AI-assisted lawtech: its impact on law firms

DECEMBER 2021

UNIVERSITY OF
OXFORD

Acknowledgements

We are grateful to UK Research and Innovation (UKRI), who funded this research under the aegis of the Next Generation Services Industrial Strategy Challenge Fund. We also appreciate the assistance of our project partners, who include Allen & Overy LLP, the British and Irish Legal Information Institute, Burning Glass Technologies, the Law Society of England and Wales, Slaughter and May, and Thomson Reuters. We thank in particular the 50+ law firm, in-house, and ALSP representatives who generously offered their insights into the deployments of AI-assisted lawtech solutions within their organisations. We also thank the Solicitors Regulation Authority for their assistance with our lawtech personnel recruitment analysis.

White paper contributors

This white paper presents a number of key findings from the research project *Unlocking the Potential of AI for English Law*, carried out by an interdisciplinary team of researchers at Oxford University in collaboration with a range of partner organisations.

The AI for English law research team at the University of Oxford, led by Professor John Armour (Law), comprises Professor Jeremias Adams-Prassl (Law), Professor Abi Adams-Prassl (Economics), Dr Stergios Aidinlis (Law), Dr Václav Janeček (Law and Technology), Professor Ewart Keep (Education), Professor Thomas Lukasiewicz (Computer Science), Dr Richard Parnham (Management Studies), Dr Alina Petrova (Computer Science), Dr Matthias Qian (Economics and Finance), Professor Mari Sako (Management Studies), Dr Adam Saunders (Education), Ms Hannah Smith (Law), and Professor Rebecca Williams (Law).

The lead authors for this white paper were Richard Parnham, Mari Sako and John Armour.

Suggested citation: Parnham, R., Sako, M. and Armour, J. (2021). *AI-assisted lawtech: its impact on law firms*. Oxford: University of Oxford. December 2021.

Contents

Executive summary

This white paper presents a number of key findings from the research project *Unlocking the Potential of AI for English Law* ('AI for English Law'), carried out by an interdisciplinary team of researchers at Oxford University in collaboration with a range of partner organisations between 2019 and 2021. The AI for English law project involved academics from the university's law, economics, management, education, and computer science departments, working across six thematic research streams. The research was funded by UKRI under the *Next Generation Services* Industrial Strategy Challenge Fund.

In Chapter One, we explain what we mean by AI-assisted lawtech, and outline the prevalence of its usage by solicitors practising in England and Wales. Our survey-driven insights suggest that around half of all English and Welsh solicitors now routinely use at least one type of AI-assisted lawtech solution – with the important proviso that usage varies enormously by solution type.

> *"We also suggest that the deployment of AI lawtech is promoting a division of work between lawyers who help produce and refine the technology, and those who mainly use it as consumers."*

In Chapter Two, we explain how AI-enabled lawtech is impacting on lawyers' work. Here, our most significant finding is that the deployment of AI-enabled lawtech solutions typically involves the creation of new tasks, new working arrangements, a new delivery infrastructure, and an association with multidisciplinary teamworking involving both lawyers and non-lawyers. We also suggest that the deployment of AI lawtech is promoting a division of work between lawyers who help produce and refine the technology, and those who mainly use it as consumers. These developments, we suggest, may sit awkwardly with traditional law firm governance and career progression models.

In Chapter Three, we explore the impact of AI-assisted lawtech on law firm organisation and business models. We document an emerging culture of law firms partnering with third parties to develop AI-enabled lawtech solutions, rather than building solutions in-house. We identify common types of partnership between law firms and lawtech companies, and common contractual mechanisms governing these relationships. We also consider whether the deployment of AI lawtech solutions is prompting law firms to move beyond their traditional "legal advisory" business model, focused on bespoke legal advice. Some law firms are starting to embrace a "legal operations"-based business model, which focuses instead on internal process efficiency and project management.

In Chapter Four, we explore the challenges posed by the need for relevant data to train AI-enabled lawtech solutions. For publicly sourced data, we observe a reluctance by some agencies to share data with commercial entities. We also identify several uncertainties for law firms and lawtech companies wishing to use client data to train AI models. These include data ownership, client consent, and the sharing between stakeholders of the training gains in AI performance associated with relevant data.

In Chapter Five, we explore the possible impact of advanced technologies such as AI on law firms' recruitment patterns, training needs and internal governance. Consistently with Chapter Four, reporting law firms' partnering with third-party organisations to deliver lawtech solutions, we find that only a very small percentage of advertised law firm jobs currently require AI-related skills. Moreover, technological skills necessary for lawtech are more likely to be sought in roles advertised for non-lawyers than for lawyers. Moreover, we find little evidence that law firms are modifying their internal governance to clarify the career paths for non-lawyers at a senior level. That said, we also found evidence that lawyers were increasingly willing to develop skills associated with AI. These skills may, in turn, facilitate more effective multidisciplinary teamworking, and career trajectories into organisations that are not law firms. ■

CHAPTER ONE

AI-assisted lawtech - definitions and deployments

CHAPTER ONE

AI-assisted lawtech - definitions and deployments

In this chapter, we first explain what we mean by AI-assisted lawtech, and how it differs from other types of legal technology. We then discuss the technology's estimated usage levels by English and Welsh solicitors, based on a recent survey. The chapter aims to allow the reader to understand the capabilities of the technology and the state of its adoption in the market.

What do we mean by AI-assisted lawtech?

"Lawtech" is a general term referring to technological solutions deployed for use cases specific to the legal system. AI-assisted lawtech is a subset of lawtech that makes use of artificial intelligence. In turn, we take "artificial intelligence" to mean the use of automated systems to perform tasks normally requiring human intelligence.

In some situations, AI is one of the distinguishing elements of a lawtech solution; that is, it would be difficult for the solution to exist without its AI component. For example, a number of lawtech products focusing on due diligence and contract reviews have incorporated AI from the product's launch. In other contexts, AI is one element of lawtech's wider purpose, and may have recently been added to an existing lawtech solution to improve its functionality. For example, eDiscovery solutions were, historically, based around keyword searches. AI has now augmented the technology, allowing users to identify materials that match desired concepts, not just exact phrases. It should therefore be appreciated that the uses to which AI can be applied in the legal sector are constantly evolving. It is possible that the way in which AI is used in the future will be different to – and more expansive than – current usage.

At present, there are two distinct approaches to AI in computer science, which are reflected in the technical underpinnings of AI-assisted lawtech. One approach, which has seen enormous advances in the last decade, is based on machine learning (ML). This seeks to identify inductively relationships existing in data. It is used in the legal sector in conjunction with natural language processing (NLP), which converts textual information into vectors that can be processed by an ML algorithm. Alongside this, some solutions make use of "expert systems" based on a logic-driven or deductive approach to AI with quite distinct computer science underpinnings. While AI lawtech solutions have existed for decades, ML-based systems have been the focal point for legal practice adoptions in recent years, reflecting advances in the underlying technology, high-profile media attention,[1] and significant amounts of new investment.[2]

Table One (left) draws on a taxonomical classification of lawtech solutions developed by project researchers to illustrate the range of uses to which ML-based lawtech solutions are currently deployed (Sako and Qian, 2021). The majority of these examples focus on managing law-related tasks, or undertaking legal work. However, others focus on the management of the legal business itself. Not all AI lawtech is legal specialism specific: for example, some solution types, such as those that focus on legal research, can be used across multiple practice areas.

How does AI-assisted lawtech work?

A typical "practice of law" ML-based lawtech solution uses what is known as "supervised learning". That is, before it can be used, the solution first requires a human user to "train" it. This training process allows the solution to recognise data points, such as a specific term in a contract. The system is trained by a user electronically "tagging" illustrative examples of what they are searching for, using a "training" dataset. For example, if a user wanted to establish how many contracts within a large dataset have an English and Welsh

Table One: types of AI lawtech

Work type	Illustrative use cases
Managing the business	People and resources management; finance and operations; managing client relations
Managing and performing law-related tasks	Knowledge management; matter management; risk management; legal rights management
Performing work	Documents & contracts / transactions (including M&A); litigation

Source: Legal Geek / Thomson Reuters (2019)[3] / Sako and Qian (2021)

1 Notably, the launch of the artificiallawyer.com website, a daily news service wholly devoted to lawtech market developments involving AI-assisted lawtech companies, in 2016.
2 For details of significant recent UK lawtech investments, see https://datacommons.technation.io/lists/14662. See also "Spark" dashboard from LegalComplex, Crunchbase, Pitchbook, and Dealroom, all of whom document lawtech investments.
3 Legal Geek / Thomson Reuters (2019). *Lawtech startup report 2019: a maturing market.*

"jurisdiction" clause, then the user would first "tag" examples of such clauses. Having been trained – and quality assurance work undertaken – the solution can identify additional clauses, which broadly correspond with the previously tagged examples, within the main dataset to be examined. Exact key word matching is not required; instead, the solution can identify similar phrases, guided by statistical probability. When results are reviewed, the findings are typically fed back into the system, so that training is a continuous process and the system's functionality continues to improve over time.

Some ML-based lawtech solutions are highly adaptable: having been deployed to perform one law-related purpose – known as a "use case" – the same solution can then be deployed for another. The amount of retraining needed to use a solution for a new use case will depend on similarities between them: for example, there may be significant overlap between an M&A due diligence exercise and contract analytics, where similar clauses are being searched for. By contrast, there are likely to be fewer similarities between a due diligence exercise and a regulatory compliance review. The amount of training data required will also vary markedly by use case, depending on what is being searched for, and the extent to which standard industry definitions are widely adopted. We return to issues relating to necessary data inputs in Chapter Four.

One of the perceived benefits for legal practices who purchase a pre-existing AI lawtech solution, as opposed to building such tools themselves, is that the solution typically comes pre-trained. For example, a contract review tool may be able to identify jurisdiction clauses in contracts "out-of-the-box". However, as several of our interviewees made clear, this out-of-the-box pre-training is typically fairly limited in scope. As a result, additional training will often be required – particularly if the client matter involves novel issues or non-standard documentation to be analysed.

> *"A lot of these systems … claim to be [trained] out-of-the box. Our experience is that really none of them are trained out-the-box – or certainly not in an environment to the extent at which you could use it on live transactions. So, you've got to invest a lot of time training the models up to do what you want them to do."*
>
> *R&D specialist, large law firm*

> *"So, we work with a third-party vendor that provides the basic algorithms for some of the normal things that you would search for in a contract due diligence exercise, but we've invested quite a lot of time in training our specific instances of those algorithms to check for things that we look for on the transactions that we do. We've also dramatically expanded the scope of those algorithms, in the sense that we've [applied them in] different contexts and in different languages, which is, you know, very, very different from [the base product]."*
>
> *Chief Legal Innovation Officer, large law firm*

> *"It was partially pre-trained in that it had been used before on leases, but it hadn't been used on the reference points that we wanted it to be used at. So we had to train it."*
>
> *Associate, large law firm*

The need to train AI lawtech solutions to perform specific tasks can sometimes render them uneconomic for certain matters: it can be quicker, easier, and cheaper to use human reviewers instead. However, AI lawtech enthusiasts also report that solutions can often process far larger volumes of data, and far quicker, than human reviewers alone. In the document review space, for example, this allows a new approach to be taken: instead of reviewing a small sample of materials, the solution can instead review a far larger dataset – often in less time and/or at lower cost, than a human-led review. The solution's speed and scope is regarded as one of its main selling points; using such tools can enable firms to win work from rival practices who have not

"As several of our interviewees made clear, this out-of-the-box pre-training is typically fairly limited in scope."

deployed the technology. This is particularly relevant in situations where time is of the essence, and large datasets require extensive investigation.

Expert systems-driven lawtech solutions harness a different form of AI. At the heart of this type of solution is a "decision tree" – essentially a defined process, which sets out how the solution will process a legal matter. Users of this type of AI-assisted lawtech are guided through a specific matter in a highly structured way – for example, the user will be asked to choose between several options which, in turn, generate new ones. The ultimate output of this technology varies. But, in a legal-specific context, the output might be a customised contract, automatically generated in response to answers given by the user. Alternatively, the output might be a highly structured piece of legal advice, also generated automatically.

Because decision tree based lawtech solutions are highly structured, they are often regarded as being "brittle": alter one aspect of the decision tree, and the entire solution may fail. That said, modern law-based expert systems are often now easy for non-technical persons to create and update, using "low code" / "zero code" "drag and drop" user interfaces. As such, they can be developed by legal professionals with comparatively little support.

Understanding these technical underpinnings of AI systems helps to appreciate the practical limits of their utility. Supervised learning or rule-based expert systems require prior examples that can be used to train or design the system. Hence, they are not useful for bespoke or low-volume work. Moreover, AI systems as yet are weak on "social intelligence" – the interpretation of the subtle combinations of verbal and non-verbal cues that make up social interaction[4]. Hence, client work is unlikely to be automated any time soon. These technical underpinnings also help us to appreciate economic considerations around AI systems' deployment. The design, training, and implementation of AI systems involves costs. But, once they are up and running, AI systems can perform the work far more rapidly than humans. From a business point of view, the deployment decision therefore boils down to whether the economies of scale are worth the initial startup costs.

Figure One: use of AI-assisted technology, by organisation type

	In-house legal dept	Law firm	Grand total
Legal research	32.3%	25.0%	27.2%
Due diligence	12.1%	18.2%	16.4%
eDiscovery / eDisclosure / technology assisted review	13.1%	14.0%	13.3%
Regulatory compliance	10.1%	12.3%	11.6%
Contract analysis	8.1%	10.2%	9.6%
Other	10.1%	5.1%	7.1%
Fee-earner utilisation analytics and / or predictive billing	2.0%	10.2%	7.9%
Predictive analytics for litigation	1.0%	2.1%	2.0%

Other	In-house legal dept	Law firm	Grand total
18	19	236	353

*'Grand total' includes all complete responses, including from respondents working at ABS and legal technology solutions providers.

Source: Sako, Armour, and Parnham (2020)

Take-up of AI-assisted lawtech

A survey of registered solicitors, undertaken by our research team in partnership with Law Society of England and Wales, suggests that – as of early 2020 – around half of all survey respondents had used one or more form of AI-assisted lawtech (Sako, Armour et al., 2020). However, the survey also revealed variances in lawtech AI usage by technology type. As Figure One (left) indicates,

4 Frey, C. B. and Osborne, M. A. (2017). The future of employment: How susceptible are jobs to computerisation? *Technological forecasting and social change*, **114**, pp. 254–280.

the most commonly used AI-assisted lawtech was one with a broad potential user-base – legal research (27.2% of respondents reporting using this technology). By contrast, the least-used AI-assisted lawtech type was predictive analytics for litigation (2%). Between these two extremes was usage of AI-assisted lawtech that aided transaction-related legal work – due diligence (16.4%) and contract analytics (9.6%) – and lawtech that aided disputes work (eDiscovery / eDisclosure / technology assisted review – 13.3%). Access all use cases, take-up averaged around 12%.

These differences in use of course reflect differences in the tasks undertaken by respondents, and so it is hard to draw clear inferences about differential levels of uptake. For example, fee-earner utilisation is likely only to be conducted by one team for any given law firm, whereas multiple teams working within a firm may be performing disclosure or due diligence tasks. Other reasons for differences in uptake may include differential access to data. For example, our research suggests data supporting tools for predictive analytics for litigation has been particularly difficult to access. We return to data access issues in Chapter Four. ■

KEY TAKEAWAYS

- AI-assisted lawtech is a sub-set of lawtech. In some circumstances, the AI element is core to the tech's functionality. In others, it may be a recent addition to it.
- There are two main types of AI-assisted lawtech – those facilitated by machine learning / natural language processing and those facilitated by expert systems / decision trees. The former typically require training data to function, while the latter are based on pre-defined rules and decision trees.
- AI-assisted lawtech can be applied to both the business of law and the practice of law.
- Around 50% of solicitors in England and Wales are now using at least one form of AI-assisted lawtech. However, deployment by use case varies considerably.

CHAPTER TWO

The impact on lawyers of AI-assisted lawtech

CHAPTER TWO

The impact on lawyers of AI-assisted lawtech

This chapter focuses on two ways in which AI-assisted lawtech is impacting lawyers' working practices. We start by outlining these impacts during the rollout of a typical AI lawtech solution. We then briefly consider the wider implications of these shifts in working practices.

The impact of AI on lawyers' work – augmentation and substitution

It is often claimed that technology takes away jobs – the key idea being that automated systems substitute for humans, rendering them redundant. However, our research insights suggest a more complex picture. First, substitution of AI-assisted lawtech for humans occurs at the level of *tasks*, as exemplified by the use cases for deployment of AI discussed in Chapter One. Second, automating some tasks means that human lawyers have more capacity available to perform those tasks that cannot yet be performed by lawtech systems. These lawyers' productivity is therefore *augmented* by the AI systems. Third, putting lawtech solutions into action itself engenders a range of *new* tasks, which require a combination of legal and technical expertise to be undertaken.

The combined effects on *jobs* of these three changes is ambiguous[5]. Substitution of tasks only displaces *jobs* where roles consist largely or exclusively of tasks now capable of automation. At the same time, human lawyers find themselves better able to focus their energies on the tasks – such as bespoke work and client interaction – for which they have comparative advantage. This augmentation of their productivity may itself stimulate demand for legal services. Moreover, the deployment and use of technology creates demand for humans capable of performing the new tasks this necessitates.

To illustrate the changes in work practices associated with an AI lawtech solution deployment, Figure Two (left) illustrates a typical workflow, as revealed by our law firm interviewees.

In terms of the law-related tasks to be undertaken during a solution rollout, labelling work is more typically undertaken by junior legal fee earners. By contrast, the quality control / output review and client advisory element of the work is

Figure Two: workflow of a typical NLP-assisted lawtech use case rollout

Process flow	Work content in practice and business of law	Multidisciplinary human-machine inputs
I. Define a problem	Predict contract clauses that are likely to lead to disputes	Lawyer expertise
II. Select and test an AI model	Develop or buy an AI model which may be pre-trained; specify the project with process steps	Expertise in data science & project/process management
III. Prepare data	Establish a pipeline of contract data, label contract documents, and specify an ontology where appropriate	Lawyer & data science expertise
IV. Generate outputs	By ingesting the data using a chosen algorithm	Machine
V. Review outputs for consistency	Review outputs, and adjust data structure and/or the algorithm if necessary	Lawyer expertise
VI. Evaluate and explain implications of outputs	Present the results of contract review to the client, as a basis for providing legal advice	Lawyer expertise

Source: Armour and Sako (2020)

5 Acemoglu, D. and Restrepo, P. (2019). Automation and new tasks: how technology displaces and reinstates labor. *Journal of Economic Perspectives*, **33**(2), pp. 3–30.

typically handled by more senior personnel, up to and including law firm partners.

> *"In terms of their … background … they're not exactly paralegals but … they're not yet lawyers. So, they've done their first degree, often in law, or have done a conversion course, and are often looking for a training contract … we call them legal support assistants."*
>
> *Innovation head, large law firm*

> *"There were paralegals, trainees and just people from different teams. We even had admin assistants on it."*
>
> *Associate, large law firm*

> *"There are first level reviews, second level reviews, and third level reviews, and we're actually putting legally qualified associates, for the first time, into our [redacted] to also assist with that quality assurance exercise … The results are checked by senior associates, and ultimately the partner in charge of the matter, to ensure the interpretation of the results [is] correct."*
>
> *Innovation head, large law firm*

"Substitution of tasks only displaces jobs where roles consist largely or exclusively of tasks now capable of automation."

Both the labelling tasks and the follow-on review process represent a change to previous working practices. For example, in traditional due diligence, junior fee earners typically spent a great deal of time reviewing documents, searching for specific materials; now, these fee earners instead identify and label contract clauses within an AI lawtech solution. Having been initially trained, the tool then searches for similar contract terms within the main dataset to be evaluated. Further along this process, senior law firm personnel previously spent time reviewing clauses potentially relevant to the matter, initially identified by junior fee earners. Now, these senior lawyers instead spend time reviewing outputs mainly generated by an AI lawtech solution. However, the final evaluation and advisory element of the AI-assisted workflow remains largely unchanged, notwithstanding the introduction of AI lawtech into the workflow; advising clients remains the exclusive preserve of senior legal practice personnel. Here, the main difference is who, and what, informed the senior lawyers' advice to clients.

Besides the introduction of new tasks, a wider supporting infrastructure often needs to be put in place, in order to drive the deployment of AI-assisted lawtech solutions. This infrastructure can go by various names, including "Innovation", "Operations", and "Delivery" teams. This infrastructure is typically required for three main reasons: first, many of the AI lawtech solutions being deployed use immature, rapidly developing, technologies. These technologies often require extensive experimentation to evaluate their effectiveness, and strategic expertise to prioritise which solutions should be deployed.

Second, notwithstanding the ability of some AI lawtech offerings to be retrained for new use cases, most remain "points" solutions, deployed to address a particular issue in a specific practice area (Armour, Parnham, and Sako, 2020a). As a result, deployments – often multiple deployments – of AI lawtech solutions may need to occur, often on a practice-area-by-practice-area basis. It is therefore helpful for firms to employ professionals who are specifically skilled in the evaluation and rollout of such solutions.

> *"[A]round about three years ago, we didn't really have any kind of resource going into innovation – and I'm not exaggerating, it really wasn't until about then that we got this off the ground. … We're going through a process at the moment where we're looking at a handful of use cases in different practice groups, and testing different tools with each of the practices to really validate which we think is the best, or which might be the best-suited for particular circumstances, because we're not convinced that there's one right answer for the entire firm. So, I would describe that as a real early stage of development."*
>
> *Partner, large law firm*

Third, infrastructure is required to facilitate lawyers' understanding of the technology being deployed. We return to the necessary *skills* this entails in Chapter Five.

> *"When we get [AI domain experts] in a room with some of our partners, it's hilarious, the first half-hour. It's like they're all talking Russian to each other... It's just a kind of 'them and us', and it's ... you wonder how it's going to go. But, within about half an hour, 40 minutes, they're all going, 'Ah, I see what you mean'."*
>
> *Innovation head, large law firm*

As with any lawtech deployment, a typical AI lawtech solution rollout will initially require a multidisciplinary team (MDT) of specialists, including lawyers, technologists, data security personnel, process specialists, project managers, and trainers, all working together. However, where an AI solution deployment appears to differ from other types of lawtech rollout is that the need for an MDT does not simply occur at initial deployment stage – rather it can be an ongoing requirement, integral to the product's continued usage. We illustrate this point in Figure Two (page 10), highlighting the role of data scientists and project / process managers within the AI lawtech workflow. Data science expertise is initially required when selecting and testing the AI models to be used, and when preparing data for analysis (stage ii). However, this same expertise is also needed later in the process, for reviewing and interpreting the outputs (stages v and vi).

Further evidence of the association between multidisciplinary teamworking and AI solutions deployments emerges from our survey findings. In our survey of English and Welsh solicitors, we asked respondents to state which other types of specialist – both legal and non-legal – they worked with on a day-to-day basis. Options presented to them included paralegals and lawyers, legal project managers, process mapping experts, data analysists / data scientists and IT / legal innovation experts. Notably, solicitors who work with legal project managers, process mapping experts, data analysists / data scientists and IT / legal innovation experts on a day-to-day basis were more likely to use AI solutions than those who only worked with other lawyers or paralegals. This evidences an association between use of AI by lawyers and multidisciplinary working.

Table Two (left) illustrates the relationship in our survey data between respondents' AI use and their propensity to work on a daily basis with professionals from a non-legal background (which we term "working in an MDT"). Of the 97 respondents who indicated they worked in MDTs, 65 (that is, 67%) reported that they used any AI-assisted lawtech systems (in the contexts discussed in Figure One). In contrast, of the 230 respondents who worked only with other lawyers or paralegals, 98 (that is, 43%) reported using any AI-assisted lawtech systems. This difference is statistically significant, and remains so when other variables such as seniority and training are taken into account (Armour, Parnham, and Sako, 2020a). However, this correlation is imperfect: it can also be seen that many of our respondents who reported using AI do not work in MDTs, and vice versa. This perhaps reflects the fact that not all MDT teamworking focuses on the deployment of AI-assisted lawtech, and not all AI lawtech requires daily MDT working. Nevertheless, our survey suggests the differences in behaviour between solicitors who use AI lawtech users, and those that do not, in relation to MDT participation was statistically significant.

We now turn to the technical experts involved in a typical AI lawtech solutions rollout. As we discuss in Chapter Three, many law firms choose to partner with lawtech companies to deploy their AI solutions, rather than employ AI specialists directly. As a result, the AI domain expertise for a

Table Two: correlation between the use of AI-assisted lawtech and participation in MDTs

		USES ANY AI LAWTECH		
		No (column percentage)	Yes (column percentage)	Row totals
Works in MDT	No	132 80.5%	98 60.1%	230
	Yes	32 19.5%	65 39.9%	97
Column totals		164 100%	163 100%	327

Source: Armour, Parnham, and Sako (2020b)

solution rollout will often reside in the AI lawtech company, rather than the law firm. But, where firms opt to directly engage data scientists to assist with their AI lawtech deployments, the impression we gained from our interviews was that there were three main sources of recruitment for this expertise. One approach was to second data scientists on a consultancy basis from other organisations – not just technology companies, but also universities via what is known as knowledge transfer partnerships. Alternatively, the firm may recruit their own data science personnel, whether straight out of university or seasoned data scientists from other sectors. Some of our interviewees indicated that hiring data scientists was challenging, for law firms in particular.

> *"[W]e're obviously a very good firm, with a good brand name associated, but in terms of access to young talent, in the software space, they normally don't want to join a law firm."*
>
> *Chief Legal Innovation Officer, large law firm*

In light of the above, it is therefore likely that law firms who wish to hire data scientists will – in all likelihood – either have to "grow their own", and/or accept that their new hires will not have substantial experience of the peculiarities of the legal practice partnership form. We explore trends in the recruitment of AI domain specialists, and why this recruitment process may prove particularly challenging for traditional law firms, in Chapter Five.

Will the rise of AI lawtech have broader implications for lawyers?

In light of the division of labour involved in rolling out AI lawtech solutions, it is perhaps helpful to think of AI lawtech affecting lawyers' roles in different ways. As we have seen previously, one group of legal experts are heavily involved in the development of the software: its design, testing, labelling, and quality control, etc. We term these individuals "lawyers as producers" because they are helping to produce legal services enabled by AI technology. Similarly, another group of lawyers interact mainly with the outputs of AI lawtech solutions, which feed into their workflow and augment their ability to perform other tasks. We characterise these individuals as "lawyers as consumers" of AI-enabled legal services because their main interaction with AI lawtech solutions is as users of the outputs of the technology, rather than developers of and/or providers of inputs to it. Figure Three (left) illustrates this distinction and maps it onto the effects of AI on lawyers' work that we have introduced in this chapter.

Figure Three: effects of AI lawtech on lawyers' work

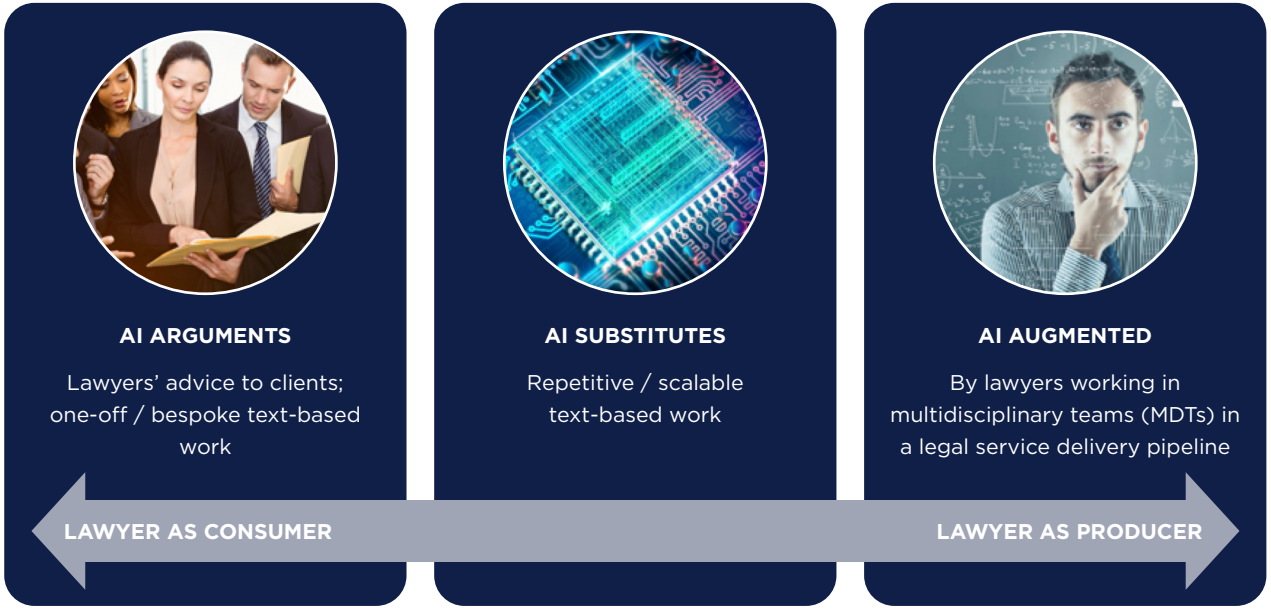

CHAPTER TWO

The impact on lawyers of AI-assisted lawtech

The distinction between "lawyers as producers" and "lawyers as consumers" tracks likely differences in necessary skills. While making use of the output of AI-enabled lawtech systems requires only an incremental acquisition of new skills for lawyers, participating in the design and delivery of outputs from such systems requires a multidisciplinary mix of skills. This raises questions about how best to acquire and combine such skills, and whether law firms are able to accommodate the divergence in working practices and skills requirements brought about by AI-assisted lawtech. We explore these questions in Chapters Three and Five. ■

KEY TAKEAWAYS

- Firms that deploy AI lawtech solutions may also require an innovation infrastructure to evaluate, prioritise, and lead their deployment.
- The deployment of AI-assisted lawtech is strongly associated with lawyers working in multidisciplinary teams. These teams do not, in themselves, need to be large, but typically involve a range of legal and non-legal disciplines working together to deploy the solution.
- It is possible that the advent of AI-assisted lawtech will drive a divergence in the work that lawyers do – resulting in a situation where some lawyers help produce lawtech solutions and others essentially consume lawtech solutions. This reflects a division in the underlying skills necessary for these types of roles.
- Those involved in producing AI solutions for the legal domain will need to acquire a higher level of technical/scientific comprehension – often by working as part of a multidisciplinary team – than those who simply make use of, or consume, the output of such solutions in the course of preparing legal advice.
- This bifurcation of work may have implications for both law firm governance and lawyer training.

CHAPTER THREE

Impact of AI-assisted lawtech on legal services business models

CHAPTER THREE

Impact of AI-assisted lawtech on legal services business models

In this chapter, we explore how AI lawtech impacts the way that law firms operate. First, we explore how law firms are partnering with third parties to develop AI-assisted lawtech solutions. We then consider how some legal practices are leveraging their use of AI lawtech to expand their business models.

Table Three: build, buy or partner?

Approach	How articulated?
Build	Developing an AI lawtech solution entirely in-house.
	Acquiring a pre-existing AI lawtech company, which then becomes part of the firm's own AI-related offering.
Buy	Buy and use the AI lawtech tool with little or no modification.
	Commissioning (or more) third parties to build an AI lawtech solution on the firm's behalf, which the firm then takes ownership of.
	Developing a unique solution that is nevertheless a heavily customised version of an existing off-the-shelf AI lawtech solution.
Partner	Co-develop an AI lawtech solution with a lawtech vendor.
	Co-develop an AI lawtech solution with a lawtech vendor, but also give them additional support (i.e. incubate / accelerate / become an investor in company).

Build, buy or partner?

In Chapter Two, we introduced the idea that lawyers could relate to AI-assisted lawtech either as "producers" or "consumers". When acting as producers, lawyers typically join a multidisciplinary team, and help build or customise an AI lawtech solution. By contrast, when acting as consumers, lawyers will typically use an AI lawtech solution – built by others – to help them do their job.

Just as individual lawyers can be either producers or consumers of AI lawtech, so too can entire legal practices. The contrasting strategic options are generally known as a "build, buy or partner". Through insights gathered from our interviews, we encountered various permutations of these broad strategic options, which we characterise in Table Three (left).

Among our law firm interviewees, a clear majority of their firms had opted to either buy (or customise) existing AI lawtech solutions, or co-develop new ones in association with a third-party vendor. Developing AI lawtech solutions entirely in-house was rare, and acquiring AI lawtech companies even rarer. Moreover, this appears to have been the result of a specific strategic preference: many interviewees articulated variants of the statement "we don't want to be a software company". A similar strategic choice was also evident among some of the ALSPs and MDPs whose personnel we interviewed.

> *"In terms of technology into legal, we'll very much partner and buy where we can. We're not a technology [company] – while we've got great technology expertise, it's more about applying that technology to right use cases. We're not looking to build where we can buy or partner."*
>
> *CTO, MDP*

> *"Now, because we are not a core software company that will licence technology out to others, we had to know … our strengths and also know our limitations. We don't have the technology budget that pure-play software companies have, so what we'll look to do is … analyse and determine how to best integrate the best software that's available. We will, however … we have almost 75 software developers on staff, and you'll say, 'Well, X, 75 software developers – that's a lot. Why? Aren't [you] really then a software company?' The answer is no."*
>
> *CEO, ALSP*

> *"I think sometimes services companies have to be careful not to confuse themselves with being a technology company and building a lot of [different] applications. They're entirely different business models. They're entirely different sets of expertise, entirely different funding."*
>
> *Digital contracting head, ALSP*

> *"Rather than starting from scratch, we thought … we can get the technology into the hands of our lawyers … [and] clients more quickly [through partnerships]."*
>
> *Department head, ALSP*

Where firms opt to partner with third parties to co-develop AI-assisted lawtech solutions, a common approach for the law firm is to assist the AI lawtech vendor with its data labelling exercise – often the core of the solution's eventual out-of-the-box functionality. For the law firm stakeholder, the obvious advantage of this approach is that the firm can avoid the cost and organisational challenges of employing their own in-house software developers and data scientists. Nevertheless, even this type of partnership may require the law firm to commit a reasonable level of internal resources: one firm we interviewed estimated that they had a team of seven personnel, six of whom were lawyers, spending half of their time labelling data on behalf of their lawtech company partner.

For legal service providers who make the (relatively) rare decision to "make" their own AI lawtech solution, a variety of reasons were offered for adopting this approach. These included a lack of off-the-shelf alternatives, or the aim that the self-built solution would have a superior functionality to those already on the market.

> *"So, it's only when there's a gap in the marketplace and there's a need in terms of providing a solution where we will delve into creating something proprietary."*
>
> *Digital contracting head, ALSP*

> *"If there's nothing that currently exists in [the firm] that we could use, the next thing is [to ask], is there something else out there in the market that could do this? Then, we go and have a look in the market and we'll test out products in the market, see if they'll fix that solution. If we still find there's nothing else on the market, that's when we start to have that thought around: is there something that we could build here to do this?"*
>
> *R&D specialist, large law firm*

> *"So, [solution name] was really the only option. And the technology was fine but, because they were trying to productise it, they couldn't tailor it enough to work for us. It wasn't a tech problem, it was … their business model wasn't what we wanted to do, and at that point, [our developer] was heavily involved in the proof of concept and looked at it and said, 'We could do all of this – we could build something that does all of this,' and built a really skinny system, to begin with, just really, really basic, and then we tried that and it worked and so we've built out from that."*
>
> *Partner, large law firm*

Where no AI lawtech solution yet exists, firms who wish to develop their own solutions can choose between hiring the expertise necessary to develop a solution in-house, or acquiring the expertise from a third party via a supply contract. Among our interviewees, different approaches were taken to this issue, even in relation to a single AI lawtech use case – for example, the development of solutions that aim to analyse historical billing data, with the aim of using this data to help predict the likely costs of future matters.

> *"We've been working with an outfit called [redacted], who are a team of data scientists who've helped us to build a bespoke tool for both cleaning up our historical billing data and then also for using that to try to make predictions about a future case, as and when it comes in."*
>
> *Partner, large law firm*

"A common approach for the law firm is to assist the AI lawtech vendor with its data labelling exercise."

> *"We actually have a team of data scientists that are looking at using machine-learning algorithms to more accurately predict the price of [a bale of] work. So, they're going through, you know, our billing system, our management systems now, our email systems, looking for patterns in the data to be able to more accurately predict, based on a relatively small set of questions, how much pay a type of legal product could eventually … cost."*
>
> *Innovation head, large law firm*

Types of partnership

While many law firms work in partnership with existing lawtech vendors to deploy, or co-develop, new AI solutions, some law firms have pursued partnerships with organisations other than established lawtech vendors.

Partnering with universities

This appears to be a favoured option for those legal practices working on early-stage, proof-of-concept use cases. Table Four (below) presents examples of this type of partnership, with a specific focus on AI-related lawtech collaborations. These partnerships are sometimes grant-funded.

Table Four: illustrative examples of recent university AI lawtech partnerships

Academic partner	Law / legal technology partner	Brief description of project	Grant awarded	Dates
University of Brighton	Family Law Partners (UK)	To embed knowledge engineering expertise to develop a rules-based decision support system to underpin a novel model of family law provision.	£97,012 (Innovate UK)	Oct 2016 – Feb 2019
University of Liverpool	Riverview Law UK	To leverage the university's expertise in areas of AI to facilitate the creation of a new service line.	£180,240 (Innovate UK)	Sept 2014 – July 2017
University of Manchester	Kennedys Law	To develop and embed an intelligent data-driven fraud prevention and detection service to support insurance claim handling utilising modern machine learning, text analytics techniques, and semantic technologies.	£79,936	Nov 2017 – May 2020

Source: Sako and Parnham (2021a)

Partnering with early-stage startups

Partnerships with early-stage startups usually do not begin before the "minimal viable product" stage. Some early-stage AI lawtech ventures are supported by law firms via incubators, accelerators, and / or funding. As Table Five (page 19) illustrates, UK-based law firm-supported incubators and accelerators have supported a large number of lawtech companies in recent years.

Partnering with general technology companies

Some law firms choose to partner with technology companies that have relevant technology skills and / or industry experience, but who are not specialist AI lawtech companies. An example of the former might be a company that specialises in data extraction, serving a range of different industries. An example of the latter might be a solutions provider that has prior experience of working in sectors contiguous to legal, such as insurance or property.

Adopting off-the-shelf solutions offered by major software companies

For example, one interviewee recalled using Microsoft's BotFramework to create an AI-powered legal chatbot; another said they used PowerBI to develop a data analytics-based offering.

The nature of the relationships between the law firm and the third-party technology supplier can also vary. For university partnerships, a favoured approach among our interviewees was to formalise the arrangement via what is known as a knowledge transfer partnership. Where firms offered support to early-stage lawtech startups, some opted actively

to guide the company's direction, including – if an investor – taking a seat on the company's board. Where firms engaged technology companies on a more straightforward contract for services arrangement, some were able to negotiate "fairly generous" licensing terms for the solution ultimately developed. This often occurred where the firm's engagement with the technology company ultimately enabled the company to launch its own AI lawtech solution.

New business models in legal services

Law firms' traditional business models can be characterised as "legal advisory" – providing bespoke legal advice to meet the needs of clients, raising revenue with input-based pricing in the form of the billable hour. That law firms appear to have outsourced much of the development work regarding AI to third parties suggests that this has not yet impacted fundamentally on their business models.

Legal operations business model

However, the deployment of AI-enabled lawtech is associated with the emergence of new business models in the legal services sector, currently predominantly adopted by organisations other than law firms. One of these is what we term a "legal operations" business model: a business-focused approach to legal service delivery, with an emphasis on efficiency. This satisfies the demand by businesses for legal service delivery that is efficient and responsive and integrated with the digital solutions being adopted in other aspects of business. Value creation is achieved by lowering costs, not only by labour cost arbitrage but also by applying business process re-engineering, process mapping, design thinking, and project management to improve the workflow and quality. AI enables the growth of this business model by automating and lowering the costs of various process steps. As a "pure play", the legal operations business model is currently predominantly associated with ALSPs. However, many law firms are also pursuing this type of business model on an auxiliary basis.[6]

Legal technology business model

We distinguish the legal operations business model from what we call a "legal technology" business model, which focuses on the design of the technical systems for lawtech solutions. For firms pursuing this model, value creation is based on product sales (licensing) or usage (so-called "software as a service"). There is considerable variation in the way in which firms adopting the legal technology business model go about charging for the use of their systems. Most offered their solution on a usage basis,[7] whereas others offered on a time-based subscription basis – for example, an annual fee. Some of the lawtech startup companies we

Table Five: selection of lawtech incubators / accelerators, which support AI lawtech companies

Name of entity	Supported by	Incubator or accelerator?	Estimated legal technology cohort size
Barclays Eagle Lab (legal technology-only cohort)	Various, including Law Society	Incubator	18* (in 2019)
Collaborate	Slaughter and May	Incubator	13
Deloitte Legal Ventures	Deloitte Legal	Incubator	14
Fuse	Allen & Overy	Incubator	22
MDR Lab	Mishcon de Reya	Accelerator	18
Scale | LawTech	PwC	Incubator	16

Source: Sako and Parnham (2021a)

6 Focusing on service innovation in general, rather than AI lawtech specifically, our SRA survey asked firms if they offered any kind of legal and non-legal services in combination with each other: just 6% of respondents said they did – although this rose to 20% among ABS and 22% among firms who acted for large corporate clients: Sako, M. and Parnham, R. (2021b). *Technology and innovation in legal services: Final report for the Solicitors Regulation Authority*, University of Oxford. July 2021.

7 However, calculation of "usage" varied considerably – sometimes by users/seats, in other cases by volume of materials processed.

interviewed freely admitted they had experimented with different charging models.[8]

Successful value capture under the legal technology business model necessitates having in-house skills in data science and software development, intellectual property (patents and copyrights), and sufficient understanding of the way in which legal workflows operate to be able to deliver an effective user experience. Design skills, human–computer interaction, legal project management, and customer liaison are all important aspects of the human capital mix.

The legal technology business model is currently predominantly associated with lawtech vendors. However, several of the ALSPs we interviewed had already developed, or acquired, their own lawtech solutions or companies, and were continuing to charge customers for software on a usage basis. They are thereby pursuing the legal operations and legal technology business models jointly.[9] By contrast, law firm interviewees, who rarely develop AI lawtech solutions in-house or offer AI lawtech solutions to clients, appeared not to be pursuing the legal technology business model even on an auxiliary basis.[10]

> *"If you look up and down our contracts, across all our customers, they're all in some form of recurring monthly payments and things like that, both for software and services."*
>
> *VP innovation, ALSP*

> *"So, the first pillar is about … the way in which our lawyers work internally, like internal processes; the second one is about … advising our clients about the tech that they use, and … delivering things for them; and then the last one is around tech products, so combining legal advice and technology together in a product, and selling that."*
>
> *Strategy and operations lead, large ALSP*

> *"We're not used to thinking of what we do as a product, and pricing that is very, very challenging. The partners certainly don't know how to do it. We have very little precedent ourselves. The market isn't really … mature enough to dictate it yet. So, it's sort of an open field, which is both exciting and can be quite daunting, frankly."*
>
> *Innovation leader, large law firm.*

Table Six: traditional and AI-enabled business models in legal services

Business models	What is sold? (What customers value)	Pricing (How value is created)	How value is captured
Legal advisory	Bespoke legal advice	Input-based (billable hour)	Trust, reputation, leverage
Legal operations	Process efficiency and project management	Output-based (fixed fee)	Process and project management capabilities
Legal technology	Technological solutions	Subscription, licensing	Intellectual property (copyright or patent) and platforms

Source: Armour and Sako (2020)

8 One acknowledged that – at the outset – their pricing regime was based on "what would [the customer] be willing to pay". Among both lawtech companies and law firms we interviewed, views on usage-based charging varied. Positively, this approach allowed firms to reduce the risk of paying for software they would ultimately decide not to use – a phenomenon commonly known as "shelfware". A usage-based approach to charging also makes it relatively straightforward for firms to pass through the cost of using AI lawtech solutions to their own clients as disbursements. Other interviewees, however, argued that a per-use charging structure for AI lawtech had the potential to discourage the solution's usage. For example, when charged on a per-seat basis, a firm has to make a conscious decision about who should be granted the right to use the solution, rather than allowing anyone to use it without limit. Equally, when an AI lawtech solution was charged for on a volume of materials basis, firms were often required to make a value judgement about the economic value of using the solution, in comparison with alternatives – typically, using a human reviewer instead.

9 See also University of Oxford / Thomson Reuters Institute / Georgetown Law (2021). Alternative legal service providers 2021. Strong growth, mainstream acceptance, and no longer an "alternative".

10 In relation to which lawtech solutions (AI or otherwise) firms are investing in, our recent survey of law firms (in conjunction with the SRA) found a clear focus on internal efficiency improvements (i.e. our previously mentioned "legal operations" business model), and a generally low prevalence of client-facing technologies (i.e. our "legal technology" business model). Sako, M. and Parnham R. (2021b). Technology and innovation in legal services: Final report for the Solicitors Regulation Authority, University of Oxford. July 2021.

CHAPTER THREE

Impact of AI-assisted lawtech on legal services business models

Will law firms change their business models?
As we have seen, most law firms have to date not significantly departed from the traditional legal advisory business model, albeit many have developed auxiliary capability in legal operations. Certainly, thanks to a benevolent regulatory regime in England and Wales, there are now few – if any – restrictions on law firms offering multidisciplinary legal services. Why do law firms nevertheless exhibit such hesitancy?

Law firms are traditionally organised as partnerships.[11] This organisational form is distinguished by senior employees being the owners, with corresponding rights to participate in business decision-making. This gives professionals relatively greater autonomy than they would encounter in a corporate enterprise, but tends to necessitate consensus-based decision-making. Organisational theorists suggest that the partnership form makes economic sense in law firms because the homogeneity of partners' disciplinary background – all being drawn from within the legal profession – helps keep the costs of consensus-based decision-making low[12].

"Many lawyers we interviewed said their firms see no need to change their business model."

However, a key distinction between the new business models legal operations and legal technology – and the traditional legal advisory model lies in the nature of the assets used to capture value. For both of the new business models, the necessary assets include technical systems and associated human capital, whereas legal advisory work is distinguished by its exclusive reliance on human capital. And within the necessary human capital, the legal operations and legal technology models require a multidisciplinary mix of expertise, whereas legal advisory success depends squarely on legal expertise.

Fully embracing a different business model would therefore require a law firm to diversify the disciplinary background of its human capital mix. In turn, this might be expected to fit better with a corporate, rather than a partnership, form. Consistently with this, ALSPs and lawtech startups are largely structured as corporations rather than partnerships. However, abandoning the partnership form would likely mean less autonomy for senior lawyers. This might make the firm less appealing to the legal human capital so crucial to the legal advisory business model. Hence there are synergies between the partnership form and the legal advisory business model; correspondingly, the partnership form generates frictions with the legal operations or legal technology business models. Together, these imply an inherent limitation on the extent to which a firm that is committed to a legal advisory business model can embrace the new business models.

Consistently with this account, many lawyers we interviewed said their firms see no need to change their business models: law firms are generally highly profitable, and have few difficulties in attracting either clients or (legally trained) employees. Moreover, our recent survey of law firms in England and Wales, conducted in association with the SRA, reported that, among barriers to innovation identified by law firms, one of the most commonly cited response was that "it is not a strategic priority" – both in relation to lawtech specifically, and innovation more generally (Sako and Parnham, 2021b). Similarly, for most of the traditional law firms we interviewed, our impression is that issues about career progression for non-lawyer technologists had not been addressed. Non-lawyers who joined law firms simply had to fit into an organisational form that is not designed with their career progression needs in mind.

Is this approach by law firms sustainable? Based on the evidence we have gathered, we suggest "quite possibly", for several reasons. Given that many law firms outsourced the technical element of AI-assisted lawtech tools to third-party developers, it is difficult to envision a situation where internal pressures within law firms compel them to update their governance regimes to facilitate the career progression of the (few) AI lawtech professions that they opt to hire. Rather, we envision only a modest impact of AI lawtech on law firms' internal governance regime, associated with a modest expansion of their traditional legal advisory business model, towards one that also embraces

11 The partnership form was formerly mandated for law firms, but these restrictions were abolished in England and Wales by the Legal Services Act 2007.
12 Hansmann, H. (1996). The Ownership of Enterprise. Belknap Press.

the (largely internally facing) legal operations business model. This expansion of activities arguably reinforces, rather than challenges, law firms' traditional monoprofessional focus. By contrast, we regard it as unlikely that law firms will develop significant new service lines in either client-facing AI lawtech solutions development, or client-facing AI lawtech deployment consulting. Because of the complementarities between organisational form and business models, that type of work, we suggest, is likely to remain largely the preserve of lawtech companies or ALSPs. ■

KEY TAKEAWAYS

- The majority of legal practices we interviewed had bought / adapted existing AI lawtech solutions, rather than developing their own in-house. This is associated with strategic decisions to avoid becoming "software houses".
- "Build" decisions tended to occur where off-the-shelf AI lawtech solutions were not already readily available.
- There are several types of third-party organisation that law firms can partner with to develop AI lawtech solutions. Potential partners include universities, non-lawtech companies, and early-stage lawtech startups. Generic tech can also be adapted to perform AI-assisted lawtech tasks.
- The production of AI-enabled legal services is associated with new business models, which we term legal operations (centring automation to scale efficiently) and legal technology (designing and marketing systems for legal applications). In contrast to the classical legal advisory business model, these new models depend crucially on nonhuman assets, and a multidisciplinary mix of human capital.
- Law firms are using AI lawtech to generate efficiencies within their existing legal advisory business models. This has not, to date, resulted in widespread change in their business models, beyond the auxiliary development of legal operations capacity.
- The partnership form may be a double-edged sword: complementary to a legal advisory business model but creating friction with the new business models. This implies an inherent limitation on the extent to which a firm that is committed to a legal advisory business model can embrace the new business models.

CHAPTER FOUR

The challenges posed by the use of data needed to train AI lawtech models

CHAPTER FOUR

The challenges posed by the use of data needed to train AI lawtech models

In this chapter, we briefly explore the role of data in training AI-assisted lawtech – what type of data can be used, and for what purposes. For supervised learning-based AI lawtech solutions, the volume of training data available must be sufficiently large, and also representative of the matters to which the model will be applied (Armour and Sako, 2020). What amounts to a sufficiently large dataset will vary by use case and algorithm deployed.

Data sources

We identify two distinct sources of data to train AI-assisted lawtech tools: publicly owned data, produced by the state, and privately owned data, collated by organisations. An example of a public data set, known to be used by AI lawtech companies in the UK, is that of employment tribunal decisions, published by HM Courts and Tribunal Service. An example of privately owned data would be an organisation's own internally generated data, such as the text of the contracts it has entered into with commercial counterparties, or a database of employee disputes or accidents at work.[13]

Data challenges

Access to publicly owned data can be problematic for AI lawtech firms. Because the justice system gathers personal data about individual participants, data privacy law – the GDPR and Data Protection Act 2018 – imposes constraints on the sharing of this data with those wishing to develop lawtech applications. Our team's analysis suggests that these constraints do not preclude sharing with lawtech firms, provided appropriate precautions are taken, but uncertainty around the interpretation and enforcement of the law tends to engender a cautious approach that may render the regime's practical application more onerous than it needs to be (Aidinlis, S., Smith, H., Adams-Prassl, A. et al., 2021).

Turning to privately owned data sources, particular challenges arise in the application of AI to contract-based work – contract analytics or due diligence, for example. These flow from the relationship between access to relevant data and the performance benefits flowing from using that data to train the AI – the "training gains".

One relevant dimension is the scale of the organisation's access to relevant data. Larger law firms or corporates are more likely to have access to larger volumes of relevant data; similarly, firms that specialise in a particular area are more likely than those that do not to get access to a larger volume of data in that specific area.

Mere access in principle does not mean that data is captured effectively. A central part of the "digital transformation" agenda in many corporates is to reorganise working and data capture, storage and transmission practices, and systems, so as to facilitate the capture of a higher proportion of potentially relevant data. However, in the legal sector, there still appears to be some way to go. Our survey of solicitors in England and Wales asked respondents whether they agreed with the proposition "my organisation captures data effectively, so it can be used by legal technology". As Figure Four (left) illustrates, the most common answer (36%) was "neutral", followed by "disagree" (28.4%), and "strongly disagree" (10.6%). Barely 3%

Figure Four: response by surveyed solicitors to the proposition: 'my organisation captures data effectively'

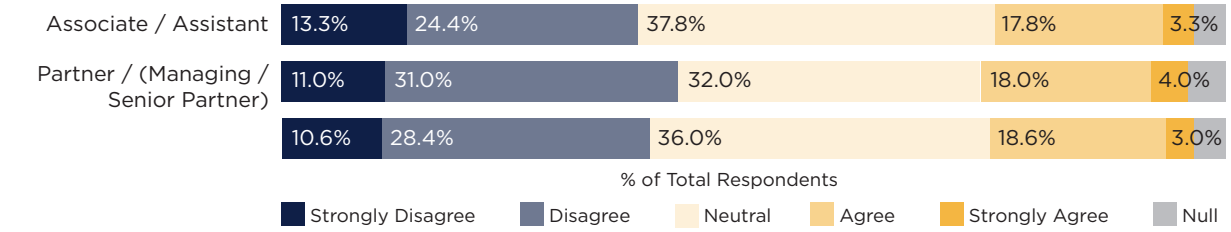

Source: Sako, Armour, and Parnham (2020)

*'Total' includes all complete responses from respondents working at Law Firms.

13 On some occasions, the distinction between the above-mentioned data categories can become blurred in a way that is relevant to AI lawtech tools. For example, it is known that the EDGAR database of company filings, maintained by the US Securities and Exchange Commission (SEC), has been used to assist with the training of AI-assisted due diligence tools. The information within this database was originally generated by companies, but then collated and made public by a state institution. In the UK, Land Registry data – typically generated by private organisations, but then collated and made public by the Land Registry – is known to have been used to train property-focused AI lawtech tools.

of survey respondents "strongly agreed" that their organisation captured data effectively.

A legal aspect of data capture concerns permission to use the data for training purposes. Commercial contracts do not generally contain personal data, so privacy issues are typically not a constraint. But there may be copyright or confidentiality restrictions if the firm does not own the data. This is not a concern for corporates with large proprietary datasets. But where data is supplied to a legal practice by its client, it is increasingly important for the firm to be clear as to what extent it is permitted to use the client data to train an AI-assisted lawtech solution. Can, for example, a practice use one client's data to train a solution, which might then be used by the practice to assist other clients – especially competitor clients in the same sector? If the practice is co-developing a solution with an AI lawtech solutions vendor, can the practice share client data with the solutions vendor – for example, when seeking to fix any data-driven bugs that emerge during the solutions testing phase?

How, and in what circumstances, legal practices were permitted to use client data was something of which many of our interviewees were aware, albeit to varying degrees.

> *"We've only used [client data] for their purposes at the moment, so we haven't started to use it in a wider sense or for other clients … So we probably haven't really thought about those purposes and actually what our restrictions might be … I think it will be something that we need to think about."*
>
> *Partner, large law firm*

> *"One answer is to build it into your engagement letters, [so] that we can use your data in an anonymised way going forward. But, with the absence of that, I mean, it's more a compliance question than a technical question, but can you use that? And, yes, that is a challenge."*
>
> *CTO, MDP*

In some situations – most notably eDiscovery – the question of whether one client's data should be used to train an AI lawtech tool for use on other matters does not arise, because the matter is very specific, and the model training obtained is unlikely to be transferable. Additionally, some of our interviewees had actively decided on a "no transferred learning" approach between clients as a matter of policy. Consequently, even if multiple client matters were similar in nature, the AI solution would be retrained each time, using only the client data that was relevant to the matter in hand.

> *"There's no real carry across from one case to the next."*
>
> *Partner, large law firm*

> *"If you are training on your own contracts and your own documents, it may be quite useless for another law firm who does it a slightly different way or has a slightly different format."*
>
> *Innovation director, Large law firm*

> *"All of the customer-specific training on their models is theirs."*
>
> *Innovation head, ALSP*

> *"If we don't have an explicit agreement in place, then the assumption is, [in our] contracts, the client owns the intellectual property. They own their insights and they … can take advantage of the insights derived from the AI."*
>
> *Digital solutions head, ALSP*

Others, though, were willing to discuss with clients the possibility of using their data for other matters – especially between matters that were likely to be repeatable between multiple clients. Perhaps unsurprisingly, some of these conversations resulted in discussions about who would benefit from using client data in this way. Sometimes, clients ask for discounts as a condition of data use. In other situations, clients took a more

"It is important for the firm to be clear to what extent it is permitted to use client data to train an AI-assisted lawtech solution."

benign approach, appreciating the wider benefits of using their data in combination with that of others – for example, when seeking to establish cross-industry norms.

> *"Clients are also realising the value of their data. So, if we say, actually, 'Can we anonymise and use your data?' naturally, the conversation then moves to, 'Well, am I getting a discount because you're getting value out of that?'"*
>
> *Innovation head, MDP*

> *"Some of them just normally go, 'Well, [that's our] data and we own it,' and we … go, 'OK … but if we have access to it … we'd be able to … [be] more consistent – we'd be able to do a better job for you,' and … that's the discussion."*
>
> *Head of client strategy, Large law firm*

> *"We've got consent from … at least three clients to be able to benchmark them against each other, on an anonymised basis. They'll know what their datapoint is, but … I don't think they know the identity of the other two. They're certainly not identified when they see the datapoints."*
>
> *Director of innovation, Large law firm*

As previously mentioned, legal practices sometimes help lawtech companies to train their AI models, in order to develop its out-of-the-box functionality. An additional question in relation to AI lawtech training is: should client-derived insights be submitted to the AI lawtech company, in order to help it expand its offering? Or should these learnings be retained for the exclusive use of the law firm and / or its own clients? For our law firm interviews, each approach had its own advantages. Where the learnings were retained within the law firm, the firm benefitted from having a solution capable of performing unique tasks. By contrast, where the firm was willing to share these learning with the AI solutions vendor, it was sometimes possible to secure a discount from the vendor for using the solution.

> *"Whenever we teach it [the tool] another change of control clause, here at [firm name], it reinforces and broadens, deepens its learning, and that's fed back to [lawtech company name] HQ, and then anyone else who buys [lawtech solution] gets the benefit of that increased learning."*
>
> *Innovation head, large law firm*

> *"So, the magic is all about how you, as a law firm … use that technology to … grow it out, your own instance, your own version, of that tool."*
>
> *R&D specialist, large law firm*

> *"It [the solution] has 30 clauses out-of-the-box … they're called 'public clauses': so, when you buy it, it will have 30 pre-trained bits of information – clauses – that it will recognise, and it can recognise quite a degree of variance for each of those clauses. And then we have … our own private [firm name] 'Secret Squirrel' clause library which we keep for our own … a couple of hundred extra clauses."*
>
> *Innovation head, large law firm*

In some situations, firms adopted a "mix and match approach" – sharing some insights with the AI lawtech solutions vendor, while also retaining other insights in-house.

> *"I think they [the lawtech software vendor] have two … versions in their licence. One is 'you … help us … train our existing model', and the other is 'you get our models as they are, but then you build your own version of that on top'."*
>
> *Innovation head, large law firm*

"Legal practices sometimes help lawtech companies to train their AI models, in order to develop its out-of-the-box functionality."

In these examples, it was generally obvious who owned the training data, and was therefore entitled to decide who should benefit from its use. But data ownership is not always clear cut. For example, in one scenario, the data that the client wished to have analysed did not initially exist in the form of structured data required to undertake the analysis. The necessary insights were scattered across a range of unstructured client-owned materials, including documents, witness statements, reports, and forms. As a preliminary step, the law firm had to review the unstructured data and turn it into a structured form, using its own personnel, expertise, and practice management system to do so. This process, in turn, begs the question: who owns, and has the right to exploit, this type of derivative dataset, where the underlying data belongs to the client but the dataset to be analysed was – effectively – created by the law firm? The law firm, the client – or both? We do not seek to answer this question here: rather, we raise it as a point of discussion, which those developing and deploying AI lawtech solutions should consider and cover in their documentation. We also suggest that is an area where guidance from professional and data regulators might be helpful.

In another scenario we encountered, a corporate client wished to have hundreds of its leases analysed. However, the law firm did not obtain these leases directly from the client, but rather downloaded them from HM Land Registry. This raises another question: if client data is already in the public domain, and available for download independently of the client's instruction, is client permission required to reuse those learnings when assisting other clients on future matters? Again, we do not seek to answer this question in this white paper – merely to raise it as an issue that firms may wish to consider at the outset of an instruction. Again, guidance from professional and data regulators on this point might also be helpful. ■

KEY TAKEAWAYS

- AI lawtech tools variously seek to source data from public or private sources. Some of this data will be available to all, and others will be only available to a limited number of market participants. Lawtech companies may be denied access to public data because they are commercial entities, rather than research organisations.
- For law firms, important training data for AI-assisted lawtech solutions are owned by their clients. Firms may find it useful to consider in what circumstances client data can be used to train AI models, and whether client permission is required. Firms may also wish to consider who the learnings from training AI models can be shared with – including other clients and lawtech solutions vendors.
- It may sometimes be uncertain who owns client-related data. Here, one area of uncertainty is where data is simultaneously available from private sources (i.e. clients) and public registers (i.e. HM Land Registry). Another area of uncertainty is where the data to be analysed is derived from client data but also created and structured by the law firm.
- A lack of availability of structured data constitutes a major barrier to AI deployments for law firms and other legal service providers.

CHAPTER FIVE

Possible impact of AI lawtech on law firms' recruitment patterns, training needs, and internal governance

CHAPTER FIVE

Possible impact of AI lawtech on law firms' recruitment patterns, training needs, and internal governance

In this chapter, we discuss the impact of AI on lawyers' future skills needs – indeed the nature of being a lawyer itself. We consider demand for technical skills in recruitment by law firms, and the nature of lawtech training currently received, and desired, by lawyers.

Is AI-assisted lawtech driving the legal profession to become more multidisciplinary in its nature?

As we saw in Chapter One, AI is increasingly being used in the legal services sector. In other parts of our research, we explored whether this is changing the skills expected of solicitors. Our findings suggest that, so far, changes in demand for lawtech skills have been modest. Using a very large dataset of the text of online job advertisements made available to us by Burning Glass Technologies,[14] we explored what fraction of positions advertised by law firms regulated by the Solicitors Regulation Authority over the period 2014–20 required skills closely associated with lawtech in general, or data science specifically. Figure Five (left) presents these findings, dividing the results into positions advertised for legal professionals (lawyers and paralegals) and those advertised for non-lawyers. The overall fraction of roles in which lawtech-related technical skills are sought is modest, but rising over the period reported in Figure Five. Clearly, the upward trend suggests a growing demand for legal-sector professionals with relevant technical skills. The modest overall level probably reflects in part the pattern discussed in Chapter Three, whereby law firms tend to work in partnership with third-party lawtech vendors to roll out AI lawtech solutions, rather than developing them in-house. The technical skills sought by law firms through the roles in Figure Five therefore likely reflect recruitment to build capacity in their legal operations units, which, as we saw in Chapter Three, are largely auxiliary to the firms' overall legal advisory business model. However, another – possibly complementary – explanation for the modest overall level of technical skill demand may be that the supply of these skills is as yet restricted through constraints on necessary education and training. We return to this issue below. Another clear

Figure Five: lawtech job adverts in the UK placed by solicitors' practices – legal professionals and non-legal professionals compared.

Note: legal professionals include solicitors, barristers and judges plus paralegals and legal secretaries
Source: Sako and Parnham (2021b)

14 The Burning Glass dataset is a very extensive data resource of job adverts posted online. Burning Glass Technologies, an analytics software company, scrapes job postings from the internet. Every day, the firm checks a corpus of more than 40,000 online job boards and company web pages to find new job vacancies. A limitation of this dataset is non-inclusion of roles that are not advertised online or are filled solely through recruitment consultants without advertisement.

finding from Figure Five is that where law firms are recruiting personnel with skills associated with lawtech, they are largely doing so into non-lawyer roles.[15] This has important implications for the way in which lawyers' working practices are being affected by AI-assisted lawtech. First, law firms appear not to be seeking significant technical skills for their legal professional roles. However, recall that approximately half of lawyers responding to our survey said they are now using AI-enabled lawtech in their work.[16] This implies that much of the engagement by lawyers in law firms with AI-enabled lawtech is as *consumers* of this technology – that is, the outputs of the lawtech augment the lawyers' legal advisory work.

Second, the fact that law firms are seeking technical skills for non-legal roles suggests that in contexts where law firms are producing outputs from AI-enabled lawtech – for example, by deploying it in the context of a legal operations unit to deliver efficiencies of scale – they are assembling multidisciplinary teams composed of individuals with distinct disciplinary backgrounds: legal expertise provided by legal professionals and technical expertise provided by professionals from other backgrounds. At the same time, it suggests that law firms are not expecting their legally qualified personnel who work in such teams to be what might be described as being "multidisciplinary individuals" (Janeček, Williams et al., 2021). Rather, the multidisciplinary emerges from co-operation within the team of individuals with different skillsets.

The modest direct impact of AI on the overall skills base within law firms was also reflected in our interviews. Although we typically interviewed some of the largest law firms in the UK, we were surprised to discover that – even among this type of firm – few employed more than 20 people in their innovation functions; even fewer were typically employed to assist with AI solutions rollouts.

> *"We've got 50 people in legal operations worldwide, of whom approximately half a dozen are specialists on commercial pricing, around about 12 to 15 are in the area of legal project management, three or four on data analysis innovation … and automation, we have a team of around about six or seven. Now, the maths probably won't add up to 50 because we have a number of individuals who … actually combine different skills … So, essentially, those are the five different components."*
>
> *Chief legal operations officer, large law firm*

Indeed, at one top 20 UK law firm, we discovered an example of an AI lawtech solution that had been developed entirely in-house with just two people – a practising lawyer and a software engineer / data scientist with a background in law. In another firm – a top 50 practice – it was considered noteworthy that the practice now employed a handful of data scientists.

Impact of AI on lawyers' skill needs – and professional identity

Working alongside professionals with different skillsets to deploy an AI lawtech solution allows all parties to play to their respective skills. But, in order for this to happen, lawyers need to be willing to accept what (for many) is a novel way of working. On this point, our survey findings paint a generally positive picture: as Figure Six (left) shows, when asked for their opinion on the statement "lawyers need to become familiar with multiple non-legal specialisms, such as data science, project management, and design thinking", a small majority of solicitors surveyed – 56.4% – agreed with this proposition. This compared with just under a quarter (24.6%) who were neutral, and around one

Figure Six: response by surveyed solicitors to the proposition: "lawyers need to become familiar with multiple non-legal specialisms, such as data science, project management, and design thinking."

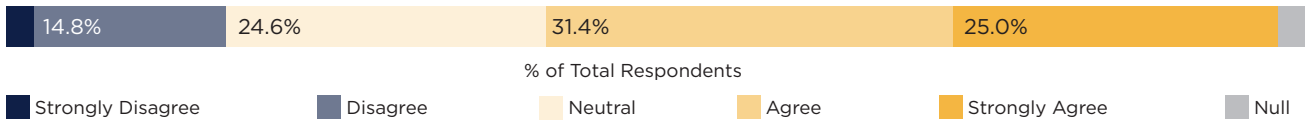

Source: Sako, Armour, and Parnham (2020).

15 Non-lawyer positions also account for the majority of posts recruited by law firms in the dataset.
16 See Chapter One.

in six (16.9%) who disagreed or strongly disagreed with this method of working.

This openness by many – but not all – lawyers to multidisciplinary working practices is also reflected in the changing nature of lawyers' training preferences, as revealed by our survey. As Figure Seven (below) indicates, lawyers' current technology-related training places a heavy emphasis on learning to use specific software packages. By contrast, it has placed almost no emphasis on skills closely associated with AI lawtech solutions development, such as data analytics and coding.

However, going forward, Figure Eight (page 32) indicates that a significant percentage of solicitors say they are willing to learn new AI-related skills in the near future (the next three years). If these intentions are acted on, the percentage of those who are familiar with software coding may rise from its current level of 3.4% to 27% during this time period. Even more remarkably, 73% of survey respondents said they intended to receive training in data analytics: currently just 1.1% of respondents are trained in this skill. This suggests that lawyers see value in their future professional careers through acquiring more advanced lawtech-related skills. These will likely unlock the ability for lawyers to move from being consumers of the outputs of lawtech solutions to producers of those outputs. Universities and professional education providers should take note of this demand.

If lawyers' skills transformation takes place on the scale that our survey indicates, one open question is whether the legal practitioners that operate in this way might continue to be regarded as "lawyers" in the traditional sense of the word. At the very least, such professionals might be regarded as being T-shaped lawyers or as "legal technologists" – that is, lawyers who, in addition to their deep legal domain expertise, also possess skills that are not normally associated with law, such as the ability to understand data (Janeček, Williams et al., 2021). Indeed, we have already seen an early example of the emergence of this form of hybrid lawtech professional in Scotland, where a specialist legal technology accreditation is now available from the local law society.

Where will lawyers who acquire such skills be able to deploy them effectively? The recruitment

Figure Seven: innovation-related skills training received by solicitors in the past three years

	Associate / Assistant	Partner / (Managing / Senior Partner)	Grand total
Software packages used by employer	49.4%	34.7%	42.4%
Project management	7.9%	11.2%	10.2%
Legal issues raised by the use of AI / technology	10.1%	8.2%	9.7%
Ethical issues raised by the use of AI / technology	5.6%	7.1%	6.8%
Innovation techniques	5.6%	7.1%	6.8%
Digital literacy	6.7%	3.1%	5.1%
Design thinking	3.4%	4.1%	4.7%
Data analytics	1.1%	5.1%	3.0%
Software coding	3.4%	2.0%	3.0%
Process re-engineering		5.1%	3.0%

Null	Associate / Assistant	Partner / (Managing / Senior Partner)	Grand total
49	89	98	236

*'Grand total' includes all complete responses from respondents working at Law Firms. ** Null indicates respondents who did not select any of the above options

Source: Sako, Armour, and Parnham (2020)

practices of law firms, as evidenced in Figure Five, suggest only modest current demand for such skills in law firm roles. However, there are a number of potential future directions. First, legal technologists may move from law firms to other legal services firms pursuing the new legal operations or legal technology business models we outlined in Chapter Three – ALSPs or lawtech startups. Second, they may work in corporate in-house legal teams, where lawyers already engage in more multidisciplinary working than those in law firms. Third, it is to be expected that demand by law firms for such skills will grow over time as these skills become more common among lawyers. ■

Figure Eight: solicitors' innovation-related skills training priorities over the next three years

	Associate / Assistant	Partner / (Managing / Senior Partner)	Grand total
Data analytics	73.0%	68.4%	71.2%
Legal issues raised by the use of AI / technology	73.0%	58.2%	65.3%
Software packages used by employer	61.8%	66.3%	63.1%
Ethical issues raised by the use of AI / technology	56.2%	42.9%	47.9%
Innovation techniques	40.4%	43.9%	43.2%
Digital literacy	42.7%	39.8%	41.5%
Project management	40.4%	29.6%	34.3%
Data thinking	30.3%	19.4%	26.3%
Process re-engineering	23.6%	20.4%	22.9%
Software coding	27.0%	12.2%	20.8%

Null	Associate / Assistant	Partner / (Managing / Senior Partner)	Grand total
49	89	98	236

*'Grand total' includes all complete responses from respondents working at Law Firms. ** Null indicates respondents who did not select any of the above options

Source: Sako, Armour, and Parnham (2020)

KEY TAKEAWAYS

- At present, legal sector job adverts that require lawtech skills in general, or AI lawtech skills in particular, only affect a small segment of the lawyer and non-lawyer roles, for either lawyers or non-lawyers working for law firms. For both lawyers and non-lawyers, the percentage of roles requiring lawtech skills has increased in recent years – but unevenly and from a very low base.
- Going forward, a far greater percentage of lawyers appear open to learning about data analytics than currently have those skills. This increase may encourage the greater adoption of AI lawtech, by allowing lawyers to envisage a broader range of scenarios when the technology may be deployed in the legal sector.
- Some law firms we interviewed are experimenting with their business models, to make them more conducive to the needs of those who pursue the unconventional careers paths associated with AI lawtech deployments. However, many more we interviewed had not yet started on this process.

Conclusions

This white paper paints a broadly positive picture of the impact of AI lawtech on the English and Welsh legal services sector. While take-up of the technology varies significantly by use case, overall usage is already high, at around 50%. Moreover, the introduction of the software is encouraging lawyers to work more efficiently, acquire new skills, and work with a more diverse group of people than previously.

What the technology does not appear to be doing, however, is prompting a radical shift in law firms' governance regimes and business models. Overall, the recruitment of AI lawtech specialists by law firms is very low, amounting to a fraction of 1% of advertised roles. Instead, much of the technical development work associated with AI lawtech is taking place within third-party providers, notably lawtech companies. On the one hand, this division of labour appears to be perpetuating the monoprofessional, lawyer-centric law firm governance and business model within law firms, arguably to the detriment of AI lawtech professionals who work for these organisations. On the other hand, this approach is also facilitating the creation of new legal market entrants, often supported by law firms via incubators, accelerators, and university partnerships.

In terms of possible problem areas associated with AI lawtech solutions deployments, this white paper has highlighted several key areas of concern. Most notably, access to and usage of the data required to train NLP-driven AI lawtech models may be a potential inhibitor of solutions development and deployment. In relation to data owned by the state, we observe that the UK government's historic preference for limiting data access to academic institutions may be inhibiting the development of certain AI lawtech tool types – for example, those which focus on predicting the outcome of disputes. We therefore advocate a change of government policy regarding access to publicly owned data. Rather than granting access on the basis of type of entity requesting the data, we suggest that the public interest in the intended use case should be the principal evaluation criterion.

More generally, we suggest that a lack of regulatory clarity regarding the circumstances in which legal practice lawtech companies can make use of privately owned third-party data to train AI lawtech models may also inhibit the usage of such tools in some situations. We therefore advocate greater clarity, from both data and legal services regulators, regarding when privately owned third-party data can be used by law firms and lawtech providers for this purpose. ■

"The introduction of the software is encouraging lawyers to work more efficiently, acquire new skills, and work with a more diverse group of people than previously."

Selected AI for English law research team bibliography

Aidinlis, S., Smith, H., Adams-Prassl, A. et al. (2020). *Building a justice data infrastructure: Opportunities and constraints,* University of Oxford. June 2020.

Aidinlis, S., Smith H., Armour J. et al. (2021). *Legitimate interests to share Justice data for Legal/Tech innovation in the UK,* University of Oxford.

Armour, J., Parnham, R., and Sako, M. (2020a). *Augmented Lawyering,* European Corporate Governance Institute. Law Working Paper N° 558/2020. December 2020.

Armour, J., Parnham, R., and Sako, M. (2020b). Unlocking the potential of AI for English law. *International Journal of the Legal Profession*, 28 (1), pp. 65-85.

Armour, J., and Sako, M. (2020). AI-enabled business models in legal services: from traditional law firms to next-generation law companies? *Journal of Professions and Organization*, 7(1), pp. 27–46.

Janeček, V., Williams, R., and Keep, E. (2021). Education for the provision of technologically enhanced legal services. *Computer Law & Security Review*, 40.

Petrova, A., Armour, J., and Lukasiewicz, T. (2020). Extracting outcomes from appellate decisions in US State Courts. *Legal Knowledge and Information Systems,* pp. 133–142.

Qian, M., Saunders, A., and Ahrens, M. (2020). Mapping LegalTech adoption and skill demand. In: Adams Bhatti, S., Chishti, S., Datoo, A. et al. (eds). *The LegalTech book*, Chichester: Wiley. pp. 211–214.

Sako, M., Armour, J., and Parnham, R. (2020). *Lawtech adoption and training. Findings from a survey of solicitors in England and Wales,* University of Oxford / Law Society. March 2020.

Sako, M., and Parnham, R. (2021a). *Technology and innovation in legal services: Annex report for the Solicitors Regulation Authority,* University of Oxford. July 2021.

Sako, M., and Parnham, R. (2021b). *Technology and innovation in legal services: Final report for the Solicitors Regulation Authority,* University of Oxford. July 2021.

Sako, M., and Qian, M. (2021). *A Taxonomy for Technology Venture Ecosystems*, SSRN. 8 April 2021.

Sako, M., Qian M., Verhagen, M. et al. (2020). *Scaling up firms in entrepreneurial ecosystems: fintech and lawtech ecosystems compared,* SSRN: University of Oxford. 15 January 2020.