Why Coding Skills Still Matter in the Age of AI: Implications for Quantitative Finance

The Python Quants with ChatGPT

May 29, 2025

Abstract

Artificial intelligence (AI), particularly large language models like ChatGPT, is rapidly transforming how we work with code and data. In quantitative finance – a field at the intersection of finance, mathematics, and programming – these AI assistants promise to boost productivity and lower barriers to entry. This paper argues that strong Python coding skills remain as crucial as ever for finance professionals. Rather than replacing the need for programming expertise, AI tools amplify the capabilities of those who can harness them. We discuss how coding proficiency enables quants to fully leverage AI assistance, examine the limits of current AI tools, and present examples illustrating a growing performance gap between experts and non-coders. Our analysis suggests that AI is best viewed as augmenting human skill, not substituting for it. In an era of powerful AI, those with robust coding and domain knowledge are poised to benefit the most, potentially widening the skills gap in industry and academia. We conclude with implications for education and workforce development in finance, advocating for continued emphasis on programming fluency alongside emerging AI competencies.

Contents

1 Introduction

The rise of advanced AI assistants has led many to question whether programming knowledge will remain important. Tools like OpenAI's ChatGPT and GitHub Copilot can now generate code, explain algorithms, and even debug programs in response to natural language prompts. In theory, a finance professional could ask an AI to implement a trading strategy or risk model without writing the code manually. This raises a provocative question: *Will AI make coding skills obsolete in quantitative finance?*

Thus far, evidence suggests the opposite. Coding is still a critical skill – arguably even more so in the age of AI. AI coding assistants have impressive capabilities, but they are not magic boxes that eliminate the need for human insight. Instead, these tools function as *force-multipliers*: they empower those who have a solid foundation in programming and finance, allowing experts to accomplish tasks faster and tackle more ambitious problems. At the same time, relying on AI without understanding the code can be perilous. In quantitative finance, where mistakes can lead to significant financial loss or risk exposure, the stakes are high for getting things right.

This paper explores why coding skills still matter and how the advent of AI is changing (and in many ways elevating) the role of the human programmer in quantitative finance. We first review the role of coding in quant finance and how Python has become a dominant tool of the trade. Next, we survey the capabilities and limits of AI assistants in coding tasks, highlighting what they can and cannot do. We then discuss the notion of *skill amplification versus skill replacement*, positing that AI more strongly amplifies skilled practitioners rather than replaces expertise. Real-world examples and case studies are presented to illustrate how expert quants leverage AI, and what can happen when those without coding background attempt to use these tools. We also examine the widening skills gap that may result: as AI increases productivity for the knowledgeable, the disparity between expert coders and lay users could grow. Finally, we conclude with implications for education and industry, recommending that institutions treat AI as an enhancement to – not a substitute for – rigorous coding and domain training.

2   The Role of Coding in Quantitative Finance

Quantitative finance ("quant finance") applies mathematical models and computational techniques to financial problems. From pricing complex derivatives to algorithmic trading and risk management, coding lies at the heart of virtually all quant tasks. In the early days, quantitative analysts were often fluent in low-level languages like C++ for performance or used specialized tools like MATLAB. Over the past decade, Python has emerged as the lingua franca of quant finance due to its balance of simplicity, extensive libraries, and strong community support. Python's versatility makes it ideal for rapid prototyping of trading strategies, processing large financial datasets, and integrating with machine learning models.

A typical quant project might involve fetching and cleaning data, implementing a mathematical model, and then backtesting it on historical data. All these steps require programming. High-level libraries (such as NumPy, pandas, SciPy, and PyTorch) simplify many tasks, but using them effectively requires understanding how to code. Crucially, quant finance often demands *customization* and *flexibility*. Every trading strategy or risk model has unique nuances, and off-the-shelf solutions rarely capture all the necessary details. A quant with strong coding skills can tailor models to specific use cases, optimize code for performance, and troubleshoot issues when results look suspicious.

For example, consider the challenge of pricing a complex financial derivative that has no closed-form solution. A quant might employ a Monte Carlo simulation to estimate its value. Writing a Monte Carlo pricer in Python is straightforward for someone with coding experience. They can generate thousands of random price paths, compute payoffs, and average them, refining the simulation as needed for accuracy. Below is a simplified code snippet illustrating a Monte Carlo simulation for pricing a European call option:

Listing 1: Monte Carlo Simulation for European Call Option Pricing

```
import math, random

# Parameters

S0 = 100.0    # initial stock price
K  = 110.0    # strike price
T  = 1.0      # time to maturity (1 year)
r  = 0.05     # risk-free rate
sigma = 0.2   # volatility
```

```
def simulate_payoff():
    # Simulate end-of-year price under GBM model
    ST = S0 * math.exp((r - 0.5 *sigma **2)*T +
        sigma * math.sqrt(T) * random.gauss(0, 1))
    # discounted payoff of call
    return max(ST - K, 0) * math.exp(-r * T)

# Monte Carlo estimation

NUM_SIM = 1000000
payoffs = [simulate_payoff() for _ in range(NUM_SIM)]
price_estimate = sum(payoffs)/NUM_SIM
print(f"Estimated call option price: {price_estimate:.2f}")
```

Even without running the above code here, a knowledgeable reader can infer that this program will output an estimated price for the option by averaging a million simulated payoff outcomes. The ability to write and understand such code empowers a quant to solve problems independently. If an AI assistant were used in this scenario, it might help by generating a template for the simulation or suggesting vectorization for efficiency. However, the expert still needs to verify the simulation converges and that the code correctly implements the model's logic.

In summary, coding remains a cornerstone skill in quantitative finance. It enables practitioners to implement novel ideas, verify results, and adapt to new market conditions quickly. Python, in particular, with its dominant presence in both AI and finance, has made programming more accessible to finance professionals. Yet, as we shall see, the advent of AI assistants does not diminish the need for this skill – rather, it changes how the skill is applied.

3 The Capabilities and Limits of AI Assistants

AI coding assistants have made remarkable strides. Large language models (LLMs) like GPT-4 can synthesize natural language queries into working code. GitHub Copilot, which uses OpenAI Codex (an LLM related to GPT-3), can autocomplete entire functions and suggest solutions inside code editors. These tools have demonstrated the ability to handle routine coding

tasks and even complex algorithms in many cases. For example, a study by GitHub found that developers using Copilot were able to complete a coding task significantly faster – **55% faster** – than those without AI assistance. This kind of speed-up indicates that AI can shoulder much of the grunt work (boilerplate code, standard library usage, etc.), allowing developers to focus on higher-level design and problem-solving.

In the context of quantitative finance, an AI assistant can help in various ways:

- **Generating boilerplate code:** e.g., setting up data loading or writing a loop for the Monte Carlo simulation as shown earlier.

- **Explaining code or math:** If given a snippet of a pricing algorithm, a tool like ChatGPT can describe in plain English what it does, which can be a learning aid.

- **Suggesting improvements:** AI might propose vectorizing a simulation with NumPy for speed, or point out potential bugs or edge cases.

- **Exploring APIs and libraries:** Instead of searching documentation, a user can ask the AI how to use a specific function from a library (say, SciPy optimization routines) and often get a quick example.

These capabilities can indeed accelerate development. However, AI assistants have notable **limits and pitfalls**:

- **Lack of Context or Domain Understanding:** LLMs do not truly understand finance theory or the intent behind the code. They generate outputs based on patterns in training data. If a request is even slightly outside the distributions of data they saw, the result may be logically incorrect or irrelevant to the actual problem.

- **Possibility of Errors or "Hallucinations":** AI-generated code can contain subtle bugs or completely fabricated approaches that look plausible at first glance. For instance, ChatGPT might provide an incorrect formula for a financial metric that sounds authoritative but is in fact not used in practice. Blindly trusting such output could lead to faulty analyses or even financial losses.

- **Outdated or Biased Information:** If the AI's training data cuts off at a certain date, it may not know about the latest regulatory changes

or newly introduced financial instruments. It might also reflect biases present in the training data (for example, focusing on certain asset classes or markets).

- **Security and Compliance Risks:** In a regulated industry like finance, using AI tools raises questions of data security and compliance. If one naively allows an AI to use proprietary data in its prompt, there could be leakage of sensitive information. Furthermore, AI-suggested code might not adhere to compliance rules or could introduce vulnerabilities (e.g. using insecure random number generation for cryptographic applications).

Crucially, current AI tools do not remove the need for human review. Even proponents acknowledge that AI-generated code should be treated as a helpful draft, not final truth. As one commentary noted, even ChatGPT itself (when asked about writing code) emphasized that it should be used as a *guide* rather than an outright replacement for human coding. In other words, the model can assist by providing suggestions and templates, but a knowledgeable human must validate and adapt those suggestions.

In quantitative finance applications, this validation step is essential. Consider an AI assistant that generates a piece of code to calculate the Greeks (sensitivities) of an option portfolio. The code might run without errors and produce numbers, but are those numbers correct? A finance professional with coding skills would know to perform sanity checks (e.g. delta hedging tests, comparison with finite differences, etc.) to verify the output. A non-coder might simply accept the result, not realizing that, say, the AI swapped delta and gamma by mistake. Such errors can have serious consequences in trading and risk management.

Another limitation is that AI models lack **true conceptual understanding**. They can't reason about the finance domain in a deep way or ensure that a solution is conceptually valid. For example, an AI might propose using a certain machine learning model for predicting stock prices that a seasoned quant would reject outright due to overfitting concerns or lack of economic rationale. The human expert provides the domain judgment that the AI cannot. Indeed, AI often works best in a *hybrid* approach, combining algorithmic pattern-matching with human intuition. A recent study on ChatGPT in finance recommends such a combined strategy: human professionals should remain in the loop to oversee and cross-check AI-driven

decisions, thereby catching flawed suggestions and preventing over-reliance on the machine.

In summary, AI assistants greatly expand what an individual can do by handling rote tasks and suggesting solutions. They can be game-changers for productivity and exploration of ideas. Yet their outputs are not automatically trustworthy, especially in nuanced domains like finance. Strong coding and domain skills are required to filter and refine AI contributions. The next section explores how these dynamics mean that AI is far more empowering for experts than for novices.

4 Skill Amplification vs. Skill Replacement

A key thesis of this paper is that AI coding assistants amplify the productivity gap between experts and non-experts, rather than leveling it. In an optimistic scenario one might imagine that AI enables anyone, even with minimal programming knowledge, to perform like an expert quant: just describe the idea in English and let the AI produce a professional-grade trading bot or risk model implementation. However, the reality is much more complex. AI tools do not eliminate the advantages of experience and skill – if anything, they enhance those advantages.

Consider what happens when an expert quant and a novice both use an AI assistant to work on a task, say building a portfolio optimization model. The expert, thanks to their coding skills and finance knowledge, will likely do the following:

- Formulate precise prompts for the AI, asking for specific functions or code segments (e.g. "Generate a Python function to compute the efficient frontier for these returns").

- Recognize when the AI's output is correct or needs modification. If the assistant returns an algorithm that is quadratic in complexity, the expert might refactor it to be more efficient for large portfolios.

- Integrate multiple suggestions: perhaps use AI to generate data preprocessing code, separately ask it for optimization code, then glue those together with custom logic that the expert writes.

- Spot mistakes: if the AI's code uses a formula that the expert knows

8

is outdated or assumes normal distributions that don't hold, they will catch it and correct it.

- Augment with domain insight: the expert might prompt the AI to incorporate specific business constraints (like sector caps in the portfolio) that a generic AI solution would omit.

Now contrast this with a novice user who has limited coding and finance experience:

- Their prompts may be vague or overly broad (e.g. "Build a trading strategy for me"). The AI might then produce a generic answer or require iterative clarification which the novice struggles to supply.

- When the AI gives code, the novice might not know if it's correct. They may run it blindly. If it produces an output, they might trust it, even if the results make no sense financially.

- If the code has bugs or errors, the novice could be stuck. They might ask the AI to fix it, but without understanding the error, they can't guide the AI effectively. This can lead to a loop of trial-and-error that consumes time and may not converge to a solution.

- Critically, the novice lacks a framework to verify the model. For example, if the AI's trading strategy shows great backtest performance, a seasoned quant would suspect overfitting or look for unrealistic assumptions. The novice might just take it at face value, perhaps leading to losses when deployed in reality.

Empirical observations back up this view. In software development, it's been noted that "AI tools are ironically way more useful for experienced devs than novices," as one discussion quipped. The experienced developer knows how to leverage the suggestions and sift out incorrect ones. Similarly, in quant finance, those with stronger skills can get more out of AI. They treat the AI as an intern or assistant: helpful but needing supervision.

Academic and industry analyses reinforce that automation rarely fully replaces expert jobs but instead changes their nature. A Medium article by Shah (2023) argues that even if code generation becomes trivial, there remains a "significant disparity in how effectively people integrate that code into valuable, stable products". In quant terms, writing code is only part of

9

the job; understanding the model, the markets, and how to operationalize the solution is the bigger picture. Those who excel in those areas will use AI to go even faster. Those who don't have a solid base may create something that technically runs, but fails in edge cases or doesn't truly solve the problem. In other words, AI can allow engineers or quants with domain expertise to "produce better, more specialized software more quickly," whereas those lacking expertise "might still produce something that compiles, but it might fail in crucial edge cases or compliance checks".

We can draw a parallel with the introduction of calculators or spreadsheets in the past. When spreadsheets became ubiquitous, a finance professional who understood finance theory could perform analyses much more efficiently, while someone without that knowledge could plug in numbers but might not realize if the outputs were nonsense. Spreadsheets didn't eliminate the need to understand finance; they amplified the productivity of those who did. The same pattern is emerging with AI coding assistants.

It is also instructive to look at how human-AI teams perform versus AI alone. In the domain of chess, it's well known that a human using computer assistance (a "centaur") could for a time beat even the best computer or human alone. The human provides strategic guidance, while the computer handles brute-force calculation. In coding and quantitative work, we see a similar synergy: the human guides and validates, the AI generates options and fills in details. The combination can outperform either one by itself. But without the human's strategic guidance, the AI lacks direction; without the AI, the human might be slower in execution. This is *skill amplification* in action.

In summary, coding- and domain-knowledgeable quants are becoming even more effective with AI at their side. Far from making everyone equal, AI is creating new opportunities for those who can wield it properly. Meanwhile, those without the requisite skills risk falling further behind, which we explore next in case studies and the broader context of the skills gap.

5   Case Studies and Examples

To ground our discussion, we present a few illustrative examples and case studies that highlight the interplay between coding skills, AI assistance, and outcomes in quantitative finance.

5.1 Example 1: AI-Assisted Strategy Development by an Expert

Imagine a scenario at a hedge fund where a senior quantitative analyst is developing a new trading strategy based on alternative data (say satellite images of retail parking lots to predict company revenues). This quant has a strong programming background in Python and is familiar with machine learning. They decide to use an AI assistant to accelerate development:

1. The quant first outlines the approach: they need to ingest image data, extract features, correlate with stock movements, and build a predictive model. They break this into sub-tasks.

2. For each sub-task, they use the AI assistant in a targeted way. They ask it to write a Python function to load and preprocess satellite images (and because the AI was possibly trained on many examples, it quickly provides code using libraries like `opencv` or `PIL`).

3. They then prompt the AI for a function to perform a certain statistical test or to implement a known algorithm (perhaps a specific type of neural network for image analysis). The AI provides a starting point.

4. The quant reviews each piece of code. In one case, the AI's suggestion for feature extraction is not optimal for their data, so they modify it heavily. In another, the AI-produced model training loop has a small bug (it forgot to shuffle the data or to scale inputs), which the quant spots and fixes.

5. Once all components are built, the quant writes custom code to tie everything together and adds their own logic for trade execution (something highly proprietary that no AI would have in its training data).

6. They use the AI again to generate unit tests for key functions to ensure reliability.

The end result is that the strategy is developed in, say, 2 weeks instead of 4. The AI assistant saved time on boilerplate and provided useful hints, but the quant's expertise drove the project. Importantly, because the quant understood every part of the pipeline, they trusted the final product and could explain it to the portfolio manager and risk team. This example illustrates

how an expert can collaborate with AI to build a complex system more efficiently. It also shows that the expert's coding and domain knowledge was crucial at each step.

5.2 Example 2: AI-Assisted Strategy Development by a Novice

Contrast the above with a less experienced individual, perhaps a finance student who knows a bit of Python but is not very confident. Excited by ChatGPT, they attempt to create a similar trading strategy solely by prompting the AI.

1. The student asks a very high-level question: "How do I predict stock prices from satellite images of parking lots? Provide code." The AI, trying to be helpful, might output a long script. This script could involve complex libraries and model training. The student copies and runs it.

2. Suppose the code doesn't run due to an environment issue or missing library. The student asks the AI for help to fix the error. They go back and forth until it runs.

3. Now the script produces a result – perhaps some numbers or a model saved to disk. The student sees this and tries to use it to trade (maybe paper trading). The performance turns out to be poor.

4. The student is unsure if the problem is the strategy concept (maybe the data doesn't actually predict returns well), or the implementation (maybe the model overfit), or usage (perhaps they applied it to wrong stocks). They ask the AI a few generic questions ("Why is my strategy not working?"), to which the AI can only give generic answers like "maybe overfitting, try cross-validation."

5. Without deeper knowledge, the student eventually abandons the approach, concluding that the idea or the AI "didn't work." In reality, an expert might salvage the project by doing rigorous validation or by bringing in complementary data.

This hypothetical story, which mirrors real anecdotes one hears in quant forums, demonstrates how a lack of coding and domain expertise limits the

value one can get from AI. The AI delivered a bunch of code, but the novice user did not have the tools to truly evaluate or adapt it. The result was a strategy that likely never had a chance to succeed, and the reasons for failure remain opaque to the user.

5.3   BloombergGPT: Domain-Specific AI Requires Domain Experts

A real-world case study reinforcing our theme is BloombergGPT, a large language model created by Bloomberg LP specifically for the financial domain. BloombergGPT is a 50-billion parameter model trained on an extensive corpus of financial data, combined with general data. The development of BloombergGPT is telling: it required a team of experts in both NLP and finance to build a model that understands financial language and concepts. General-purpose models (like base GPT-3) sometimes falter on finance-specific tasks or terminology. By creating a domain-specific AI, Bloomberg effectively encoded a lot of expert financial knowledge into the model.

However, using BloombergGPT still requires expertise. Financial professionals interacting with it need to know how to pose the right questions. The model outputs need to be interpreted by someone who understands finance. For instance, if BloombergGPT is asked to summarize a complex SEC filing or extract key risk factors for a company, a finance expert is needed to judge whether the summary is accurate and complete. If errors or hallucinations occur, only a knowledgeable user would notice. BloombergGPT's existence underscores that domain knowledge (which includes coding ability to harness the model via APIs, etc.) remains vital. It was built to augment finance experts, not to replace them.

5.4   Morgan Stanley's AI Assistant for Advisors

Another pertinent example comes from the wealth management industry. Morgan Stanley, a leading global financial services firm, implemented a GPT-4 powered assistant for its financial advisors, integrating it with the firm's vast knowledge base. This internal chatbot, aimed at helping advisors answer clients' questions faster, was not simply deployed out-of-the-box. Morgan Stanley's technology team (which included programmers and prompt

engineers) worked extensively to fine-tune the AI and, importantly, to devise an evaluation framework to ensure the AI's answers are accurate and meet compliance standards.

The result was a tool that saw rapid adoption (reportedly 98% of advisor teams used it) and allowed even non-coder professionals (the advisors) to benefit from AI. But behind the scenes, it was the coding and AI expertise of the development team that made this possible. The system was carefully curated and is continuously monitored. This example shows how organizations can bridge the gap: those with technical skills create AI-driven solutions that others can use via natural language. Even so, the domain experts (advisors in this case) still need training on how to phrase queries and how to assess the AI's answers in a financial advisory context. If the AI suggests a particular investment based on some analysis, the advisor must apply their own judgement and knowledge of the client's situation. Morgan Stanley's approach highlights that successful integration of AI requires significant investment in skills and infrastructure, again emphasizing that human expertise (both technical and domain) is a critical component.

6 The Widening Skills Gap

The differential benefits of AI for experts versus novices raise a concern: AI could widen the skills gap in quantitative finance and related fields. Those who have strong coding skills and can effectively use AI will pull even further ahead in productivity and capability. Those who lack these skills might find themselves even more left behind than before.

Industry reports are noting this trend. In financial services broadly, there is worry that while AI adoption is accelerating, a "widening skills gap" is hindering firms from fully realizing AI's potential. The gap refers to the shortage of professionals who both understand the business/domain and have the technical skills (data science, coding, AI) to implement AI-driven solutions. As AI tools become more prevalent, simply having surface-level knowledge (like being able to operate an Excel spreadsheet) may no longer suffice for higher-value roles. Firms might start expecting even traditionally non-technical roles to have some familiarity with AI and coding concepts to collaborate with technical teams or to leverage no-code AI tools effectively.

On the quantitative finance front, we can imagine a bifurcation among practitioners:

- **"Super Quants":** Those who combine deep financial knowledge, excellent coding abilities in Python (and perhaps other languages like C++ when needed), and proficiency in using AI tools. These individuals or teams will be extremely productive. A single super quant armed with AI could possibly do the work that used to require several junior analysts, at least when it comes to prototyping and analysis.

- **Traditional Quants:** Those with financial knowledge and perhaps moderate coding skills, but who do not adapt to using AI assistance. They might continue to work in the old paradigm. They might still be effective, but relative to the super-quants, their output could be slower or less innovative. They may need to team up with others to cover the same ground.

- **Non-Coding Domain Experts:** Professionals in finance who rely on others to do the technical implementation. For example, a portfolio manager who understands an investment strategy but needs a quant team to implement the models. These individuals could start to feel more pressure as organizations encourage self-service AI tools. If they don't upskill, they might be limited to standardized tools that others prepare for them, potentially constraining their creativity or agility.

- **Non-Domain Tech Experts:** Another category are AI and tech specialists who lack finance background. With AI, some of these folks might jump into finance projects by virtue of their technical prowess. However, as discussed, lacking domain knowledge can lead to misapplication of AI. So, while they add to the mix, they typically need to partner with domain experts to be effective.

The net effect in the job market could be increasing returns to the combination of skills. Being good at just one (only finance or only coding) might not be enough for many roles; the premium will be on those who can do both, or who can work in cross-functional teams effectively. There is a risk that those who don't adapt will find their opportunities limited. For instance, a junior analyst who in the past might have spent their days manually munging data in Excel might find that task automated by AI. If they haven't learned Python or how to use the AI tools to add value in new ways (like designing better analyses or interpreting results), they could struggle to justify their role.

Education and training institutions are starting to recognize this. We see more programs aimed at upskilling finance professionals in coding and AI, and conversely, teaching data scientists about finance. The goal is to create more of the hybrid skill sets that the future seems to demand. Companies are also investing in internal training. As one AI skills report puts it: technology alone isn't enough – comprehensive training is essential to unlock AI's transformative power. This often means training in both the AI tools themselves and the underlying technical skills.

There is also a broader societal implication. If advanced AI tools require advanced skills to use properly, there is a danger of a divide where only large firms or well-educated individuals can harness the best technology, potentially increasing inequality. On the flip side, if done right, AI could help reduce the gap by making learning easier (e.g., AI tutors to teach coding) or by providing smarter interfaces that guide users. But as of now, the pattern in quantitative finance seems to be that those who proactively learn and embrace coding and AI are accelerating ahead.

In practical terms, to avoid a widening gap, organizations may need to focus on **upskilling** their workforce. This might include:

- Introducing Python and data science training for finance employees who historically didn't code.

- Encouraging a culture of learning where using AI assistants is standard practice, but coupled with checks and balances.

- Updating job descriptions to reflect the importance of both domain and coding/AI skills.

- Hiring from diverse backgrounds (e.g., hiring software engineers and training them in finance, and vice versa) to build teams that collectively have all necessary skills.

Ultimately, the presence of powerful AI tools raises the bar for human skills. In quantitative finance, as in many fields, it's the integration of domain expertise with technical prowess that yields the strongest results. Those prepared to meet that bar will find AI a powerful ally; those who aren't may find it a daunting competitor.

7 Conclusion

The advent of large language models and AI assistants marks a new era for quantitative finance. These tools can write code, interpret data, and even generate new ideas to some extent. However, the core message of our exploration is clear: **coding skills still matter – perhaps more than ever – in the age of AI**. Rather than rendering programming ability obsolete, AI has made such skills the key to unlocking its full potential.

Quantitative finance has always been a field that rewards a combination of mathematical acumen, financial insight, and technical skill. AI does not remove any of these pillars; it adds a new layer on top. A practitioner who is fluent in Python and understands how to work with models will use AI to supercharge their workflow, testing more ideas and processing more data in less time. Conversely, someone who tries to skip the learning and rely on AI as a crutch will quickly encounter the limits of their understanding. As we discussed, AI can produce working code, but knowing what to ask for and what to do with the answer remains the province of the human expert.

We argued that AI acts as a force multiplier for those with expertise. The case studies illustrated a growing divide in effectiveness between those who have coding and domain proficiency and those who do not. Far from democratizing the field to the point where anyone can be a quant, the current generation of AI tools may actually increase the advantage held by top performers. This doesn't mean that AI has no benefit for novices – it certainly can help them learn and perform simpler tasks. But to reach the frontier of what's possible, the guidance of human knowledge is indispensable.

There are important implications for education, hiring, and professional development in finance:

- **Education:** Universities and programs in finance and business should continue to emphasize programming and data science skills. In fact, they should incorporate AI tool training as well, but framed as tools that work in tandem with human skills. A student should graduate knowing how to code *and* how to use AI to help code, knowing the strengths and pitfalls of each.

- **Continuous Learning:** For those already in the field, staying current is crucial. This means not only keeping up with financial markets and theories but also keeping technical skills sharp. Learning how to effectively use new AI APIs, how to prompt ChatGPT for specific research

tasks, or how to validate AI outputs should become part of a quant's continuing education.

- **Collaboration:** The most successful teams will be those that blend skill sets. We may see more quants embedded in AI research teams and vice versa. The silos between "the quants" and "the tech people" in finance will need to break down, as the best solutions often require both in concert.

- **Ethics and Accountability:** Another dimension to consider is that with great power comes responsibility. AI can generate errors at scale. Firms will need clear guidelines on oversight. Just because an AI suggested a trade or a risk assessment doesn't absolve the human team from responsibility if it goes wrong. Maintaining rigorous validation and governance processes will be critical. This is another reason that having skilled people in the loop is non-negotiable.

In conclusion, AI is transforming quantitative finance, but not by making human quants irrelevant. Instead, it is changing the tools and amplifying what skilled humans can do. Coding is the language by which we communicate with these AI tools and ensure they do what we intend. Those who are fluent in that language will find that the AIs are powerful collaborators. Those who are not will find that simply having AI translate plain language into code is not a substitute for true understanding. The divide between expert coders and lay users may well widen before it (if ever) narrows.

The central thesis we put forward is that strong Python coding skills enable users to fully leverage AI assistants and models, and this amplifies the divide between experts and others. The evidence and examples we discussed support this view. Quantitative finance, being a high-stakes field that marries theory with implementation, serves as a compelling microcosm for observing this phenomenon. The lesson for professionals is clear: invest in your skills, learn to ride the AI wave, but do not abdicate your critical thinking and expertise. In the age of AI, the human coder is not obsolete; on the contrary, the human coder is the conductor that directs the AI orchestra, and with a good conductor, the music can reach new heights.

References

[1] Hall, David (2021). "2020 AI Finance Year in Review." *Medium.*

[2] GitHub (2024). "GitHub case study: Enhancing customer support with AI." *GitHub.*

[3] OpenAI (2024). "Morgan Stanley uses AI evals to shape the future of financial services." *OpenAI.*

[4] Scale AI (2023). "Guide to AI in Finance." *Scale AI.*

[5] Mishra, Sagarika, Michael T. Ewing, & Holly B. Cooper (2022). "Artificial intelligence focus and firm performance." *Journal of the Academy of Marketing Science.*

[6] Ali, Hassnian, & Ahmet Faruk Aysan (2023). "What will ChatGPT revolutionize in the financial industry?" *Modern Finance.*

[7] Khan, Muhammad Salar, & Hamza Umer (2024). "ChatGPT in finance: Applications, challenges, and solutions." *Heliyon.*

[8] Wu, Shijie, *et al.* (2023). "BloombergGPT: A Large Language Model for Finance." *arXiv.*