actions/2023/10/30/executive-order-on-the-safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence/

Artificial Intelligence: Do the Advantages Outweigh the Risk?

Artificial Intelligence (AI) and its profound manifestation in ChatGPT are rapidly transforming the financial landscape, providing financially savvy investors with a trove of opportunities while introducing a set of unique risks and challenges. The integration of AI in finance is not just a fleeting trend; it is a seismic shift reshaping how decisions are made, risks are managed, and opportunities are seized.

AI-driven risk management tools are empowering investors to navigate the complex financial markets more effectively. By monitoring market sentiment, detecting anomalies, and suggesting timely portfolio adjustments, these tools enhance diversification and optimize risk-reward profiles. Investors are also leveraging AI for algorithmic trading, tapping into high-frequency trading, arbitrage opportunities, and quantitative analysis to gain a competitive edge. AI's prowess in predictive analytics is proving invaluable, offering insights into market movements, and helping investors make data-driven decisions. Enhanced due diligence capabilities enable the analysis of vast volumes of financial reports, news articles, and social media data, uncovering investment opportunities that might elude human analysts.

The financial landscape has witnessed AI-powered hedge funds like Renaissance Technologies consistently outperforming their traditional counterparts, showcasing the technology's potential. The rise of decentralized finance (DeFi) platforms, propelled by smart contracts and AI-driven algorithms, is challenging the status quo of banking and investments, providing novel opportunities for lending, borrowing, and yield farming in a decentralized ecosystem. However, the GameStop saga of early 2021 serves as a stark reminder of AI's double-edged sword, as trading algorithms fueled market volatility and led to unpredictable outcomes.

Algorithmic trading, while offering unparalleled speed and efficiency, has introduced risks of "flash crashes" and increased market fragility, as evidenced during the 2020 pandemic-induced market turbulence. Data privacy and security concerns loom large, with AI systems handling massive datasets and potentially sensitive information. Investors must be vigilant and ensure that AI platforms comply with stringent privacy standards amid increasing calls for transparency and data protection. The danger of overreliance on AI is palpable, as investors who depend solely on algorithms risk being blindsided by unprecedented events. Ethical concerns also abound, with the potential for AI to manipulate public perception and stock prices through disinformation.

In the fast-evolving world of AI-driven finance, investors are presented with a paradox: a wealth of opportunities juxtaposed with significant risks. The question is not whether the advantages of AI outweigh the risks, but rather how investors can judiciously harness AI's potential while mitigating its dangers. AI should be viewed as a powerful tool, not a cure-all, requiring informed decision-making, prudent risk management, and a steadfast commitment to ethical practices. In this transformative era, those who adapt wisely, balancing the promise and perils of AI, will be best positioned to thrive, safeguarding their investments, and capitalizing on the opportunities at hand. As we navigate this complex and interconnected market, the onus is on us to ensure that the advantages of AI in finance truly outweigh the risks.

Loranne van Lieshout,
committee Risk Management

We took a different approach to the column for the winter issue of the Risk Committee this time. Given the theme Artificial Intelligence and ChatGPT, we felt it was fitting to have the column written by ChatGPT this time. The column you see here is also fully written by ChatGPT.

What is noticeable is that not only is the writing style visibly from ChatGPT, but there are also passages where the accuracy can be seriously doubted. The examples given seem plausible, but when you examine the subject matter better or have more knowledge of it, it turns out to be much more nuanced than the column suggests. This is especially apparent in paragraphs where the subject matter gets more complex.

It is therefore important not to read this column as the opinion of the Risk Committee, but as an example of how uncritically adopting utterances from ChatGPT can lead to misinformation. It thus remains important to remain critical of generated texts from ChatGPT and similar technologies.

The current state of AI for investment management

Mike Chen, Iman Honarvar and Harald Lohre

When OpenAI released ChatGPT in November 2022 it created a wave of excitement around artificial intelligence (AI) among the general public, media, and across industries worldwide. Such generative AI models (also called generative large language models or LLM) have sparked a new race in the tech world regarding who has the best AI offering. On one hand, the Microsoft-OpenAI partnership was leading the way with incorporating generative AI into many applications, while other big firms like Google were stepping up with their own AI tech in the form of BARD. More recently, open-source versions of generative AI such as Meta's Llama are surging in popularity, because they enable end-users to finetune generative AIs for their own specific applications without the massive cost and data required to train a generative AI from scratch.

The excitement around generative AI is rooted in the vast opportunities they present to transform various aspects of our personal and professional spheres. Not only is generative AI seen as a new companion for casual conversations, but also as a versatile tool capable of authoring poems, offering legal opinions, brainstorming marketing strategies, and even selecting the best investment stocks (Lopez-Lira and Tang, 2023). Like in other industries, finance practitioners are fascinated by the possibilities. While some finance practitioners are pioneers in the use of AI, the vast majority of the industry is either in the early stages of exploring its potential or has recently begun integrating it into their workflows.

In this article, we touch on salient advances in artificial intelligence and its applications in investment management, highlighting the opportunities it offers across a broad spectrum

Harald Lohre (l)
Executive Director of Research at Robeco, Quantitative Investments division in Rotterdam, The Netherlands, and an Honorary Researcher at Lancaster University Management School in Lancaster, UK

Iman Honarvar (m)
Director of Research at Robeco, Quantitative Investments division in Rotterdam, The Netherlands

Mike Chen (r)
Head of Alternative Alpha at Robeco, Quantitative Investments division in Rotterdam, The Netherlands

of use cases. For example, AI can be used to uncover non-obvious relationships in massive quantities of data or to detect sentiments expressed in textual or even audio data. We start with a discussion on the key developments in AI over the years and highlight how the finance industry have been incorporating these advances. We then discuss three specific sub-branches of AI: machine learning (ML), natural language processing (NLP), and generative AI.

A BRIEF HISTORY OF AI AND THE CURRENT LANDSCAPE

Researchers have relied on fundamental principles and mathematical derivations throughout history to make predictions and decisions. But the increased complexity of tasks today calls for new methods and algorithms to identify more intricate relationships between many variables. To this end, more advanced systems and algorithms have been developed to conduct heavy data-processing tasks and make human-like intelligent decisions. The umbrella name for all such systems is artificial intelligence, the ultimate goal of which is a system indistinguishable from human intelligence, the so-called artificial general intelligence (AGI).

While early AI systems were unable to learn and adapt as new data came in, in time they flourished, developing into various subfields such as machine learning (ML) and natural language processing (NLP) (see Figure 1). For instance, in ML, problems are solved using algorithms that can adapt, learn, and discover their own rules and relationships, while in NLP computers are given the ability to quantify and comprehend textual and/or speech data.

To flexibly learn data patterns and nuances, ML and NLP algorithms often use deep learning models which are neural networks with multiple layers. One of the most prominent recent innovations in artificial intelligence are generative AI models, wherein deep neural networks learn the nature of the input training data to then generate new data that has similar characteristics. Well-known examples of generative AI models include GPT4 for textual applications and DALL-E for image applications.

Figure 1
Schematic AI landscape,

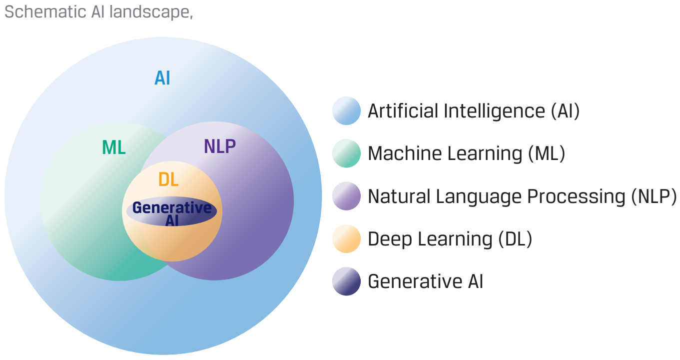

Source: Robeco

In field after field, artificial intelligence techniques have made their way into various tasks such as facial recognition, internet searches, medicine, autonomous driving, or playing chess and Go. In this last example and indeed many other fields, machines have become fully autonomous and outperformed their human contenders.

MACHINE LEARNING IN FINANCE

While ML excels at many complex tasks, investment brings a whole new level of difficulty because it relates to predictions and decisions on the economy and financial markets. These are influenced by millions of human beings who are in turn influenced by a myriad of variables. The finance industry began adopting elements of AI as early as the 1990s, with varying degrees of success. For instance, consumer credit default probability and fraud detection models were among the first to use ML techniques. Conversely, the practical use of ML for stock selection and market timing beyond simple quantitative algorithms has only taken off in the last decade.

There are at least three reasons why investment management applications have adopted such techniques only recently. First, ML models require a massive amount of training data, which is more abundant in some fields than others. For example, two computerized chess players can play endlessly against each other, thus generating unlimited training data.[1] The amount of data in investment applications tends to be more limited. Take the example of predicting asset returns during different macro-economic periods. The most commonly used dataset in asset pricing, the US Stock Database of CRSP,[2] has less than one hundred years of sample data and only covers US companies and sixteen recession periods (as defined by the NBER[3]). Predicting recessions based on sixteen unique observations is inherently more challenging than predicting the best next move in chess.

Secondly, finance is dynamic and adaptive unlike fields bound by rigid rules. For instance, pawns can only move forward in chess, but in investing, market participants have a larger opportunity set but less information. Unlike chess or Go, where all pieces are visible, investors are typically unaware what other market participants hold. This opacity, coupled with the diverse reactions of investors to similar decisions at different times, adds complexity. Investors are heterogeneous. Moreover, the unpredictability of events like the sudden abandoning of the gold standard in the US or the COVID-19 pandemic make it hard for ML to add value in investment applications.

Thirdly, stock price patterns are not obvious. Otherwise, they would be quickly arbitraged away by smart market participants. Because financial asset returns result from aggregate inputs of thousands of investors each incorporating a plethora of considerations, from macro to micro and from sentiment to positioning and liquidity needs, they are more driven by human behavior, biases, and heuristics. Put another way, the signal-to-noise ratio in finance is famously low compared to other fields in which ML has excelled as a result.

Nevertheless, there are many instances where machines can help humans in investment decisions, owing to the observation that "history never repeats itself, but it often rhymes." While ML models do not have domain knowledge per se, the practitioners developing these algorithms do. As a result, ML models can be directed to learn the relevant features by going through many examples in the training data. These algorithms generally become more robust and accurate as machines are fed with more data, ideally arriving with a higher signal-to-noise ratio.

This highlights the benefits and shortcomings of ML models. On the one hand, while humans might be swayed by behavioral biases, machines can see the data as it is. Take the unduly high valuation of internet companies during the dot-com bubble as an example: while a human investor might have overreacted to the bright future of the internet, a proper ML model would have been less likely to have rationalized these company valuations. On the other hand, when predicting an event, machines are generally mostly successful for events with similar precedence in the training set. When the event is completely new, machines are likely to struggle. For example, ML models could barely predict the stock market swings in 2020, as they had not seen a widespread global pandemic before.

Hence, although a machine can execute a task faster than a human and bring sensible second opinions, in finance, it is not about 'man versus machine' . Instead, "AI power and human wisdom are complementary in generating accurate forecasts and mitigating extreme errors, portraying a future of 'Man + Machine' (Cao et al, 2021). Machine-learned patterns are useful only if they can be rationalized by economic intuition, something that is arrived at by human judgment.

INVESTMENT MANAGEMENT APPLICATIONS OF MACHINE LEARNING

In this section, we illustrate a few concrete examples of how machine learning can benefit the practice of investment management, focusing on risk and return prediction. Over the past few years, academic research has investigated the efficacy of ML in stock price prediction (see, for example, Gu et al. 2020, Hanauer and Kalsbach 2023) or bond price prediction (see, for example, Bianchi et al. 2021). Machine learning models can identify complex patterns that might easily go unnoticed by the human eye, especially regarding asymmetries and interaction effects.

Short-term price reversal: A robust pattern in stock prices is the short-term price reversal effect, which was discovered in the 1980s. When stocks experience a large negative price move over a short period (a couple of days), this price move tends to partially reverse over the next few days. Indeed, short-term price reversal is often the most prominent feature in ML models that predict the cross-section of equity returns, see Gu et al. (2020), Leung et al. (2021), or Blitz et al. (2023a). Furthermore, ML models often find that short-term price reversal interacts with other features.[4]

For instance, short-term reversal is much weaker after news or earnings announcements are released. Such machine-learned patterns can be economically rationalized, as price reversal happens mostly after a large market participant liquidates a large number of shares and moves the firm's stock price away from its fair value (abstracting from any change to the firm's fundamentals). However, when a big price movement is driven by news or earnings announcements , the price does not revert because the fundamental stock value has shifted to a new level.[5] Based on Chen's (2023) analysis, Figure 2 depicts how this interaction of returns and number of news items in the previous ten days can be leveraged to forecast future stock returns.

Distress prediction: While the ML prediction of stock prices suffers from a low signal-to-noise ratio, this ratio is more benign when it comes to risk and distress prediction. For example, we intuitively assume that increasing the cost of debt can have a negative impact on a stock's riskiness. However, the ML model of Swinkels and Hoogteijling (2022) reveals that this effect is not linear and homogenous; increasing cost of debt is especially detrimental for heavily indebted companies with shorter distance-to-default. Overall, the authors demonstrate how modelling such non-linear patterns together with interaction effects can benefit the prediction of stock price crash risk. Figure 3, as presented in Swinkels and Hoogteijling's (2022) analysis, demonstrates that the prediction of stock returns involves non-

Figure 2
Stock price pattern conditioned on the presence of recent news

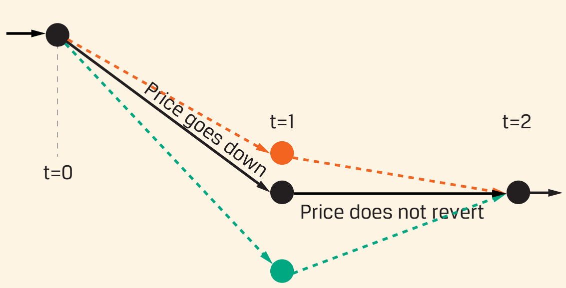

Predicted return over the next 10 days

Return over past 10 days	News count over the past 10 days: =-3	=-2	=-1	=0	=1	=2	=3
=-3	16%	13%	10%	8%	5%	2%	-1%
=-2	11%	9%	7%	5%	3%	1%	-1%
=-1	6%	5%	4%	3%	1%	0%	-1%
=0	1%	0%	0%	0%	0%	0%	-1%
=1	-4%	-4%	-3%	-3%	-2%	-1%	-1%
=2	-10%	-8%	-7%	-5%	-3%	-2%	0%
=3	-15%	-12%	-10%	-8%	-5%	-3%	0%

Source: Chen (2023)

Figure 3
Non-linearity and interaction effects, detected by ML models, for stock distress prediction

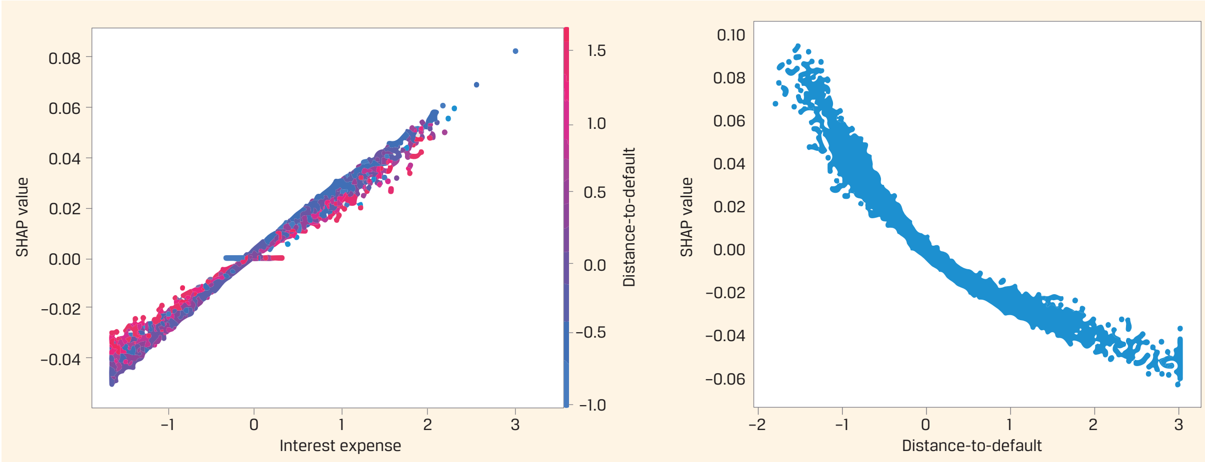

Source: Swinkels and Hoogteijling (2022)

linearity and interaction effects, for example with respect to distance to default and interest expenses.

Volatility surfaces: As machine learning leaps forward in other fields, its financial applications follow suit. For example, image recognition is an everyday use case of machine learning, wherein a machine is tasked to identify the main objects in an image. The first layers of the neural network that are used in these models are often convolutional ones that identify the primitive shapes in the image, such as lines, circles, and squares. In the later layers, these primitive shapes are combined to construct more complex patterns, such as a human portrait, with a circle for the face, two small circles for the eyes, two lines for the eyebrows, and so on. While this use case might seem farfetched, it can readily be translated into a finance application. Kelly et al. (2023) show how the volatility surface can be seen as an image, with pixel colors as the implied volatility of the options at different moneyness levels and time-to-maturity (Figure 4), and then use convolutional layers to predict stock returns with the data embedded in this image. In this setup, the machine learns the relationship between the option-implied information and the future returns of the corresponding security.

Company clustering and theme identification: All of the above are examples of supervised learning, where the target (either stock return or risk) is clear. However, ML techniques are also well established for unsupervised learning, such as company clustering, where the aim is to form clusters of similar companies based on a large number of characteristics. These clusters can then be used to identify emerging market trends and themes, or for peer group comparison. Conventionally, similar companies are grouped based on their sector or industry classifications. However, these classifications are often too crude, bucketing companies in some groups without considering the closeness of companies in two distinct buckets.

Figure 4
Volatility surface of NVIDIA on October 9th, 2023

The Convolutional Neural Networks allow for the extraction of valuable insights from the options volatility surface, enabling predictions on the risk and return associated with stocks.

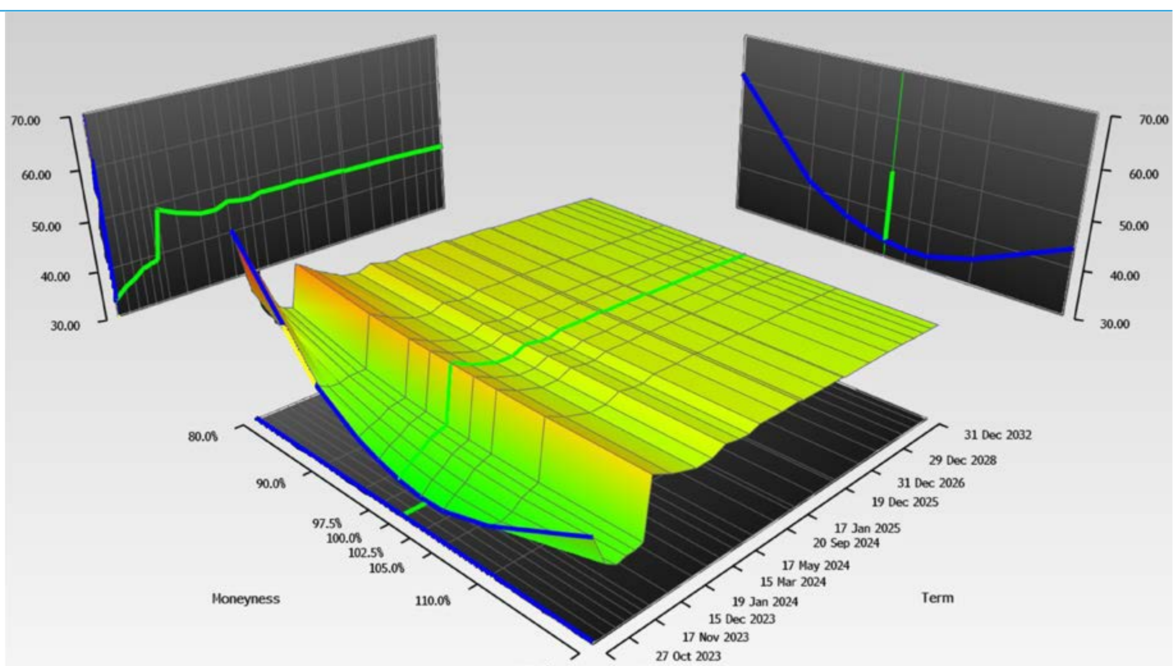

Source: Bloomberg

For example, consider a lithium mining company and an electric vehicle producer. Although such companies are classified in two distinct sectors (mining and consumer discretionary), they have exposures to similar risk factors. In order to improve stock peer groups in this way, one can use Principal Component Analysis (PCA) or Uniform Manifold Approximation and Projection (UMAP) to reduce the data dimensionality, and subsequently use clustering techniques such as K-Nearest-Neighbors (KNN) or Density-based Clustering (DBSCAN) algorithms to group the companies based on their similarity. Figure 5 displays clusters of companies based on the similarity of company descriptions.

Trading: Once the suitable stock has been identified, it becomes crucial to employ the most effective trading strategy. For example, revealing the intent to trade a substantial volume of securities can inadvertently trigger significant price fluctuations, often detrimental to one's objectives. Thus, traders or brokers frequently utilize advanced machine learning models to ascertain the most beneficial trading approach, such as determining the optimal participation rate, the best trading venue, etc.

The availability of more independent observations at intraday frequencies allows machines to be furnished with a larger volume of data. This, in turn, enhances their ability to identify robust patterns in the data, rendering trading as one of the most promising domains for the application of AI in finance.[6] However, the application of ML in trading is not confined to optimal trade execution. It can also be extended to other significant use cases, such as broker selection. For instance, a study conducted by ter Braak and van der Schans (2023) demonstrated how reinforcement learning algorithms can be employed to identify the most effective broker execution strategies.

THE VIRTUE OF NATURAL LANGUAGE PROCESSING

Along with the development of ML models, natural language processing (NLP) models have made great strides in performing textual analysis. The first NLP models such as bag-of-words were statistical and only based on the coappearance and frequency of words in a text. These primitive models treat each word as an independent observation. For example, the words "car" and "vehicle" are treated independently, although these two words are similar.

These models cannot gain much understanding of the meaning and similarity of two distinct words, and they cannot provide a deep understanding of the context and the relation of different words and sentences in a document. To overcome these shortcomings, deep learning models have been developed that can embed each word within a vector of numbers. Words with more similar meanings would have a smaller cosine distance between their embedding vectors, and one can even perform simple calculations with these vectors, such as "King – Man + Woman = Queen."

The invention of attention layers and transformer models brought another giant leap to natural language processing

Figure 5
Clusters of companies based on company descriptions

In this visualization, every small dot on the graph represents a company, distinguished by different colors based on the GICS (Global Industry Classification Standard) sector they belong to. The dots that are close together indicate groups of companies with similar business descriptions or lines. The larger circles on the graph indicate the top 50 companies that are closely related to the field of "Electric Vehicles". This graph reveals that the companies involved with "Electric Vehicles" are spread across different GICS sectors including Industrials (light green), Utilities (dark green), Consumer Discretionary (yellow), Information Technology (purple), and Materials (light blue).

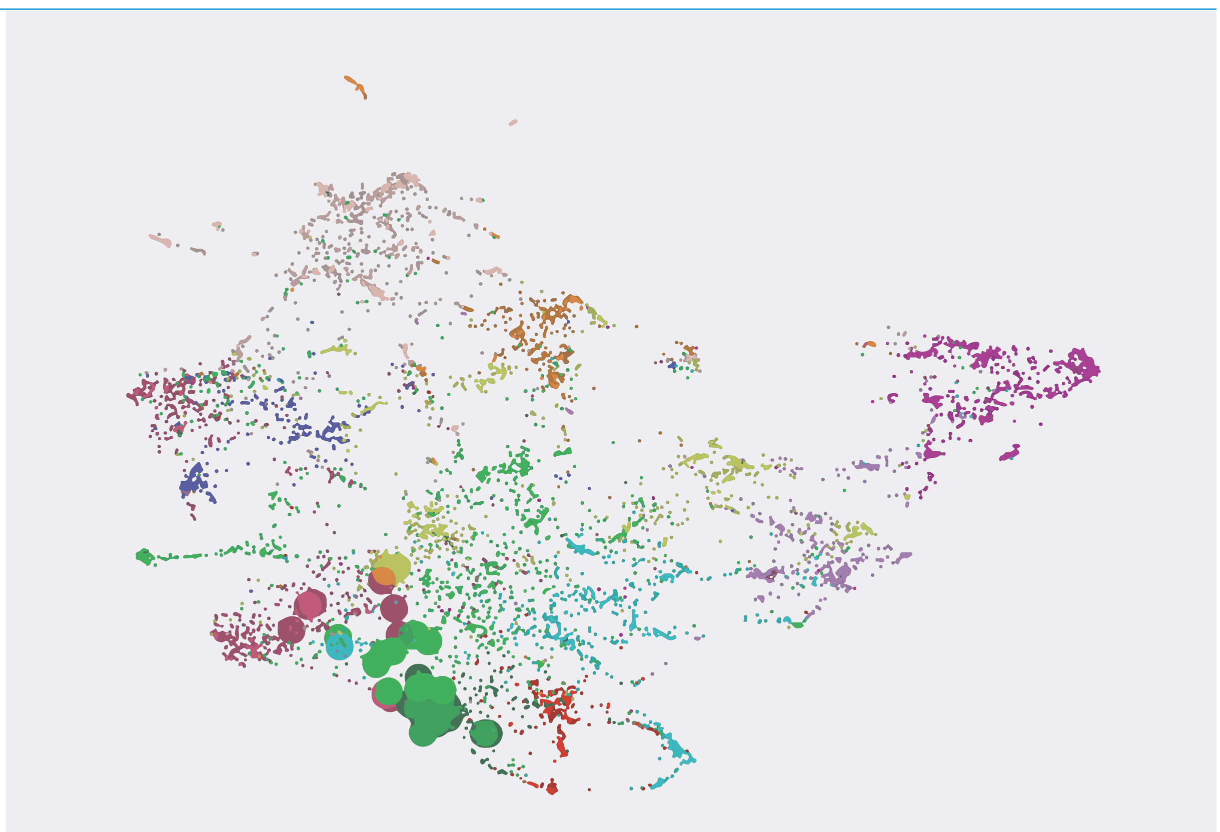

Source: Robeco

models, see Vaswani et al. (2017). Such models can learn the word's context within a context window (for example, within the sentence or paragraph), providing a much deeper textual understanding through the machine. Below, we highlight how NLP models can be applied to sentiment analysis as well as help advance sustainability investing.

Sentiment Analysis: The abundance of financial textual data makes it a unique hunting ground for testing NLP models. For instance, sentiment analysis helps to gauge the tone of management speech or company news. A simple way to identify text sentiment is to measure the proportion of positive and negative words. However, this naïve approach can be improved upon by leveraging more recent NLP developments as presented in context-aware models tailored for financial texts such as FinBERT (2020).

Indeed, Heston and Sinha (2018) show that news sentiment can help predict stock returns at different horizons, and Jiang et al. (2019) show that higher manager sentiment (extracted from company 10-K fillings, 10-Q fillings, and conference call transcripts) precedes lower aggregate earnings surprises and greater aggregate investment growth.

Sustainability investing: Sustainability is a fruitful area of AI application. For instance, in the past, company ESG scores were primarily based on self-declared surveys filled by the company management. Today, company news (media perception), earnings calls, financial disclosures (management perception), or employee reviews (employees' perception) can be rich data sources for company profiling. For example, one could score companies on ESG aspects based on how the media perceives them in each dimension. Similarly, Amel-Zadeh et al. (2021) apply NLP models on Corporate Social Responsibility (CSR) reports of Russell 1000 companies to measure their alignment with UN Sustainable Development Goals.

To perform specific sustainability investing tasks, LLMs have also been trained and fine-tuned to perform specific sustainability investing tasks. For instance, ClimateBERT, ESGBERT or ControversyBERT can aid in investigating a company's climate disclosure to spot greenwashing or identify companies linked to controversial behavior, see Webersinke et al. (2023), Schimanski et al. (2023) and Lohre et al. (2023).

THE RISE OF GENERATIVE AI

The most prominent recent innovations in artificial intelligence are generative AI models, particularly GPT models in natural language processing. These models are trained on vast text datasets and can interact with the user, answer questions, edit given documents, and assist in ideation by generating sensible sentences and paragraphs. As such, they can be used in the automation of tasks, not just in finance but in many industries. They can thus reduce the burden of repetitive labor-intensive tasks, such as drafting regular reports, requests for proposals, or emails. It is important to caution that generative AI models (e.g., ChatGPT) are not flawless and require interactive supervision by a human. In other words, humans function as domain experts and validate rather than create the content.

But the use cases of GPT are not limited to repetitive tasks. With the release of generative models tuned for programming, such as Github CoPilot, one can enjoy the machine as a sparring partner in coding functions or when translating them from one programming language into another one. GPT models can assist lawyers by analyzing large volumes of legal documents, case laws, and regulations and providing a second opinion. In customer relationship management, GPT models can power chatbots to manage initial customer inquiries, provide automated responses, and offer basic troubleshooting guidance. They can understand and respond to customer queries based on pre-trained knowledge and historical data from the company.

Finally, GPT can function as a helpful research assistant in various ways. It can help generate ideas, conduct basic literature reviews, summarize academic papers, compare arguments between two papers, edit drafts, and answer basic questions like identifying peer companies of NVIDIA or determining which sectors are impacted by energy price inflation.

IMPLICATIONS AND CONCLUSIONS

The introduction of AI in finance is a significant advancement that can revolutionize various office tasks. While concerns about personnel redundancy and unemployment are certainly valid, history has shown that humans can adapt and find new avenues for growth and development alongside automation. For example, the rise of the steam engine replaced horse cart drivers, but created jobs such as train engineers, railroad builders, conductors, and so on.

The rise of computers replaced people that did manual computations, data entry, and record keeping, but created new jobs such as programmers, database engineers, or system architects. Similar to the way in which robots automated factory jobs in the past, AI has the potential to free up human resources to focus on higher value-added and creative activities, such as idea generation and business development. As we witness thousands of humans collaborating on NASA's "unmanned" Artemis program to explore the moon, it becomes evident that humans and machines can work together synergistically.

More specifically for finance, the introduction of AI can automate many of the repetitive and labor-intensive processes currently carried out by human practitioners. Examples include summarization of news and broker reports for consumption by portfolio managers and writing strategy performance reviews for asset owners to keep track of their investments. Other, more intellectually involved aspects of finance may also be speeded up or completely automated by AI, such as finding unique and causal relationships within massive amounts of data or deciding in which asset class to invest, depending on macro-economic data and market sentiments.

What about disruption posed by AI? For the finance industry, on a macro level, implementing artificial intelligence will allow the industry to further evolve. Fundamental investing will experience a particularly large disruption. Until recently, AI was mainly used by advanced quantitative investors. As AI tools become more user friendly, they will allow fundamental investors to acquire the capabilities previously only available to quantitative investors. With AI, fundamental investors will be able to systematically analyze large datasets, perform highly complex calculations, and so on.

However, this does not mean quant investors' traditional advantages will dissipate in the age of AI. Many AI programs are trained for general tasks. For finance specific applications, further trainings (called "fine tuning" in AI lingo) are necessary. Skilled quant investors will have the tools and know-how to customize general AI programs or develop completely custom AIs from the ground up. These specifically customized AI will have an advantage over the more generic AI employed by fundamental investors.

At the individual level, certain skills and jobs will disappear while others will evolve as AI rewrites the rules of the game. While certain human capabilities are not in danger of being replaced for the foreseeable future – think of empathy, creativity, and the creation of new knowledge – adopting and harnessing the possibilities of AI will be key for investment professionals to stay relevant. For today's practitioners, it is beneficial and perhaps even essential to become comfortable in working with AI, at the very least. Learn the basics of AI – there are many useful and easily accessible courses available – and think about which parts of one's daily work can be made better by incorporating AI into the process.

References

- Amel-Zadeh, A., Chen, M., Mussalli, G., & Weinberg, M. (2021). NLP for SDGs: Measuring corporate alignment with the sustainable development goals. Columbia Business School Research Paper.
- Bianchi, D., Büchner, M., & Tamoni, A. (2021). Bond risk premiums with machine learning. The Review of Financial Studies, 34(2), 1046-1089.
- Blitz, D., Hoogteijling, T., Lohre, H., & Messow, P. (2023a). How can machine learning advance quantitative asset management? Journal of Portfolio Management, Quantitative Tools 2023, 49(9), 78-95.
- Blitz, D., Hanauer, M. X., Hoogteijling, T., & Howard, C. (2023b). The Term Structure of Machine Learning Alpha. Journal of Financial Data Science. Forthcoming.
- Cao, S., Jiang, W., Wang, J. L., & Yang, B. (2021). From man vs. machine to man+ machine: The art and AI of stock analyses (No. w28800). National Bureau of Economic Research.
- Cartea, Á., Jaimungal, S., & Sánchez-Betancourt, L. (2021). Deep reinforcement learning for algorithmic trading. Available at SSRN 3812473.
- Chen, M. (2023). Eight things you should know about next-gen quant, Robeco white paper, https://www.robeco.com/en-int/insights/2023/08/8-things-you-should-know-about-next-gen-quant
- Dai, W., Medhat, M., Novy-Marx, R., & Rizova, S. (2023). Reversals and the returns to liquidity provision (No. w30917). National Bureau of Economic Research.
- Gu, S., Kelly, B., & Xiu, D. (2020). Empirical asset pricing via machine learning. The Review of Financial Studies, 33(5), 2223-2273.
- Hanauer, M. X., & Kalsbach, T. (2023). Machine learning and the cross-section of emerging market stock returns. Emerging Markets Review, 55, 101022.
- Heston, S. L., & Sinha, N. R. (2017). News vs. sentiment: Predicting stock returns from news stories. Financial Analysts Journal, 73(3), 67-83.
- Jiang, F., Lee, J., Martin, X., & Zhou, G. (2019). Manager sentiment and stock returns. Journal of Financial Economics, 132(1), 126-149.
- Kelly, B. T., Kuznetsov, B., Malamud, S., & Xu, T. A. (2023). Deep Learning from Implied Volatility Surfaces. Swiss Finance Institute Research Paper, (23-60).
- Leung, E., Lohre, H., Mischlich, D., Sheah, Y., & Stroh, M. (2021). The promises and pitfalls of machine learning for predicting cross-sectional stock returns. Journal of Financial Data Science, 3(2), 21-50.
- Lohre, H., Nolte, S., Ranganathan, A., Rother, C., & Steiner, M. (2023). ControversyBERT: Detecting social controversies and their impact on stock returns. Journal of Impact and ESG Investing, forthcoming.
- Lopez-Lira, A., & Tang, Y. (2023). Can ChatGPT forecast stock price movements? Return predictability and large language models. arXiv preprint arXiv:2304.07619.
- Ritter, G. (2017). Machine learning for trading. Available at SSRN 3015609.
- Schimanski, T., Reding, A., Reding, N., Bingler, J., Kraus, M. and Leippold, M. (2023). Bridging the Gap in ESG Measurement: Using NLP to Quantify Environmental, Social, and Governance Communication, Working Paper.
- Swinkels, L., Hoogteijling, T. (2022). Forecasting stock crash risk with machine learning, Robeco white paper, https://www.robeco.com/en-za/insights/2022/06/forecasting-stock-crash-risk-with-machine-learning
- ter Braak, L., & van der Schans, M. (2023). Optimal Order Routing with Reinforcement Learning. The Journal of Financial Data Science.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. (2017). Attention is all you need. Advances in Neural Information Processing Systems, 30.
- Webersinke, N., Kraus, M., Bingler, J. A., & Leippold, M. (2022). ClimateBert: A pretrained language model for climate-related text, arXiv preprint arXiv:2110.12010, AAAI Fall Symposium 2022 (Conference Proceedings).
- Yang, Y., Christopher, M., & Huang, A. (2020). FinBERT: A Pretrained Language Model for Financial Communications. Available at: https://github.com/yya518/FinBERT.

Notes

1 For example, AlphaZero AI model of DeepMind for learning and playing chess, Go and Shogi.

2 Center for Research in Security Prices, LLC, affiliate of Booth School of Business at University of Chicago.

3 National Bureau of Economic Research, an American private nonprofit research organization.

4 While short-term price reversal looks strong on paper, models that predominantly rely on short-term models exhibit high turnover and the associated transaction costs may reduce the effective value-add in actual investment portfolios, see Leung et al. (2021). Blitz et al. (2023b) further document that short-term price reversal is particularly important when the forecast horizon is short.

5 For more discussions on reversal and its interactions, see Dai et al. (2023).

6 See for example Ritter (2017) and Cartea, Jaimungal, Sánchez-Betancourt (2021).

De sombermansen en de optimisten

Tijdens een etentje met vrienden ontstond vorig jaar een verhitte discussie over ChatGPT, dat toen net gelanceerd was. Er waren twee kampen: de sombermansen die in ChatGPT de ondergang van het kritische denken zagen, de doodsteek voor de creativiteit, het einde van vele mooie beroepen. Aan de andere kant zaten de techoptimisten – waar ik mezelf toereken – die van een mooie toekomst droomden waarin ChatGPT de saaie routineklusjes overneemt waardoor er meer tijd overblijft voor het diepe en echt creatieve denkwerk. Het werk zou alleen maar leuker worden, de wereld mooier. De sombermansen lachten schamper over zoveel naïeve onnozelheid. De techoptimisten wezen er fijntjes op dat de grammofoon de musicus niet werkloos had gemaakt en de digitale camera de fotograaf niet overbodig.

Dit blijkt niet een heel unieke discussie te zijn.

De komst van een nieuwe technologie kent namelijk een vrij voorspelbaar reactiepatroon. Van 'grote beloften en grote angstbeelden' aldus de Wetenschappelijke Raad voor het Regeringsbeleid (WRR) die er een aantal rapporten over schreef. Ik citeer: *'Enerzijds circuleerden er utopische beloften van vrijheid, welvaart en modernisering, anderzijds angstbeelden dat de technologie de huidige maatschappij en manier van leven diepgrondig zou ontwrichten, bestaande ongelijkheid zou versterken en dat nieuwe partijen een ongebreidelde macht krijgen.'*

Zo heeft elektriciteit ons ontegenzeggelijk veel gebracht, maar destijds klonk er ook kritiek. Zo zagen sommigen elektrische straatverlichting als een gewiekste manier van de staat om zijn burgers beter in de gaten te kunnen houden. Elektrisch licht en klokken werden beschouwd als machtsmiddel om arbeiders beter te kunnen controleren en eronder te houden.
Bij de komst van de stoomtrein vreesden de boeren dat hun koeien zure melk zouden geven. De arme beesten zouden immers helemaal van slag raken door zo'n langsrazend voertuig.

Ai zal als *general purpose technology* veel veranderen, zoals ook de stoommachine, verbrandingsmotor, elektriciteit en het internet hebben gedaan. Wat dat precies is, kan helaas niemand overzien, en ik al helemaal niet. Krijgen we een periode van *creative destruction*, mogelijk met een recessie, en daarna een bloeiperiode? Of wordt het *steady* groei en krijgt de kwakkelende productiviteit eindelijk een impuls?

Het zal ons leven en werk veranderen, dat is zeker, maar hoe en in welke mate? Voor de beleggingsprofessional snijdt het mes aan twee kanten. U wilt natuurlijk weten waarin u moet investeren. U wilt het nieuwe Google van de ai uiteraard niet missen, terwijl u de ai-versie van de dotcombubbel graag overslaat.

Ook belangrijk: welke bedrijven en sectoren staan onder druk en dreigen door ai ge-*disrupt* te worden? Dat kan natuurlijk ook heel goed voor uw eigen sector gelden: wellicht bent u straks overbodig. Want misschien was u nooit zo bang voor robots als concurrentie, maar ai is een serieuze bedreiging voor de banen van hoogopgeleiden. Maar ook daarover ben ik een optimist: de geschiedenis leert dat er altijd nieuwe banen bijkomen. De *lump of labour fallacy*, ofwel het geloof dat er een vaste hoeveelheid werk in de wereld is, heet niet voor niets een *fallacy*.

Bedrijfsmatig is natuurlijk de vraag wanneer je welke technologie gaat gebruiken. Je wil niet te vroeg (wet van de remmende voorsprong!) maar ook zeker niet te laat instappen. En er komt vast nog een hele berg regels (en discussies) op ons af om ai in goede banen te leiden.

Wat ik wel zeker weet: we staan aan het begin van een enerverende tijd. Op een *inflection point* voor de curve steiler omhoog gaat. En het leuke is, hoe dat allemaal precies zal gaan en welke kant zich dat op ontwikkelt; dat kan ai dan weer niet voorspellen. Hebben we toch nog onze eigen denkkracht nodig.

Anna Dijkman is financieel journalist en columnist bij het FD

Smart Sustainability: How Artificial Intelligence helps to enrich Sustainable Investing

Valentijn van Nieuwenhuijzen

In the dynamic domain of investment management, a discernible paradigm shift is underway, compelling practitioners to reassess traditional approaches within an evolving landscape. Central to this evolution are dynamics that extend beyond customary metrics, prompting investment managers to synthesize considerations at the intersection of technological advancement and ethical responsibility. Against the backdrop of market adaptability to transformative forces, astute managers find themselves positioned at a juncture where exploration of emerging technologies converges with a conscientious appraisal of broader societal and environment imperatives.

Yes indeed, this machine-generated introduction (OpenAI, 2023) was prompted with instructions to generate a text that leaves the reader guessing if the article is about artificial intelligence or sustainable investing.

A DISCERNIBLE PARADIGM SHIFT IS UNDERWAY

Both artificial intelligence (AI) and sustainable investing (SI) are top of mind for investors. Breakthrough technological advances in the field of Natural Language Processing (NLP), like the great success of ChatGPT and other large language models, have heightened interest of investors in AI this year. For a bit longer already the growing demand for sustainable investments has caught investors' attention with equal force. In this article we explore what is happening at the intersection of AI and SI when it comes to the work of investment professionals. Taking inspiration from recent publications on environmental, social and governance (ESG) topics will hopefully trigger further thinking on the opportunities that novel AI offers to investment professionals and asset owners.

NATURAL LANGUAGE PROCESSING

The concept of AI spans a broad universe of technologies and has a rich history that includes periods of widely held enthusiasm as well as the proverbial AI winter. Despite AI's depth and breadth, many sustainable investing use-cases specifically leverage NLP techniques. This is typical for the field of sustainable investing and is at least partially explained by the specific data challenges that this rapidly evolving domain is facing. Where investors typically leverage tabular datasets in their research, they find themselves without those for many concepts when it comes to sustainability themes. Even when consensus grows on how to measure a specific sustainability theme and self (or regulatory) reporting starts (for example using tonnes of CO2 emissions to express climate impact), the next topics of interest present themselves already. The rapid evolution and diversity of sustainability topics make it fertile ground for investment researchers that can navigate alternative datasets using NLP.

Valentijn van Nieuwenhuijzen
Co-Head of Multi Asset Solutions (MAS) and Global Head for Sustainability for Public Investing at Goldman Sachs

The value of NLP as a technique to capitalize on alternative datasets is illustrated in recent research on sustainability. Consider Sautner et al. (2023) as a first of three examples handpicked from the 2023 vintage of papers that span both the E and S space of ESG. Sautner et al. (2023) create a novel firm-level climate change exposure measurement from earnings calls and investigate its relationship to real world impact (green hiring, green patent generation) and financial market outcomes (risk, risk- premia). To this end metrics needed to be created from things like earnings call transcripts, job postings and patent filings. Often textual sources of information lack the identifiers that practitioners typically use to navigate financial datasets. NLP can play a vital role in matching document to company identifiers. In this case attributing patents to the right organisation by connecting those organisation names to the names on patent filings is essential, but applications go beyond that, just think about social media or news articles in which companies' names are mentioned. Indicators like this can be helpful across the investment management value chain. They can aid investors in idea generation, portfolio construction and risk management.

A second illustration is on measuring biodiversity. With attention for biodiversity on the rise, it is important for investors to explore how biodiversity exposure can be quantified. Related to that: any relationship between exposures and financial returns is important to learn about. Giglio et al. (2023) demonstrate how biodiversity risks can be quantified. To this end, they first construct a news-based measure of biodiversity attention using newspaper articles and Google search activity. Thereafter, amongst others, they construct a firm specific measure of biodiversity attention based on corporate disclosures (financial reports). Studying the relationship between company and sector level biodiversity attention and the news-based biodiversity attention, they find that equity markets are already pricing biodiversity risks. Next to using classical NLP methods like keyword-based search, they apply attention-based transformer model to determine sentiment. Sentiment classification is a critical technique for differentiating between biodiversity risk and opportunity. Thanks to the recent shift to attention-based models that incorporate context, sentiment classification can be performed with improved precision. They no longer must rely on a 'bag of words' and run the risk of missing valuable nuancing in preceding or following words. Quantifying biodiversity (risk) is high on many investors' agenda; hence some level of practical ingenuity is required. NLP is likely to play an important role here.

Outside the environmental domain of climate and biodiversity, sustainable finance professionals are potentially even more challenged by the lack of availability of data in the social domain. Lohre et al. (2023) set out to measure social controversies and their impact on stock returns. Against a backdrop of earlier research that found negative ESG news to be detrimental to financial returns, they zoom in on social controversies. However, until now only a few datasets for social controversies exist, and for the widely known ones there is a high level of disagreement (low overlap) amongst their data vendors. To solve for this the researchers trained a bespoke classification model: ControversyBERT. When constructing their model that classifies news articles about a firm as containing a social controversy or not, they utilize an attention-based model to bring the context awareness discussed in the example above. With classical NLP approaches a sentence of the form "Company XYZ launches policy to tackle gender inequality" would be classified as a social controversy due to word matching on 'gender inequality'; however, with contextual awareness added this is unlikely to be flagged as a controversy (Lohre et al., 2023, p3.). Both deepening investors' understanding of the relation between social controversies and financial returns, as well as improved accuracy of classifying sentiment are important to the development of sustainable investing. The former as it answers one of the most fundamental questions for investors, the later as improved accuracy increases the level of comfort that practitioners can have with the model output.

INVESTMENT RESEARCH AND BEYOND

By now the usage of AI already has a decent history in investment management, an industry that is hyper competitive and focused on finding new ways to outperform the market or better tailor to evolving client needs. In contrast to the long history of analysing classical financial datasets that include prices, volumes and the fundamental data originating from accounting statements, the usage of textual analyses is a nascent phenomenon. The adoption of natural language processing techniques in their tool kits gives investment professionals an ocean of new opportunities for finding unique insights that can complement classical financial analyses at a scale. Suddenly information that was not systematically available in traditional datasets can be harvested at scale and be the ingredients for an additional layer of investment analyses.

USAGE OF AI IS INCREASINGLY EMBRACED IN THE INDUSTRY

For analyses of the sustainability profile of a firm, it is increasingly important to also consider forward looking elements like intentions, commitments, targets, etc. Historically, this would take researchers a substantial amount of time as they first need to find the information and then interpret it. AI has now reached the level of sophistication at which it can support researchers in both stages. It can reduce the time spend on search, thereby freeing up capacity for the more creative and expertise-based activity of interpretation of information. In this interpretation phase the scalability of AI can help researchers again, this time by enhancing their analyses via increased objectivity and completeness. Information extracted with help of NLP can for example be shown against a larger, more complete peer group, mitigating typical human biases in the process. When it comes to news or other local content, machine translation has now reached the level of maturity that local language articles can be incorporated into the analyses. As a result, more timely and complete insights on developments concerning investments internationally can be considered.

LOOKING AHEAD

NLP unlocks alternative datasets that contain a magnitude of relevant sustainability information for investors. It enables refined analyses via contextual awareness that can improve NLP model accuracy. It facilitates the joining of these new insights to classical finance datasets so that both real world and financial impact can be researched.

NLP UNLOCKS ALTERNATIVE DATASETS

So, does this imply all is perfect? Obviously not. Working with new tools and datasets also introduces new risks that need to be managed carefully. For tools, model suitability is an important consideration: the limited transparency of the popular context aware models might not be blocking for a researcher that is using it to speed up information gathering. However, for feeding directly into trade-signals a different degree of explainability is often required. Practical ways of navigating explainability are evolving. For example, in today's world of Large Language Models (LLMs) their powerful generative capabilities (they can 'generate' new text), summarization is a popular application. Summarizing an annual report of an investee firm on its sustainability information can be a timesaving exercise. However, the tolerance for model hallucination will be low if such summaries will be used for direct decision making.
To not fully scope out summarization from investors' toolkits, a differentiation can be made between extractive and abstractive summarization methods. Extractive NLP preserves the original text and therefore has a high degree of explainablity, a summary created with extractive NLP will act as a highlighter; each bit of information is traceable to its source and no new text is created. Abstractive NLP will be more eloquently written, but at the risk of hallucination text, hence not the most suitable for direct decision making.

When it comes to datasets there are several biases at play. Reporting biases, that might make available sustainability data not representative for the overall population of investable securities, as well as size and geographical biases that might skew portfolio construction unintendedly, just to name a few. Although considerations around model suitability and dataset biases are not new to finance professionals, domain expertise in both sustainable investing as well as data science will prove critically important to avoid the pitfalls on this journey and enable investors to capitalize on opportunities along the way. So more than enough focus needed on safeguarding prudent human oversight and effective quality control, but with these aspects in play a much smarter way of navigating sustainable investing opportunities and risks can be identified.

References

- OpenAI. (2023). ChatGPT (November 2023 version) [GPT-3.5]. https://chat.openai.com/chat
- Giglio, S., Kuchler, T., Stroebel, J., & Zeng, X. (2023). Biodiversity Risk. National Bureau of Economic Research. https://doi.org/10.3386/w31137
- Lohre, H., Nolte, S., Ranganathan, A., Rother, C., & Steiner, M. (2023). ControversyBERT: Detecting Social Controversies and Their Impact on Stock Returns. In The Journal of Impact and ESG Investing (Vol. 4, Issue 1, pp. 87–108). With Intelligence LLC. https://doi.org/10.3905/jesg.2023.1.076
- SAUTNER, Z., VAN LENT, L., VILKOV, G., & ZHANG, R. (2023). Firm-Level Climate Change Exposure. In The Journal of Finance (Vol. 78, Issue 3, pp. 1449–1498). Wiley. https://doi.org/10.1111/jofi.13219

Elicitation of sustainability preferences under MiFID II – Influence on the dynamics of financial advice

Sustainability preferences often do not translate into action on the stock market. A recent study reveals that Swedish households with a pro-environment stance, who value recycling, do not exhibit a higher likelihood of holding assets aligned with their environmental inclinations (Anderson & Robinson, 2022). A key reason for this disparity: a lack of knowledge about sustainable financial products. In a survey with 3,098 German retail investors, I found surprisingly low reported knowledge about sustainable investment with 40.9% stating to not be knowledgeable at all, while only 0.8% claimed to be very knowledgeable.

One solution to this issue could involve holding financial institutions responsible for informing clients and eliciting sustainability preferences. An exemplary case is the approach taken by *Pensioenfonds Detailhandel*. The Dutch pension fund proactively provided its members with a straightforward explanation of sustainable investment strategies, accompanied by the question whether these should be applied more rigorously in the allocation of their pension money. The outcome was remarkable: a majority of households supported enhanced sustainability in their investments, even with the potential for lower financial returns (Bauer, Ruof, & Smeets, 2021).

So, should other financial institutions follow suit and elicit their client's sustainability preferences to help them in selecting investments in line with these preferences? As of last year, this is no longer a matter of choice in the European Union. With the implementation of an amendment to MiFID II, all financial institutions and insurance companies are now mandated to elicit their clients' sustainability preferences. In my dissertation, I explored how this regulation change affects the dynamics of financial advice.

I first examined pricing implications in a study involving 415 professional financial advisors across the United States and Europe. Advisors managed investment portfolios on behalf of clients, who submitted either a conventional or a sustainable investment mandate. Advisors then set an advisory fee to each client. The results show that financial advisors charge a premium for sustainable investment mandates, with the study design ruling out differences in effort, skill, and costs as explanations. Instead, the results are consistent with the notion of price discrimination, where advisors exploit clients' sustainable investment preferences to extract additional profits.

Furthermore, this premium is primarily imposed on sustainable clients with low or unknown finance knowledge and diminishes when advisors perceive high client finance knowledge. Providing advice to sustainable investment clients enables advisors to earn higher fees, with sustainable investors being as likely as conventional clients to pay for advice, even at a premium.

In a separate study, I distributed information to clients of a German universal bank to observe its impact on investment behavior. Clients in my sample received a simple explanation of sustainable investments and a subgroup also received information about peers' inclination towards sustainable investing. The results show that peer information received during a buying decision significantly increases investments in stock funds labeled as sustainable. However, an analysis of account-level portfolio holdings over time reveals that the provided information does not affect later investment decisions.

Finally, I analyzed a large dataset from a European bank to explore whether retail investors respond to ESG scandals by divesting. For example, a shareholder may re-evaluate the sustainability of *Volkswagen*'s business practices after learning about the emissions test manipulation scandal. The results show that retail investors do not divest in response to ESG scandals. Interestingly, investors show a consumption response to scandals, consuming about twice as much out of dividends associated with negative ESG news sentiment, compared to income from companies without negative ESG news.

Overall, what do the results say about the merit of mandated elicitation of sustainability preferences? On the one hand, identifying investors with such preferences can allow financial institutions to increase sustainable investments by providing information at the right moment. On the other hand, financial advisors also leverage their clients' sustainability preferences through price discrimination that may threaten the long-term attractiveness of sustainable investments.

References

- Anderson, A., & Robinson, D. T. (2022). Financial literacy in the age of green investment. *Review of Finance*, *26*(6), 1551-1584.
- Bauer, R., Ruof, T., & Smeets, P. (2021). Get real! Individuals prefer more sustainable investments. *The Review of Financial Studies*, *34*(8), 3976-4043.

Marten Laudi
Postdoctoral Researcher at the University of Bremen

Demographics Unravelled:
How Demographics Affect and Influence Every Aspect of Economics, Finance and Policy

Review door Tjitsger Hulshoff

Mensen vormen de kern van elke economie. Hun gedrag als consumenten en werknemers is economisch veel relevanter dan hun aantal of leeftijd. Dat is het paradigma van Amlan Roy, in zijn boek uit 2021: "Demographics Unravelled: How Demographics Affect and Influence Every Aspect of Economics, Finance and Policy".

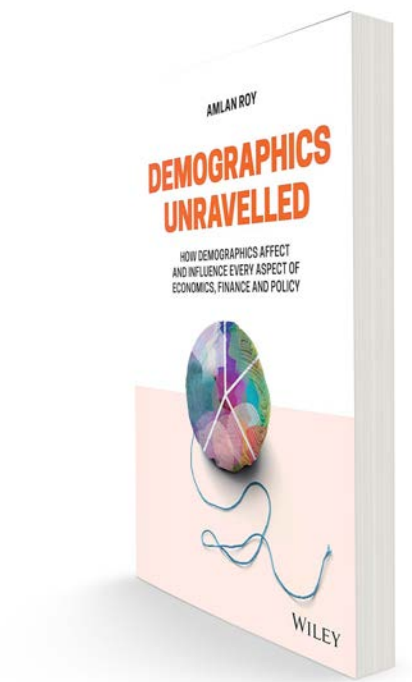

Auteur: Amlan Roy
ISBN: 9781119799139

Bij demografie denken economen aan het tellen en voorspellen van het aantal mensen per leeftijdscategorie, aan bevolkingspiramides, sterftecijfers en *fertility*-rates. Amlan Roy betoogt dat demografie veel dieper gaat en meeromvattend is. Het woord Demografie is een samenstelling van de Griekse woorden *demos* (mensen) en *graphos* (karakteristieken). Roy stelt dat demografie gaat over het gedrag van mensen als consument en arbeider. Dit gedrag is relevant voor economische analyses; en complexer dan het voorspellen van een bevolkingspiramide van tien jaar in de toekomst. Demografie beïnvloedt prijsstelling van activa, gezondheid, pensioenbeheer en politiek beleid.

Amlan Roy verdeelt zijn boek in zes themagebieden: i) Demografische grondslagen; ii) bevolkingsdynamiek, iii) de invloed van demografie op de macro-economische omgeving, iv) het verband tussen demografie en activaprijzen, v) problemen met gezondheid en levensduur bij verschillende bevolkingsgroepen, pensioenen en pensionering, en vi) het effect van demografie op levenskwaliteit, bestuur en duurzaamheid. Elk onderwerp is gekoppeld aan lange termijn rendementen en relatieve prijzen in beleggingscategorieën en marktsectoren,

In de hoofdstukken duikt Roy zowel de diepte als de breedte in. Hij zoekt een holistisch verband in economische theorieën en koppelt macro-economie aan asset-pricing, aan psychologie en beleidszaken. Het geheel is ruim gelardeerd met data, grafieken en voorbeelden. Het resultaat is leesbaar als proza alsook als naslagwerk. En een must-read voor beleggersprofessionals.

Demografie beïnvloedt de waardering van assets. Maar wat is de precieze dynamiek? Roy gaat hier op in. Demografie beïnvloedt de rente, en dan vooral de meest fundamentele vorm van rente: de *natural rate of interest.* En demografie beïnvloedt de aandelenrisicopremie. Asset prijzen worden bepaald door de voorkeuren van individuen. En de voorkeuren van individuen worden bepaald door de karakteristieken van mensen. Bijvoorbeeld de snelheid waarmee nieuwe technologie (zoals Artificial Intelligence) omarmd wordt, of de behoefte aan betere zorg bij stijgende welvaart en inkomens. Alles is verbonden en hoewel lastig te kwantificeren biedt Roy een denkkader en een stevig raamwerk voor analyse, onder de aanname dat demografie een onderliggende drijver is van de economie.

De management goeroe Peter Drucker zei ooit '*Demographics is the single most important factor that nobody pays attention to, and when they do pay attention, they miss the point*". Dat geldt in ieder geval niet voor Amlan Roy. En dankzij Demograhics Unravelled hoeft dat voor niemand meer te gelden.

COMMISSIES/VOORZITTERS 2023

COLOFON

VBA Journaal
is een uitgave van CFA Society Netherlands.
Het VBA Journaal verschijnt vier keer per jaar.

Hoofdredacteur
Dr.ir. Gerben de Zwart CFA

Adjunct Hoofdredacteur
Vacature

Redactie
Drs. Mark Geene RBA RMFI
Dr. Dirk Gerritsen
Dr. Roy Hoevenaars
Mr.drs. Manon Hosemann
Tjitsger Hulshoff MSc RBA
Ronald Kok MSc RBA
Frederik Mannaerts CFA
Dr. Sander Nooij CFA
Sebastiaan Reinders MSc RBA
Drs. Michel Wetser RBA

Auteursinstructie
zie: www.cfasociety.nl

Fotografie en illustraties
Cor Salverius fotografie
Max Kisman

Redactieadres & opgave advertenties
CFA Society Netherlands – Irma Willemsen
telefoon: 020 - 299 6590
e-mail: irma.willemsen@cfasociety.nl

Abonnementen
CFA Society Netherlands
Professor J.H. Bavincklaan 7
1183 AT Amstelveen
telefoon: 020 - 299 6590
e-mail: irma.willemsen@cfasociety.nl
Abonnementenprijs 2024:
€ 53 inclusief btw en verzendkosten

Opmaak en realisatie
az grafisch servicebureo bv,
Den Haag. www.az-gsb.nl

Cover ontwerp
Max Kisman

ISSN-nummer 0920-2269


MK2023