

A Story of Two Holy Grails: How Artificial Intelligence Will Change the Design and Use of Corporate Insolvency Law

*Felix Steffek**

* * *

I. Two Holy Grails

When I started to engage with artificial intelligence (AI) in the context of law, Måns Magnusson, professor of statistics at Uppsala University, asked the lawyers in the room: “What is the holy grail of lawyers? What is the best thing we could do using artificial intelligence in law?” At the time, my answer was: “Predicting what a court will decide.” With this answer, I did not only intend to capture the situation in which parties are in court litigating. I also meant “knowing the law,” in other words, knowing the legal consequences of actions and events should a court be invited to apply the law. What other holy grail could there be than anticipating legal outcomes as a result of decisions, independent of whether such decisions were actually judged in court or—as in most of the cases—part of [bargaining in the shadow of the law](#).

A little later, in a different context, I faced this question again and, having explained the first holy grail, was challenged to identify the second-best thing AI could achieve in law. This time my answer was: “Drafting a contract such that in cases of dispute or unforeseen events it will order the relationship between the parties such that the result matches the parties’ interests.” I thought that supporting private ordering would be another holy grail we should search for.

Another holy grail, inspired by the topic of this conference, “How AI Will Change the Law,” might be “predicting the law that will maximize welfare.” This concerns using AI for public ordering. However, at this point, using AI for public ordering in the form of drafting laws is only in its infancy. While it will not remain there, at least currently, I would find it very challenging to conceptualize and

* Felix Steffek is Professor of Law at the University of Cambridge. He serves as Director of the Centre for Corporate and Commercial Law (3CL) and holds a JM Keynes Fellowship in Financial Economics. He is Global Distinguished Professor of Law at the University of Notre Dame, SGRI Visiting Professor at Singapore Management University, and Senior Member of Newnham College. The author is grateful for the very helpful comments received at the conference on ‘How AI Will Change the Law’ organized by the University of Chicago Coase-Sandor Institute for Law and Economics, the University of Chicago Law Review, and the Oxford Business Law Blog, in particular from Tony Casey who contributed a formal comment.

implement a public ordering project using AI to write laws, in particular since AI cannot legitimately set the policy to be implemented.

Hence, for this Essay, I will run with the two grails I have identified: (i) predicting court decisions and (ii) predicting contracts. I am aware that there is some overlap between the two. In order to draft contracts, one needs to know the law. However, both issues can be functionally distinguished, and I hope you find some interest in following my quest to find both grails in the area of corporate insolvency law.

II. The First Grail: Knowing Corporate Insolvency Law

Starting with the first quest, I will explore (i) how AI can be used to predict the decisions of courts in insolvency law, (ii) what the limitations of such predictions are, and (iii) what needs to be improved to optimize the use of AI in predicting court decisions.

A. How AI can be used to predict court decisions

AI can be used to anticipate court decisions in at least two ways: first, predicting the general outcome in terms of “win,” “partly win,” “lose,” and “other” and, second, predicting the detailed outcome, e.g., the amount of damages awarded for insolvent trading or for how long a director is disqualified for misbehavior in financial distress. In terms of data input, there are also at least two core aspects: the facts and the application. The model is then required to predict the general and/or detailed outcome(s) on this basis. Possibly, this is combined with a prediction of the reasons for this outcome. This task aims to predict the outcome (and explanation) that a real court will provide, not the outcome that an ideal court would find. In reality, parties are mainly interested in knowing the law that will be applied to their decisions, not the law that should be applied.

Court outcome predictions have been focused on in AI research and practice for at least five years, and at this point, there are more than two hundred papers attempting to predict court outcomes in various jurisdictions for a range of courts based on different kinds of inputs and predicting different kinds of court decisions. At a first look, the results are promising. Depending on which score is used—accuracy, precision, recall, or F1—the results range between 50% and close to 100%, with many papers achieving precision, recall or F1 scores in the range of 80% and higher.

In our [research project](#), Helena Xie, Joana Ribeiro de Faria, Christine Carter, Jonathan Rutherford and I established baseline results for human experts and different AI models (BERT, T5, GPT-3.5 and GPT-4), predicting decisions of the [UK Employment Tribunal](#) in the format of “win,” “partly win,” “lose,” and “other.” The human experts achieved an F-score of 67%, while the best AI model (T5) attained an F-score of 56%. GPT-4 got close with an F-score of 55%. Some of the recall scores were impressive, such as the human’s recall score of 82% for “wins” and BERT’s recall score of 83% for “wins.” When the “win” and “partly win” results are aggregated, the F-score of the human experts rises to 81%, while T5 and GPT-3.5 achieve F-scores of 71% and 70%, respectively.

I am not aware of any paper predicting decisions in the area of insolvency or bankruptcy law. However, such predictions are neither conceptually nor from a data perspective more challenging than other areas of law, and it is only a matter of time until such a paper is published.

B. The limitations of such predictions

I have used the term “prediction” loosely thus far. More recently, many papers have been challenged on the grounds that they purported to predict court decisions while in fact classifying them. [Masha Medvedeva and Pauline McBride \(2023\) argue](#) that of the over 150 papers claiming to predict judgments that they investigated, only around 7% really do so. Their argument is based on comparing the information that parties have in reality when anticipating court decisions and the input data that the projects predicting court decisions use. While parties have the information, or parts of it, that one would find in a court docket before the judges have made any determination, many papers use facts extracted from court decisions to predict what the court will decide. The criticism is also that facts extracted from court decisions may be contaminated by information leakage. For example, judges summarize the facts in their decisions only after having formed an opinion, which may lead the facts to be selected and presented in a way that supports the final decision, and their description may be colored by the sentiments of the judges.

My own verdict is not as strict as Medvedeva and McBride’s since the term “prediction” is used widely in legal AI, and I consider research valid as long as a paper is clear and consistent about what it does and does not do. Nevertheless, I agree that many prediction papers use inputs that differ very substantially from the information that is available to the parties when they ask, “What will a court

decide?” The reason for this, however, is not an unwillingness of researchers or project teams to tackle the challenging task of curating and inputting such data. Instead, the reason is researchers’ lack of access to such data.

In my own research on AI and law, for example, there was no large legal dataset of UK court judgments available when we started. As a result, we spent two years creating a dataset of UK court decisions and have [published](#) it in the autumn of 2023.¹ It now offers around 320,000 UK court decisions, which is a good start. However, there is currently no dataset publicly available for the UK which contains court files or similar information that would allow a prediction using AI in a situation similar to the one parties find themselves in. The dearth of legal data for use in AI is similar or even worse in many other jurisdictions. In short, the lack of legal datasets stands in the way of finding the first holy grail.

C. How to improve the use of AI in predicting court decisions

Public ordering is needed to improve the data problem facing legal AI. Private actors have no say on which information concerning court litigation and similar processes such as arbitration and ombud proceedings is made available to the general public. In other words, the law on the provision of legal information to the public needs reform.

As regards *court decisions*, they need to be made available more comprehensively and freely for bulk download and use. I do not ignore the challenging issues concerning copyright and data protection. However, these issues can be solved—by public law reform, not the private actors. Laws can determine the copyright as regards court decisions and, as far as necessary, the anonymization of court decisions. In the end, the law applied and made by the courts belongs to the people, and it is difficult to legitimize withholding access to it. Also, instead of limiting access to data, restricting use of data is a less intrusive and, therefore, more proportionate approach.

Bulk access to *legislation* is easier, as it does not involve the tricky data protection questions. However, legal AI models need access to legislation in its evolution over time—not just a snapshot of today’s statutes. Therefore, I would like to recommend making legislation available for bulk download not only in its actual form but in its evolution over time.

¹ The dataset is available for researchers to apply for at *A Corpus of UK Court Decisions for Legal and AI Research*, UNIV. OF CAMBRIDGE, <https://perma.cc/V2JA-5C6B>.

Access to *court files*, i.e., information that is more similar to the type of information parties have when anticipating the law, is even more challenging than access to court decisions. However, if one shares my conviction that the law and legal consequences need to be predictable for those affected by them, then ways need to be found to provide large datasets to enable useful AI applications. One way could be anonymous case summaries provided by the parties or the judges for public access and written before any decisions are taken, thus avoiding data leakage. Possibly, such summaries can even be created by AI, thus lowering the cost involved in creating them. Another way could be the creation of synthetic data. However, synthetic data faces the criticism of not being real data, which is the kind of quality data models need.

III. The Second Grail: Contracting About Financial Distress

For the second quest, using AI to contract about corporate financial distress, I will apply the following structure: (i) how can AI be used to improve contracts about financial distress, (ii) which parts of corporate insolvency law will become default, (iii) what are the limitations of the use of AI, and (iv) how can the use of AI in contracting about corporate distress be optimized.

Before engaging with these issues, a short note on why we would still expect corporate insolvencies in a world of AI. Corporate financial distress would persist due to risk-taking incentives on the side of both the company and the lenders. Companies would continue to take risk as the necessary condition for the chance to make profits. Lenders would continue to lend with default risk present, since they are primarily interested in the overall return of their credit portfolio as opposed to the rate of insolvency in that portfolio (which is of secondary relevance).

A. *How AI will improve contracts concerning financial distress*

AI will make contracts concerning corporate financial distress more efficient. This applies to contracts both before financial distress sets in (*ex ante*) and after it has occurred (*ex post*). This is because AI decreases the cost of coordination. For example, AI can support the drafting of contracts in order to increase their expected utility, and it solves valuation deficits by improving the accuracy of expected values. The expectation of increased utility holds independent of which economic contract theory we apply: AI will reduce transaction costs (transaction cost theory), agency costs (agency and principal cost

theory), and will make contracts more complete (incomplete contract theory).

Using AI in contracting about financial distress will not necessarily mean that contracts will always regulate more specific details more comprehensively. It may well be, for example, that some optimal solutions of ex ante insolvency contracting will consist of assigning decision rights rather than predetermining automatic action for defined future events. As mentioned, the efficiency advantages of AI are not limited to ex ante contracting. They also improve ex post decision making, i.e., decisions in financial distress, such as collective decisions whether to rescue or to liquidate.

For the public ordering of insolvency law, the improved efficiency of contracting means that legislation and case law will become less mandatory and more default. If the actors are, overall, in a better position to contract such that their interests are maximized, then less public intervention is necessary. Corporate insolvency law will then facilitate the relevant actors to adapt their contracts to the specific governance and capital structure of the relevant company, an argument similarly developed by Professors Zohar [Goshen and Richard Squire](#) in the context of corporate law. Hence, a key doctrinal change that AI will bring for corporate insolvency law is that it will be less mandatory.

Less mandatory corporate insolvency law and more efficient contracting about corporate financial distress will breathe new life into the ideas of contractualizing corporate insolvency that started to be discussed since the late 1980s. Just to mention a few, Professor Lucian [Bebchuck suggested](#) establishing a system under which claimants could buy the assets of a company in distress, starting with the lowest-ranking group and then proceeding towards the higher-ranking group in case the lower-ranking group did not exercise their option to buy. Professor Alan [Schwartz provided](#) probably the most radical impulse, suggesting to generally contractualize corporate insolvency law by allowing debtors and lenders to enter into contracts about the actions to be taken in financial distress. Professor Robert [Rasmussen argued](#) for more private autonomy in international corporate insolvency law.

To sketch just one example along the lines of Schwartz's idea to contractualize corporate insolvency, AI could be used to predict future states of the company and draft contracts that fully reflect the preferences of the lenders. In addition, AI would be very well placed to automatically update such contracts as new information on the company becomes available. Similarly, AI could be expected to very well manage the challenge of large numbers of creditors.

B. Which parts of corporate insolvency law will become default?

Those arguing against contractualizing corporate insolvency law at the time—for example, [Warren and Westbrook](#)—based their criticism on a taxonomy of different kinds of actors. They distinguished strongly adjusting, partly adjusting, and nonadjusting actors. Strongly adjusting actors, such as banks, have the capacity to perfectly adjust their lending conditions to their interests, in particular as regards the insolvency risk. Partly adjusting actors, such as suppliers, are in a position to adjust only the entirety of their contractual relationships to their interests, e.g., the insolvency risk they are willing to accept. As a result, they adjust their contracts to the combined insolvency risk of the entire group of their debtors. Nonadjusting actors do not adjust the terms of their relationship at all. They are either involuntary creditors, such as tort creditors, who do not make an active choice to enter into a relationship with the debtor corporation, or they are voluntary creditors without the capacity or intention to adjust their contracts such as consumers.

This taxonomy is also useful in understanding which parts of corporate insolvency law will become default and which parts will remain mandatory. This will depend on the availability of AI services as regards corporate financial distress, and it will be driven by two factors: (i) the specific tasks AI services will offer to solve and (ii) who will be in a position to use these services efficiently. Partly and nonadjusting actors affect outcomes in two ways: contracts are not optimal (an efficiency aspect), and they may divide the available value in an unfair way (a justice aspect).

First, the future-specific, AI-driven financial distress services are difficult to predict in terms of their breadth, quality, and regional spread. As mentioned already, AI services or research focusing on corporate insolvency law is not particularly intense at the moment. However, some AI tools, such as the prediction of loan defaults or the prediction of recovery rates in litigation, are firmly or at least increasingly established. Also, the lower visibility of financial-distress related AI services is not due to structural problems or a lack of market demand compared to other areas of law. What is difficult to predict is whether further advances will be driven by generalist large language models becoming so powerful that further fine-tuning for corporate insolvency law is unnecessary or whether specific AI services will be built on top of the large language and other large data models.

Second, access to these services will be decisive for the reclassification of the relevant actors as fully, partly, or nonadjusting as regards corporate financial distress. [Eidenmüller and Wagner have investigated](#) the undesirable effects of one-sided access to AI-driven

contracting services. Against this background, the question arises how AI will impact the distribution of the fully, partly and nonadjusting classes of actors. Overall, the efficiency advantages will lead to an increase of the fully and partly adjusting groups and a shrinking of the nonadjusting group. For example, I expect small and medium-sized businesses to have cost-effective access to AI services, improving their contracting about financial distress with counterparties, leading to them being qualified as fully rather than only partly adjusting creditors. Beyond this, an important role will be played by start-ups and similar market players offering (B2)B2C services. In other words, it will be key whether high-quality services will be available at low cost to those usually classified as nonadjusting creditors, such as consumer creditors, employees, and micro businesses.

One group, however, will always remain nonadjusting: the involuntary creditors. Their characterization as nonadjusting is not based on their capacity or resources to adjust terms to the insolvency risk, but due to their not becoming creditors by choice. The classic example is a tort creditor having become the victim of a car accident, assuming that AI will not be able to prevent all types of accidents and other events leading to liability in tort in the foreseeable future. Hence, there will always remain a group of involuntary creditors, which need to be dealt with by corporate insolvency law.

So, what would happen if all relevant actors, in particular creditors and debtor companies except involuntary creditors, would become fully adjusting as a result of AI and other digital services? Then, there would indeed be a strong case to make corporate insolvency law default rather than mandatory except for those decisions that would negatively impact the involuntary creditors. Perhaps we already see some foreshadowing of such a development in those restructurings where nonadjusting groups such as consumers, employees, and micro businesses are paid in full while the professional lenders negotiate the restructuring of their claims.

C. The limitations of using AI to contract about financial distress

In addition to the availability of the relevant AI services and their costs, a serious problem of data availability needs to be solved to let the future scenario just sketched become a reality. Finding the second holy grail requires even more data than locating the first. AI tools helping to write contracts require not just legal data; in addition, they require nonlegal data. For example, whether the counterparty will breach a loan contract cannot be predicted on the basis of just legal data. Whether a contractual clause will be breached in the future, in most cases, cannot be predicted just by looking at the specific clause

and a dataset of loan contracts. Instead, nonlegal data—in particular on the business to be financed, its assets, business model, economic outlook, and key personnel—is necessary.

Bulk access to the relevant nonlegal data, in addition to the relevant legal data for AI-driven services for financial distress, is currently very limited for both factual and legal reasons. This is probably the main reason why public research and datasets on AI in corporate insolvency are so scarce.

D. What needs to be improved to optimize the use of AI in contracting about corporate distress?

Other than with legal data, private actors have some capacity to solve the problem of access to nonlegal insolvency-related data. Financial institutions such as banks and other professionals hold relevant data. However, legal restrictions such as data protection law and competition-driven self-interest make it unlikely that such data will become more widely used without public intervention. Limited data pooling by privately organized syndicates might be a way forward, but this is still to be witnessed. Perhaps private data intermediaries will emerge, but they will suffer from the restrictions just mentioned.

Hence, the question arises whether there is a case for legal reform to make nonlegal insolvency-related data publicly available. On a limited level, this is not unheard of. Many governments, regulators, and state agencies collect and make available some insolvency-related data. Take, for example, government-issued statistical series on insolvency and bankruptcy data or case studies in regulatory or register reports. In many jurisdictions, the collection and availability of nonlegal insolvency data has been significantly increased. Consider, for example, the EU best practice of public data collection mandated in Article 29 of the [EU Directive on Restructuring and Insolvency](#).

If one agrees that the use of AI in contracting about corporate insolvency law holds significant welfare benefits and is, at the same time, impeded by the lack of nonlegal data, then it is worth exploring the publicly mandated provision of such data to the general public. Doing so will involve innovative solutions and compromises in order to respect legitimate interests to keep certain data undisclosed for both personal and commercial reasons.

IV. Conclusion

Artificial intelligence holds immense potential to support those affected by corporate financial distress in understanding the legal consequences of their actions and negotiating contracts to solve the coordination problems caused by such distress. Both potentials are currently restricted by the lack of data, both legal and nonlegal. I recommend exploring legal reform to make both types of data publicly available. In this sense, access to corporate insolvency law will include access to legal and nonlegal data concerning corporate financial distress. If such access is provided, I expect corporate insolvency law to become less mandatory and more default. However, a group of involuntary creditors which require specific protection by corporate insolvency law will remain. The degree to which, and the areas in which, corporate insolvency law will become default will depend on the availability of related AI services and which groups have efficient access to such services.

* * *

Felix Steffek is Professor of Law at the University of Cambridge.