Bloomberg Law

Professional Perspective

AI in M&A Transactions: A Step-by-Step Drafting Guide That Comes With an Agenda

Stuart Irvin, Covington & Burling and
Adam Beguelin

AI in M&A Transactions: A Step-by-Step Drafting Guide That Comes With an Agenda

*Contributed by **Stuart Irvin**, Covington & Burling and **Adam Beguelin***

Artificial intelligence technologies are becoming increasingly important assets in acquisition transactions. In some recent transactions, almost the entire value of the enterprise was based on the target's AI assets. Although AI technologies make use of software and data, the unique attributes of AI assets (the attributes that give them value and create risk for the purchaser), may not be adequately captured by the standard representations and warranties on intellectual property in a merger agreement or asset purchase agreement.

There is no form or model for "how things get done" with respect to the allocation of risks in an AI transaction. Eventually, these forms and models will be developed and there will be standard practices and contract terms that are familiar to the lawyers and business people who do these transactions. But none of that exists now.

In the sections below, the authors will provide this model, including specific suggestions for how to draft representations and warranties in an acquisition transaction where AI is a key asset of the target company. This part of the article will avoid generalities of the type that are frequent in similar articles and provide specific terms that can be included *in haec verba* in acquisition agreements.

Before providing the step-by-step guide, the authors will address some of the larger issues that arise in the context of risk allocation in AI acquisition transactions. In particular, the authors believe that market forces can be used to create incentives for the ethical development and use of AI, and that the representations and warranties in purchase agreements are an effective way to reward the responsible use of AI and to punish its irresponsible use. We have an agenda that we are trying to advance with these contract terms that goes beyond the typical "how-to" guide that lawyers love.

Why Do We Need This?

AI assets are mostly software and data, and the risks associated with both software and data are covered in detail in any reasonably well-drafted purchase agreement. So the question naturally arises, why is any of this needed? Can't you just tweak the definitions of "Intellectual Property" and "Company Data" in a standard purchase agreement to include references to AI and then let the normal representations and warranties in the deal documents allocate risk between the buyer and the seller?

The short answer to that question is a definite "maybe." It may be possible to adapt the existing terms of an off-the-shelf agreement to allocate the risks associated with AI assets, but that is probably not the best way to do it, especially for those transactions in which AI assets are core to the business that is being acquired.

The entire purpose of requiring the seller to make representations and warranties is to isolate the particular risks associated with the business of the target company and to allocate the risk of loss between the buyer and the seller. The representations and warranties in the buyer's draft of the contract are designed to force the seller to make disclosures before the transaction is consummated.

When used in combination with a due diligence request list that requires the disclosure of key AI agreements and a description of the business practices used by the target company in connection with the use of AI, the representations and warranties, if properly crafted, will give the buyer a more-or-less accurate picture of how the target uses AI, which can help the buyer asses the risks associated with the transaction and adjust its acquisition strategy accordingly.

A breach of a representation or warranty in a purchase agreement typically gives the buyer the right to terminate the agreement between signing and closing. Such a breach will also typically give the buyer remedies in the form of indemnification and/or damages for breach of contract. By these means, the buyer and the seller essentially push money back-and-forth across the table based on the risks (or perceived risks) associated with the business of the target company.

It's important to note that the seller may not have any better information on the subject matter of the representation or warranty than the buyer does. But the seller may be forced to make the representation or warranty anyway so that the buyer will have a remedy if the representation is false. These contract terms are tools that buyers and sellers use to allocate risk of loss.

As a general rule, if a type of asset is particularly important to the business of the target (either due to its value or to the hazards associated with its use), then the representations and warranties relating to that asset type will be more detailed in the purchase agreement. For example, the representations and warranties on environmental hazards and groundwater contamination will be relevant and important in a contract for the purchase of a petroleum production facility and will be mostly irrelevant and unimportant in a purchase agreement for a software business.

As new types of risks are identified, the models that people use are expanded to take account of these risks. For example, as the risks associated with open source software or data breaches became better known and the losses associated with those risks became more dramatic, the representations and warranties in purchase agreements evolved to allocate the economic risks of such issues more clearly between the buyer and seller. Similarly, as risks become less important, the representations relating to those risks get smaller or disappear altogether. The "Year 2000" representations that were common in purchase agreements in the 1990s are ancient artifacts that only appear in an agreement these days when a lawyer forgets to update a very old form.

In the section on model terms below, the authors propose a set of representations and warranties that could be used to allocate the specific types of risks associated with AI assets between buyers and sellers.

Urgent Need for Carrots and Sticks

Before jumping into a lawyerly disquisition on representations and warranties, the authors felt it was important to address the threshold question of motivation. When discussing this article with colleagues, the initial reaction was almost invariably the same. Why, they asked, are you giving all of this away for free? The drafting of representations and warranties is part of the stock-in-trade for a transactional lawyer, so why publish these terms and give away your competitive advantage?

The answer is that the authors are trying to create a model that will help society avoid the doomsday scenarios associated with AI that Ray Kurzweil, Nick Bostrom, and Max Tegmark have described and that nobody can truly rule out.

Some have predicted an "intelligence explosion" where computers re-write and improve their own software without human assistance to become progressively more and more intelligent. If it happens, this process will likely take off in an incredibly fast manner. As Elon Musk has observed in a much-quoted blog post, "the risk of something seriously dangerous happening is in the five-year timeframe. 10 years at most." The authors share his concerns, and in particular those about an AI fast-takeoff that leaves government unable to effectively police the unethical or "seriously dangerous" use of AI technologies.

But even if one discounts the possibility of an intelligence explosion or a singularity, there is the risk identified by Yuval Noah Harari that "[l]ong before the appearance of super intelligent computers, our society will be completely transformed by rather crude and dumb AI that is nevertheless good enough to hack humans, predict their feelings, make choices on their behalf, and manipulate their desires."

Lawyers are incrementalists by training. The changes to society that will be brought by the widespread adoption of AI are not likely to evolve along a timeline that suits the business model of an A-List law firm. Something has to be done now to make the market police the ethical development and use of AI.

If the representations and warranties that we have proposed below (or something like them) are widely adopted and become part of how things get done, then unethical and irresponsible use of AI technologies will be identified and the prices paid to those who developed the technologies will be discounted. The incentives to take shortcuts in the race to develop new AI technologies will be reduced. This is the motivation that prompted the authors to draft the contract terms in the section that follows.

Model Terms for an AI Purchase Transaction

The representations and warranties and related defined terms (shaded in gray below) are intended be part of the section of the purchase agreement that deals with intellectual property. They make use of standard conventions for the drafting of contract provisions in M&A transactions, but will obviously have to be customized for the particular agreements to which they are added.

Not all of these representations and warranties will be appropriate for all transactions. They are offered with the usual caveats that the authors are not providing legal advice to anyone and accept no responsibility for their use.

The authors have provided "Negotiation Notes" at the bottom of each provision to explain the background of the terms and how to negotiate these provisions in the context of a purchase transaction.

Defined Terms

The following defined terms should be added to the definition section of the purchase agreement. Capitalized terms that are not defined have the typical meanings assigned to them in purchase agreements. For example, references to the "Company" below are references to the target company in the acquisition.

"AI Technologies" means deep learning, machine learning, and other artificial intelligence technologies, including any and all (a) proprietary algorithms, software or systems that make use of or employ neural networks, statistical learning algorithms (like linear and logistic regression, support vector machines, random forests, k-means clustering), or reinforcement learning, and (b) proprietary embodied AI and related hardware or equipment.

"Company AI Products" means all products and services of the Company that employ or make use of AI Technologies.

"Scraped Dataset" means Training Data that was collected or generated using web scraping, web crawling, or web harvesting software or any software, service, tool or technology that turns the unstructured data found on the web into machine readable, structured data that is ready for analysis.

"Third-Party AI Product" means any product or service of a third party that employs or makes use of AI Technologies.

"Training Data" means training data, validation data, and test data or databases used to train or improve an algorithm.

X AI Products and Services

X.1 Company AI Products. Part [X.1] of the Disclosure Schedule accurately identifies and describes (a) all Company AI Products that have been licensed, sold or offered for license or sale by the Company, [including those that have been (i) installed on customer networks or systems, or made available for download by customers; (ii) embedded in robotic, automotive, aeronautic or other hardware that is sold to customers, or otherwise provided to customers in tangible form; or (iii) hosted on a software-as-a-service, cloud services, AI-as-a-service, or similar basis for remote access and use by customers]; (b) all Company AI Products that are under development by or for the Company and planned to be offered for license or sale by the Company within the next [three years]; and (c) all Company AI Products that are not licensed or sold to third parties but are used internally (i) in connection with the design, development, manufacture or delivery of any [material] products or services of the Company; or (ii) to generate sales leads, analyze customer purchasing data, customer complaints and product returns, schedule equipment maintenance or replacement, analyze supply chain disruptions and vulnerabilities, monitor Company compliance obligations or for other similar technical, administrative and compliance functions of the Company.

> Negotiation Note: This representation is intended to identify the AI products of the Company and put them into three buckets, namely: (1) AI products that the Company sells or licenses to third parties now,

(2) AI products that the company has in the development pipeline and will sell or license soon, and (3) AI products that the company does not sell or license to third parties but uses in its own business. The bracketed language is optional but makes it clear that not all AI software is operated on a server at the Company. Some of the AI may be installed in customer systems, operated in the cloud or installed on hardware that is sold to a customer. Knowing where AI is operating may be important for assessing risks and understanding the costs and mechanisms required to monitor such risks.

X.2 Verification, Validation and Control. For each Company AI Product that has been commercially released by the Company, [except as disclosed in Part [X.2] of the Disclosure Schedule], there have been no interruptions in use of such Company AI Product in the past [twenty-four] months arising from or as a consequence of (a) the failure of the software, including any software embodying an algorithm, used in the Company AI Product to (i) satisfy its expected requirements or execute its specification correctly (e.g., verification errors); or (ii) fulfill the intended use and goals of customers or other stakeholders (e.g., validation errors); (b) the use of a "kill switch" or "circuit breaker" to prevent the Company AI Product (or component of a Company AI Product) from executing or completing a particular function; or (c) the use of any other emergency or failsafe mechanism (including human intervention as a failsafe) to prevent the Company AI Product (or component of a Company AI Product) from executing or completing a particular function.

Negotiation Note: AI products are not like other software. There are risks associated with AI products (like verification, validation, security and control issues) that may be more significant than with typical software. This representation is intended flag risks that are particularly problematic in an AI transaction so that the costs and losses associated with these risks can be allocated by the parties to the transaction. The representation reflects the thinking of MIT professor Max Tegmark and others on AI safety issues (see Research Priorities for Robust and Beneficial Artificial Intelligence by Stuart Russell, Daniel Dewey, Max Tegmark (2015)) and attempts to use the same taxonomy that these academics have developed. The ultimate goal of the representations in this Section is to enable the buyer gain a better understanding the errors that have occurred in the past. This understanding will help the buyer to assess (i) the value of the technology asset at issue, (ii) the likelihood that a given error (or type of error) will occur again in the future, and (iii) the Company's ability to evaluate an error and respond appropriately.

X.3 Security. The Company maintains [industry standard] access control protocols and capabilities that secure access to the Company AI Products and Training Data and, [except as disclosed in Part [X.3] of the Disclosure Schedule], there has been (a) no unauthorized access to the algorithms or software used in a Company AI Product, or to the Training Data used to train or improve a Company AI Product; (b) no unauthorized access to the IT systems used in the development, improvement or operation of Company AI Products; and (c) no use of the Company AI Product by a third party to engage in unlawful activity or any activity that violates the Company's license terms or terms of service for the Company AI Product.

Negotiation Note: The representations in (a) and (b) may be covered in other parts of the agreement that deal with IT systems. The drafter should check for duplication before adding these. However, the representation in (c) is likely not covered elsewhere. AI products, like other software, can be used for unlawful purposes, but the risks with AI are significantly higher than with other types of software. These risk and losses should be isolated and allocated between the parties to the transaction.

X.4 Improvements to Company AI Products. For each Company AI Product that has been (a) developed or improved pursuant to any specifications provided by a customer or partner of Company; (b) developed or improved using any Training Data provided by a customer, partner or other third party; or (c) customized in any material respect for any customer or partner of Company, the Company owns [or has an exclusive license to] all intellectual property rights in and to any such developments, improvements or customizations; and there are no restrictions on the Company's exploitation or commercialization of such Company AI Product or on the Company's ability to enforce its Intellectual Property Rights in such Company AI Product arising from or as a consequence of any of the foregoing. The Company maintains a technical description of any neural networks used in or with any Company AI Products (including a description of the learning rates selected for each such neural network) that is a sufficiently detailed so that the neural network can be modified, debugged and improved from time to time by programmers skilled in the development of AI Technologies.

Negotiation Note: This provision tracks a fairly standard representation on ownership of improvements to software, but note that some AI companies (particularly early stage companies) may have given IP rights to customers that are more extensive than a traditional software company would give to its customers. The business models for AI companies are still being developed, so big customers are in a position to ask for, and receive, IP rights in improvements. It's standard operating procedure for most large companies to insist on ownership of improvements as a cost of doing business, even when such ownership may not make sense in the context of the transaction. The real problem arises if a customer provides the target company with access to its internal data, and that data is used to train the target's AI algorithm. If the contract between the target company and its customers provide that the customer owns "improvements" to the target's software it may not be possible to separate the algorithm that has been trained with a customer's data from the improvements to that algorithm. These risks need to be identified, isolated and allocated between the buyer and seller.

X.5 Use of AI Products in High-Risk or Regulated Environments. For each Company AI Product that is used to make (or facilitate the making of) decisions in a hazardous, high-risk or regulated environment (e.g., credit worthiness, prison sentencing), the Company (a) retains information in human-readable form that explains or could be used to explain the decisions made or facilitated by the Company AI Product, and the Company maintains such information in a form that can readily be provided to regulators upon request, and (b) has complied with all the laws, regulations, and industry standards applicable to the Company AI Product.

Negotiation Note: This provision gets at the heart of one of the most difficult and problematic feature of AI implementation, namely the inability of humans to understand the decision-making process of the AI algorithm. As algorithms increasingly develop the capacity to improve themselves without human intervention, even the engineers who created them are frequently unable to explain their decisions. That is particularly problematic in regulated industries where understanding the decision-making process, and being able to explain that process to regulators, is important.

X.6 Training Data. Part [X.6] of the Disclosure Schedule accurately identifies and describes all third-party Training Data that is material to (a) the development of a Company AI Product, or (b) the ongoing operation or improvement of a Company AI Product (each, a "Third-Party Dataset"). The Company has complied with all license terms applicable to each Third Party Dataset disclosed or required to be disclosed in Part [X.6] of the Disclosure Schedule, including (i) the end user license agreement or other terms that govern the Company's use of any application programming interface used to collect Training Data, and (ii) the website terms or other terms that govern the Company's collection and use of each Scraped Dataset.

Negotiation Note: The more data that is available to train the AI algorithm, the better it becomes. This fact has lead one prominent technologist to observe that data is to the AI economy what petroleum was to the industrial economy. The value of data is increasing and some AI companies (particularly early-stage companies) are tempted to take shortcuts when it comes to collecting the data that is used to improve AI assets. This representation gets at this problem, and in particular seeks to identify the risks associated with two of the most common ways that data is harvested by AI companies; namely extracting unstructured data from a web service using an application programming interface or API and data scraping. The APIs used by most web-based businesses have license terms that restrict the use of the data extracted using the API. The terms of use for most web-based businesses have provisions that prohibit the use of web scraping or web crawling technologies to harvest data. Therefore, the data that is collected using these technologies can be tainted by breach of the applicable license terms or terms of use.

X.7 Ethical Use of AI Technologies; Transparency and Bias. The Company maintains or adheres to [industry standard] policies and procedures relating to the ethical or responsible use of AI Technologies at and by the Company, including policies, protocols and procedures for (a) developing and implementing AI Technologies in a way that promotes transparency, accountability and human interpretability; (b) identifying and mitigating bias in Training Data or in the algorithmic model used in Company AI Products, including implicit racial, gender, or ideological bias; and (c) management oversight and approval of employees' use or implementation of AI Technologies. [Except as disclosed in Part [X.7] of the

Disclosure Schedule], there has been (i) no actual or alleged non-compliance with any such policies, protocols and procedures; (ii) no actual or alleged failure of a Company AI Product to satisfy the requirements or guidelines specified in any such policies, protocols and procedures; (iii) no complaint, claim, proceeding or litigation alleging that Training Data used in the development, training, improvement or testing of any Company AI Product was falsified, biased, untrustworthy or manipulated in an unethical or unscientific way; and no report, finding or impact assessment of any internal or external auditor, technology review committee, independent technology consultant, whistle-blower, transparency or privacy advocate, labor union, journalist or academic that makes any such allegation; and (iv) no request from regulators or legislators concerning any Company AI Product or related AI Technologies.

> Negotiation Note: This provision seeks to allocate what is perhaps the greatest risk currently associated with the use of AI, bias in the training data. In a recent example, Amazon discontinued an experimental program that was designed to review job applicants' resumes and make hiring recommendations based on that review. The program used machine learning tools to give job applicants scores ranging from one to five stars. Amazon's algorithms were trained to score applicants by observing patterns in resumes submitted to the company over a 10-year period. Most of those resumes were submitted by men, so the internal dataset used by Amazon to train its AI was itself tainted with gender bias. The AI algorithms, reflecting the bias in the dataset, downgraded resumes containing the word "women's." The outputs of the AI (the number of stars awarded to an applicant) reflected the bias in the training data. ("Amazon scraps secret AI recruiting tool that showed bias against women," by Jeffrey Dastin, Reuters, October 9, 2018).
>
> This is an example of how ethical business practices adopted by Amazon resulted in the cancelation of a project could have created significant risk for the company had it been implemented. But it's easy to imagine a similar set of facts with an AI technology that was developed by a less scrupulous company. The representations and warranties above seek to capture this hypothetical set of facts.s

X.8 Third-Party AI Products. For each Third-Party AI Product that is used in the Company's business, the Company (a) has complied with all license terms applicable to such Third-Party AI Product; (b) owns any improvements to the Third-Party AI Product that are developed at the expense of the Company [or has exclusive or non-exclusive licenses to any such improvements]; (c) owns the model that is created by use of algorithms applied to the Company's owned or licensed Training Data or has exclusive or non-exclusive licenses to any such model; and (d) owns the outputs generated by use of the Third-Party AI Product at the expense of the Company [or has exclusive or non-exclusive licenses to any such outputs].

> Negotiation Note: There are vendors of technology solutions (like Google DeepMind and IBM Watson) that offer pre-built AI application services on a turn-key basis to corporate customers. The representations above seek to flag instances in which the target company may not own the AI model that is created, at least partly, at its expense or the AI outputs that are created by the vendor but used in the target's business.

X.9 Insurance for AI Risks. Part [X.9] of the Disclosure Schedule accurately identifies and describes (a) any insurance policy that covers risks or losses associated with (i) the marketing, sale, license or use of Company AI Products; (ii) the use of Third-Party AI Products by the Company; or (iii) the use of AI Technologies by the Company; (b) all claims made by the Company under any such insurance policy; and (c) any insurance policy of the Company that expressly excludes (by rider or otherwise) coverage for any of the risks or losses in the foregoing clauses (a)(i) through (a)(iii).

> Negotiation Note: When seeking to allocate the risk associated with the use of AI, it's important to know if any of that risk is covered by insurance or expressly excluded from insurance coverage. The authors are not aware of any policies that are written with the express purpose of covering or excluding AI risks, but it seems likely that these policies and exclusions will develop over time. It may also be useful to add a provision that requests information on caps in customer-facing contracts that limit the Company's liability to its customers. These liability caps could have an impact on the premiums paid for AI insurance, if and when it becomes available.

Enlisting Support for a Market-Based Approach

The representations and warranties outlined above are not, of course, a silver bullet that will fix all societal problems associated with AI. But they do make use of existing legal and market-based structures in a way that can be implemented quickly, on an international basis, and without waiting for government action.

The authors hope that others will further develop and improve the representations and warranties that we have proposed. We see this as part of an ongoing process that will require the efforts of many talented people, including lawyers in Hong Kong, Shanghai, London, and New York.

Representations and warranties in purchase agreements have been used for decades to identify illegal and unethical conduct. These contract terms are ideally suited for the purpose of allocating risk between the buyer and seller of AI assets and for the larger purpose of helping to create a system of carrots and sticks that reinforces lawful and ethical AI business practices.

The authors could like to thank Peter A. Stephan, a computer scientist of great ability, for his comments and corrections to the text.

