COSMOS
Navigate the Compliance Universe

The Complete Compliance and Ethics Manual 2024
Ethics, Compliance, and Artificial Intelligence

by Matt Kelly[1]

Artificial intelligence (AI) is likely one of the most transformational technologies ever invented. Perhaps that prediction is hyperbole. But for our purposes here, we can accept it as a given. How are compliance officers supposed to manage something so sweeping?

Indeed, the single biggest challenge that AI poses for corporate ethics and compliance officers today is its very name—we have too many concepts squeezed under the umbrella term "artificial intelligence." As a result, compliance officers struggle to identify all the risks that AI can pose—and rest assured, that list is long.

Without that precise understanding of AI's risks, companies can't respond to those risks efficiently and effectively. Compliance officers are left overwhelmed, management teams left uncertain, and all the while, the risks of AI keep piling up.

That is where we are today. The corporate world has been flooded with AI tools, with more coming onto the market every day. Compliance officers need to find a better approach to AI somehow, and they need to find it fast. Let's consider what that better approach might look like.

The Wide World of AI

First, we should understand just how wide the world of AI really is. The temptation is to associate the phrase "artificial intelligence" with flashy new arrivals, such as ChatGPT, Bard, Claude, and DALL-E—that is, generative AI, giving humans entirely new answers to the questions we ask.

In truth, AI has been around for decades. The first AI program was invented at Oxford University in 1951. It learned how to play checkers. Businesses have already been using various forms of AI for years. This includes:

- **Robotic process automation**: A finance team might use this to match invoices to purchase orders and then automatically send payment to a vendor.
- **Predictive analytics**: Marketing teams run algorithms to study customer behavior and then deduce what other products the customer might want to buy. When you shop on Amazon and see "You might also like…" product offerings at the bottom of your screen, that is predictive AI at work.
- **Automated filtering**: HR teams have programs perform keyword matches to screen out low-quality job candidates from high-quality ones.

Indeed, many compliance programs themselves use AI. For example, if you have a due diligence tool that flags high-risk third parties, and then those high-risk parties automatically receive additional compliance training modules from your training library, that's AI at work.

What makes generative AI different is *how* it works. Until recently, almost all AI tools were *process*-driven. That means they were designed to execute a certain process over and over (compare invoices, offer a new product, screen out certain parties) to produce the same sort of result every time. In contrast, generative AI is *data*-driven.

It studies a large trove of data to learn a "good" answer to the user's request. But the exact process that the generative tool uses to give you an answer . . . well, we're not always sure how that process works. (Then again, when you ask another person a question, do you see the exact sequence of neurons firing in their brain before they answer?)

Now, the floodgates of AI have opened. By mid-2023, a popular chart was racing around the internet depicting more than 120 AI tools people could use.

Is the difference between "traditional" AI and generative AI important to understand? Yes and no. Yes, because generative AI's different approach to answering questions means you have different risks to consider, especially around data privacy, cybersecurity, and unintended ethical dilemmas. On the other hand, the long history of AI in corporate life is a powerful reminder that as dazzling as generative AI might seem, it's just another technology meant to help humans. To that extent, AI is no different from previous transformational technologies, such as social media, mobile devices, cloud computing, and even the internet. The best way to manage the compliance risks for any new technology is to start with governance. As we explore the specific compliance risks and challenges that arise from AI, we'll keep coming back to the imperative for good governance over and over again.

What a Risk Assessment for AI Might Look Like

One intimidating point about AI is the sheer number of issues that it poses to corporations. This isn't like switching from Oracle to SAP to run your business systems or moving from an in-house email system to one managed by Gmail. Those business processes are already mature and well-understood; you're simply switching around the technology people use to run those systems.

AI will let corporations design entirely *new* business processes. It's more akin to adopting cloud computing or the arrival of mobile devices. It will allow an enterprise to set new strategic goals, change financial targets, and redefine its human capital needs. That said, AI will also change how a company interacts with customers, employees, and third parties, which, in turn, will create new operational and compliance risks.

Simply put, compliance officers will need to think about how their company might use AI and how others will use it. You'll need to consider how others' use of AI affects you and how your use of AI affects them.

To that end, we could follow this simple risk-reward matrix:

The above matrix is one example of how an in-house risk committee could start to game out the implications of AI. Bring together the people within your enterprise who'd have good insight into each of those squares, such as:

- IT
- Sales and marketing
- HR
- Product development
- Finance
- Compliance and privacy
- Legal

Then, start brainstorming or assign people to the squares most relevant to them. They can work up a list of potential risks and benefits. For example, compliance teams would presumably have lots to say about risks the company poses to itself and others. Sales would have better insights about the benefits of your company using AI and the risks of others using AI. Then, the committee should reconvene to compare notes. They can see where risks and benefits overlap or which risks and benefits are the largest and, therefore, should get the most attention. Start to develop a process to manage AI's arrival in your enterprise and your broader world.

When we talk about a risk-reward matrix and in-house risk committees, we're really talking about establishing a governance process to manage your company's embrace of AI. Why bother with that effort? Two points come to mind.

First, you should establish some sort of governance process because that is something the board will want to see. Technically, the board doesn't establish that governance process itself; it is there to assure that the management team has established a sensible governance process. If you haven't, and your company slowly finds itself outflanked by competitors embracing AI smartly, it's not the board's job to step in and develop that AI governance process. It is the board's job to replace the management team with new managers who can.

Second, establish a governance process because, without one, employees in your enterprise will start implementing AI on their own. That creates the one risk that senior managers hate most of all: They are surprised by something they didn't know their company was doing.