JUSTICE IN THE DIGITAL ERA - MINUTES

Speakers

Lord Bach – Chair

Peter Chamberlin – Public Digital

Ellen Lefley – JUSTICE

Jasleen Chaggar – Big Brother Watch

Martha de la Roche – The Access to Justice Foundation

Mia Leslie – Institute for the Future of Work

Lord Bach – Chair

Lord Bach welcomed everyone in attendance and noted the relevance of the topic of discussion. He announced that the APPG on Access to Justice will be hosting a meeting with David Gauke on Wednesday June 18, starting at 14:30, to discuss the Independent Sentencing Review, and encouraged those in attendance to come.

Lord Bach then laid out the agenda for the panel event and introduced each of the speakers in turn. He explained that the panel would take a first round of questions after hearing from the first three speakers, before hearing from the final two speakers and concluding with another round of questions or discussion.

Peter Chamberlin – Public Digital

Peter thanked Lord Bach for his welcome. He began by noting that AI is a broad topic and a complex one, and that it is often difficult to stay grounded in the topic at hand. Peter outlined that he would be looking at AI in practice, drawing on examples from the Justice sector, setting out why and how AI is transformative, where we are in relation to the use of this technology at present and what steps we might take next to make the best use of this technology. Finally, he would make some concluding remarks about policy and regulation in the space

Peter noted that there is a lot of jargon in this space, and for most of us there is a fuzziness in our understanding. He set out that, rather than unpacking the technical specifics, he wanted to first discuss what these systems actually do in terms of:

- Comprehension: they comprehend natural language - you prompt them with language, and they make sense of it
- Reasoning: they can perform reasoning over the information they are given - this a technology you can feed a library of information, and it will provide an analysis of it. And do a good job.
- Scale: The technology does this at scale, and at speed: a scale greater than a human can - they can process far more information, much faster

- Tools: You can give AI access to the Internet, and it will work out for itself when and how to use it to respond to whatever instructions it is given

Peter then suggested that it helps to understand this in terms of information technology, like the telephone introduced instant communication over distance, and computing and the web have revolutionised our lives since the 1980s. He highlighted how successive generations of technology have changed the dynamics of information: the scale and speed of information creation, access, distribution, concluding that AI provides new capabilities to comprehend and reason, at unprecedented scale.

Turning to the emergent properties of AI in practice, Peter drew on an example from a published case study in the Justice sector: the Criminal Injuries Compensation Authority. He explained that the CICA has created a prototype, not a service in live production, which uses AI to triage cases, which:

- Process documents: categorising them, digitising them, converting them to useful formats
- Validate the case: is it complete? Are there missing elements, and if so, what?
- Verify it against policy: guidelines, compliance, consistency
- And then provides decision support to case workers: providing summaries and insights that support the human case worker in their decision making process

Peter summarised by saying that AI is being used to comprehend and reason, at scale and speed, using tools. In terms of outcomes, he recalled that they report good accuracy and that they also report that this has been done within CICA's information governance regime, so meeting significantly elevated data governance requirements with a potential 80% saving in time spent triaging cases.

Peter then turned to the questions posed by the CICA example, in particular, what does it mean if you automate away 80% of a person's role? He notes that it may be the case, in this instance, that this would improve the quality of work for those individuals, but automation does not always benefit workers. He highlighted a number of other questions:

- How do you know it's achieving what you intended it to?
- How do you measure the quality of these automated steps over time?
- How do you measure the way the automation colours the casework?
- What is the framework for evaluation?
- And considering that we know AI models drift, how do you monitor the quality of performance and outcomes over time?

Peter then emphasised the need to clearly define the purpose behind adopting AI in justice services. He highlighted that alongside identifying goals, it is crucial to develop robust systems for ongoing monitoring, evaluation, and governance. Peter noted that the justice sector has historically struggled in these areas, often lacking a strategic approach to understanding whether services effectively meet the needs of users and failing to manage or evaluate them over time. He stressed that the adoption of AI must be accompanied by broader institutional capabilities to support its integration.

To illustrate how AI can address practical barriers, Peter shared an example from the United States involving SNAP (Supplemental Nutrition Assistance Program), formerly known as Food Stamps. This welfare program provides food support, but like many similar systems, recipients

often encounter situations where their benefits are revoked or denied, requiring legal assistance.

Peter described how a private company called Propel created an app to help SNAP recipients manage their benefits. One of the key challenges faced by users was knowing where to seek legal aid, especially given that the legal aid system in the US is highly fragmented and lacks a centralized registry of providers.

In response, a Propel team member built an AI tool that automatically identifies and retrieves contact details for legal aid providers based on location. Peter pointed out that this tool was created quickly, with no significant cost, and provided a clear, tangible improvement: users were given direct contact details rather than vague instructions to "contact your local legal aid provider." Peter highlighted this as a practical and scalable application of AI that removed a significant barrier for vulnerable individuals.

Peter was careful to address the risks and limitations of AI. He emphasized that AI operates on statistical models, not deterministic rules, and so errors are inevitable. However, he noted that in the Propel example, the potential error—such as referring someone to a legal provider who couldn't help—was a low-impact mistake compared to leaving the person without any information at all.

Peter described this as a matter of "calibration to the impact of error", encouraging organisations to adopt AI solutions where the risk of error is acceptable and manageable. In more high-stakes scenarios, however, Peter cautioned that the consequences of AI failures could be significant. In these cases, policy, regulation, and guidance are essential in shaping ethical and safe AI deployment in the public sector.

Peter was clear in his stance that the use of AI is inevitable and that resisting it is not a viable option. He called on policymakers and practitioners to confront the challenges posed by AI, particularly around ethics, transparency, trust, and truth. These foundational concepts, Peter argued, must inform the regulatory and operational frameworks for AI use in justice and government. He acknowledged that other panellists would explore these issues further during the discussion.

Peter also pointed out that many people interact with government systems that still feel like they were built in 2005. Yet, people's expectations of digital services are shaped by the platforms and tools they use every day. Those expectations don't go away when they engage with public services, often they matter even more, especially when someone is in crisis or facing a complex problem.

Finally, AI continues to push the boundaries of what's possible, Peter warned that public services risk being left behind unless they modernise. He stressed that this isn't just about upgrading the tech, it's about recognising that this is a whole-of-organisation challenge, and leadership at all levels needs to respond.

Peter challenged the idea that public sector organisations can't take advantage of AI because they don't have the budget. In his view, the bigger issue is a lack of leadership and imagination. He encouraged public bodies to start small, pick real, concrete problems, and experiment with solutions that can scale.

Peter also praised the UK Government's AI Opportunities Action Plan, particularly its "scan, pilot, scale" framework. He said this kind of approach can help organisations learn, iterate, and improve in a way that's responsible and sustainable. Peter concluded by saying that the biggest variable in whether we make the most of AI is leadership.

Ellen Lefley – JUSTICE

Ellen explained that her organisation, JUSTICE, is a law reform charity committed to promoting a fair and equal justice system accessible to everyone in the UK. As the UK section of the International Commission of Jurists, the organisation has been active since 1957, with a dedicated branch, JUSTICE Scotland, established in 2011. The organisation brings together experts, conducts research, advises policymakers, and works to build public understanding in order to improve the justice system and reinforce the UK's commitment to human rights and the rule of law.

Ellen outlined two key themes in her remarks: first, how her organisation approaches the topic of artificial intelligence in the justice system; and second, the direction of future work and the big questions that lie ahead.

She described JUSTICE's approach to AI in the justice system as part of a wider, multi-year programme focused on AI, human rights and the law. However, she noted that this isn't the first time the organisation has considered the role of technology in justice. Over the years, JUSTICE has built a strong reputation for championing the modernisation of the justice system. For example, the organisation was deeply involved in the HMCTS reform programme, where it consistently emphasised the need to redesign digital processes around users, rather than simply digitising existing analogue systems.

JUSTICE has also made numerous recommendations across its reports for improving data collection within the justice system, to better understand who is accessing justice and whether that access is equal. During the COVID-19 pandemic, the organisation ran a fully virtual mock jury trial and partnered with academics to study its impact on procedural fairness.

Ellen stressed that the organisation's approach is always principled—focused on delivering a justice system that upholds human rights and the rule of law—but also pragmatic, advocating for workable solutions using available tools. AI, she noted, is just one more tool in the toolkit. However, its scale and the speed at which it is advancing are why JUSTICE has created a dedicated programme of work around it. The potential impact across all public systems is expected to be significant, and it is essential to get it right.

She pointed out that AI is a central part of the Government's vision for growth, with the Prime Minister expressing a desire for the UK to become a "great AI superpower." AI is also seen as a key mechanism for achieving efficiencies in public services, with the Labour leader recently describing plans to "mainline AI into the veins" of the nation by revolutionising public service delivery.

Justice, Ellen emphasised, is one of those public services. She explained that recent Government announcements, such as the AI opportunities action plan launched in January, are not the beginning of something new but rather a continuation of developments already well

underway. For instance, police forces have been using AI for years to enhance surveillance, allocate resources, assess risk, and triage emergency calls.

She highlighted that political figures like Shabana Mahmood have acknowledged the potential of AI to increase efficiency and productivity in the criminal courts, and that recent independent reviews, such as Jonathan Fisher KC's work on prosecution disclosure in the digital age, are exploring how AI can tackle delays and streamline labour-intensive tasks.

Ellen also noted that a recent survey suggests four out of five lawyers in the UK have already adopted, or plan to adopt, AI tools in their practice. These tools support a range of functions, from legal research and document review to drafting and litigation prediction.

Ellen then referenced existing judicial guidance on AI, which cautions against relying on it for legal research in unfamiliar areas but acknowledges its usefulness in reminding judges of familiar legal materials, summarising large volumes of text, drafting presentations and emails, and transcribing meetings.

Ellen spoke about the current integration of AI into the UK justice system, noting that while this development is well underway and not inherently problematic, the stakes are extremely high. The justice system, she emphasised, holds people's lives in its hands and plays a vital role in upholding both human rights and the rule of law. She observed that AI has the potential to support these principles, but it could also fundamentally undermine them. Neither outcome is inevitable, and society must take deliberate, responsible steps to shape the impact of AI.

She stated that this was the focus of JUSTICE's first report on the subject, published in January 2025 and led by Sophia Adams Bhatti. Ellen explained that this report was, in part, a response to the current direction of AI policy discussions in the UK, which have largely mirrored global trends by focusing on ethics. While ethical discussions have been useful in building global consensus, they often lack specificity and enforceability, and can be open to broad or even misleading interpretations. As Ellen pointed out, virtually every AI-producing company now makes commitments to AI ethics—yet the practical significance of these commitments remains questionable.

In contrast, JUSTICE advocates for the rule of law and human rights to play a central role in guiding AI development and deployment within the justice system. The report proposes a human rights and rule of law-based framework, intended to embed these principles throughout the design, development, and deployment of AI. The underlying belief is that if AI undermines human rights and the rule of law rather than upholding them, then the justice system itself risks losing its purpose and legitimacy.

Ellen described the framework as intentionally simple and practical, designed for use by policymakers, developers, and deployers of AI alike. It is built around three core steps:

1. Set clear goals for how AI can strengthen human rights and the rule of law. If the goals do not serve these purposes, Ellen asked, "what are you doing?"
2. Understand the risks—consider how AI might undermine these principles at technical, individual, and societal levels.
3. Act responsibly, by identifying how different actors, at different stages, can safeguard human rights and the rule of law.

Ellen summarised this as being goal-driven, understanding the risks, and taking responsibility, setting out three principal goals for AI that would advance human rights and the rule of law:

- Improving equal and effective access to justice;
- Improving the fairness and lawfulness of decision-making; and
- Helping open up the justice system to public understanding and scrutiny.

To illustrate the first of these, Ellen delved deeper into the goal of improving access to justice. She underscored its fundamental role in the rule of law and its status as an indispensable human right ensuring that individuals can vindicate their rights and defend themselves when accused. However, she noted that this is far from the reality for many in the UK today.

Cuts to legal aid have significantly reduced the availability of legal assistance. Ellen cited data indicating that while two-thirds of people in the UK will face a legal issue in the next four years, only half will have access to legal help. Additionally, backlogs in the courts, exacerbated by the COVID-19 pandemic, have caused widespread delays.

She highlighted that while those with financial means can still access legal representation and privately funded resolutions, others are left navigating legal aid deserts, complex rules, and an often alienating experience as litigants in person, or may abandon valid claims altogether.

Ellen then highlighted the clear potential for AI to help address these challenges. For instance, AI could streamline administrative tasks in courts and advice centres facing resource constraints. Translation tools and drafting assistance could support lay people and advisors. Of particular significance, she said, is the potential of AI systems, particularly those powered by large language models (LLMs)to democratise access to legal information. Chatbots using LLMs can interpret user prompts and generate answers across a range of legal topics. These tools are far cheaper than lawyers and available 24/7 and could therefore play a transformative role in improving access to justice.

Ellen continued her remarks by applying the JUSTICE framework to a practical example: an LLM-powered chatbot. At this stage of the framework, Ellen highlighted that the most obvious concern is the possibility of individuals receiving poor quality or even demonstrably incorrect legal information. This kind of misinformation has the potential to actively undermine access to justice, especially if people make decisions based on flawed outputs.

She noted several causes for this, including poor or incomplete training data, and referenced the now well-known issue of LLM "hallucinations." While it is often said that such technology will improve over time, Ellen urged scrutiny of any such claims. As an example, she pointed to overstatements made by tech developers who promote retrieval-augmented generation (RAG) as a solution that “eliminates” hallucinations or guarantees “hallucination-free” legal citations and notes that recent studies have effectively debunked this claim

Ellen also stressed that there are more fundamental challenges than just technical accuracy. In a common law jurisdiction like the UK, what constitutes “accurate” law is not always clear-cut. She explained that judgment data is not a definitive “ground truth”: judgments can be overturned on appeal and must always be interpreted within their specific legal and factual contexts. She noted that whilst some judgments are authoritative, others are not. Ellen warned that legal chatbots risk flattening these complexities, potentially ossifying a single interpretation of the law. If such tools are widely adopted, the impact could be broad—but if that impact plays

out primarily in settled and unreported cases, she asked, how would we even know it is happening?

Privacy risks were also flagged as a significant concern. Ellen referenced the recent data breach at the Legal Aid Agency, where reportedly over two million data points, including sensitive personal and financial information, were accessed by hackers, resulting in the system being taken offline. In her view, this underlines the need for robust data governance. She pointed out that direct-to-consumer legal chatbots carry additional risks, as users must navigate their own privacy protections—along with those of any other individuals involved in their legal issue—while simultaneously trying to understand the law.

Turning to the third and final stage of the JUSTICE framework—responsible action—Ellen made clear that this does not mean avoiding innovation simply because risks exist. On the contrary, a rights-based approach demands acknowledgement of the unacceptable status quo: deep inequality in access to justice and an urgent need for improvement.

However, she emphasised that responsible action does require taking the risks seriously. Existing governance mechanisms need to be stress-tested to determine whether they are fit to handle the challenges posed by LLM-powered legal tools. Currently, she noted, there is limited regulatory oversight of legal AI in the UK legal services market. Ellen suggested the need for strategies to incentivise good practice, such as regulatory sandboxes for legal innovators, and stronger, more proactive support from Government-backed initiatives like LawTech UK—particularly when it comes to tools aimed at improving public access to justice, not just equipping lawyers with more tech. (At this point, Ellen acknowledged Christina Blacklaws' leadership in advocating for this.)

She called for efforts to proactively establish standards and ensure access to reliable, high-quality data resources. Alongside this, Ellen argued that Government should already be thinking about how to monitor and evaluate the impact of AI-powered legal services on the wider market and access to justice more broadly. She posed an important question: how will we know where AI is making a positive difference, and where it is failing? Without attention to this, there is a real risk that people could find themselves excluded from both traditional legal services and emerging AI-enhanced alternatives. Legal support strategies, she concluded, must evolve to ensure that digitally excluded individuals are not left even further behind.

Ellen closed her remarks by posing a broader question: where are we going? She offered a provocation for future discussion, suggesting that there is a fork in the road ahead when it comes to access to justice. While there is real potential for AI to be a transformative force, and early signs of innovation such as the Contend app and the AI-powered firm Garfield (recently approved by the SRA), Ellen noted that it is clear the law tech market for lawyers is far more advanced in terms of development and adoption and that the reason for this is simple: it’s where the money is.

Ellen pointed to the fact that legal services for those with legal needs but limited means represent a failed market. She challenged the audience to consider whether we are heading toward a future in which those with ample resources continue to access justice—now supercharged by AI-enhanced legal teams—while those with few or no resources are left with low-quality tools or no assistance at all. Rather than narrowing the access to justice gap, she warned, AI could entrench and widen it. Ellen ended by asking: if that is not the future we want, what strategic interventions are needed to ensure we take a better path?

Jasleen Chaggar – Big Brother Watch

Jasleen began by speaking about the growing role of AI and technology in the 'pre-crime' space, particularly in relation to predictive policing tools. She explained that predictive policing uses data to try to anticipate crime and profile individuals based on the perceived risk of them committing criminal behaviour in the future. Jasleen warned that these tools have the potential to violate a wide range of rights, including the right to privacy, freedom of expression and association, the right to a fair trial, the presumption of innocence, and the right to equality.

As an example, Jasleen highlighted the use of the Harm Assessment Risk Tool (HART) by Durham Police. This algorithmic tool was designed to assess the risk of recidivism for individual suspects and inform decisions about whether someone should be charged and prosecuted, or instead diverted into a 'Checkpoint' rehabilitation programme. Jasleen pointed out that HART included postcode data as a factor in its risk assessments. This, she said, effectively acted as a proxy for sensitive data—such as an individual's race and socio-economic background—and risked reinforcing existing biases within policing. Jasleen noted that one in four individuals assessed by HART were categorised as 'high risk' of reoffending and were more likely to be prosecuted rather than offered rehabilitation. She also cited a study by the Metropolitan Police which found that this predictive tool was accurate in just 53.8% of cases.

Jasleen identified several fundamental problems with predictive policing. She argued that biased data entrenches existing racial inequalities in the justice system. Furthermore, in order for predictive tools to function, they require the ongoing monitoring, collection, and recording of personal data—even where there is no suspicion of crime. This, Jasleen stressed, contributes to the erosion of the presumption of innocence.

Turning to automated decision-making (ADM), Jasleen expressed concern about changes to the data protection framework which would expand the use of ADM in law enforcement and criminal justice. She reminded the panel that under current law—specifically sections 49 and 50 of the Data Protection Act 2018—solely automated decisions in a law enforcement context are prohibited unless expressly required or authorised by law. However, Jasleen noted that Clause 80 of the proposed Data Use and Access (DUA) Bill would reverse this safeguard. If enacted, it would permit solely automated decisions in all scenarios, unless the data being processed involves special category data.

Jasleen explained that in practice this would mean that automated decisions could be made about individuals based on factors such as socio-economic status, postcode or regional data, inferred emotions, or even regional accents. She warned that this significantly increases the potential for bias, discrimination, and a lack of transparency in the justice system. Jasleen also highlighted how non-special category data can still act as a proxy for protected characteristics when used in ADM—for instance, a person's name or occupation may be a proxy for sex, and postcode data can often function as a proxy for race.

Jasleen voiced particular concern over the weakening of safeguards in the law enforcement context. Under current legislation, individuals subjected to solely automated decisions must be notified of the decision and are entitled to obtain human intervention to contest it. Clause 80, however, would allow data controllers to ignore these safeguards where they consider ADM necessary for reasons such as public or national security. Jasleen stressed that, given the already minimal levels of transparency and the difficulties individuals face in challenging automated decisions, this would leave many people with no, or severely limited, avenues for redress.

Jasleen also addressed the use of AI-powered surveillance in policing. She highlighted recent government plans for mass surveillance of bank data to detect fraud and administrative error. These powers, she said, are being described as routine "data pushes," but in reality they represent quite extraordinary levels of surveillance.

She further raised the issue of live and retrospective facial recognition, which is increasingly being used by police forces in England and Wales. Jasleen noted that seven forces are now using live facial recognition technology on a regular basis, with others deploying it for major public events or as part of trials. Jasleen focused on the use of live facial recognition (LFR) technology in policing and surveillance, highlighting the serious implications for rights, privacy, and democratic participation in the UK.

LFR technology operates by matching faces captured on live surveillance camera footage against a pre-prepared police watchlist in real time. She noted that before any deployment, police create a watchlist which can include both police-originated images, such as custody photographs from the Police National Database, and non-police-originated images, such as publicly available open-source material or information shared by other public bodies.

Jasleen explained that LFR software detects human faces, extracts facial features, and converts them into biometric templates, and that these templates are then compared against those on the watchlist. The system generates a numerical similarity score to assess how closely a captured face matches any on the list. She continued, explaining that a threshold is set by the police, and matches that exceed this threshold are flagged. Officers may then choose to stop the individual in question.

Jasleen also drew attention to the growing use of LFR by private sector actors, including retailers, to manage issues that would ordinarily fall under criminal law, without sufficient scrutiny or judicial oversight. She noted that this represents a radical departure from long-standing UK policing principles. Traditionally, individuals are not required to identify themselves to the police unless there is a suspicion of criminal activity. However, LFR reverses this presumption, effectively treating every passer-by as a potential suspect subject to mass identity checks.

She described how the technology undermines the presumption of innocence. During an LFR deployment, everyone who passes the camera is subjected to biometric vetting by law enforcement, regardless of any individual suspicion. This is deeply problematic. As Jasleen argued, the police would never randomly stop members of the public to check their fingerprints against a database—so why is it acceptable to perform a similar function using automated facial recognition?

The technology threatens privacy not just at the individual level, but as a societal norm. Jasleen shared examples where LFR was used by police to target protestors—for instance, during an anti-arms fair demonstration in Cardiff, against climate activists at the British Grand Prix at Silverstone, and during high-profile events like the Coronation of King Charles. The presence of this surveillance technology has a chilling effect, deterring individuals from participating in lawful democratic activities. Jasleen framed this as an interference with Article 10 (freedom of expression) and Article 11 (freedom of assembly) rights.

While LFR is often compared to other biometric identification tools such as fingerprinting and DNA testing, Jasleen emphasised that it is far more invasive. It targets a deeply personal and public identifier, our faces, and is deployed without consent.

She noted that even though human officers are involved when a match is flagged, there is a worrying over-reliance on the software. In observed deployments, officers tended to place significant trust in the system, even when it produced errors.

Jasleen expressed concern about the broader implications of mass surveillance. LFR changes what it means to move through public space and interact with justice systems. The idea that police can obtain a biometric scan of someone's face without suspicion raises pressing questions about the direction of democracy and civil liberties in the UK. She highlighted that people modifying their behaviour due to surveillance is already happening, even in relation to lawful, democratic activities like peaceful protest.

She pointed to real-world consequences of false identification, citing the example of Shaun Thompson, who was misidentified and detained by police using the Metropolitan Police Service's LFR system, despite having done nothing wrong. The burden was placed on him to prove his identity and innocence which is an alarming reversal of legal norms. This, she argued, illustrates that the claim "if you've got nothing to hide, you've got nothing to fear" simply doesn't hold up.

Crucially, all of this is occurring outside a clear legal framework. Jasleen called for primary legislation, such as an AI Bill of Rights, to regulate the use of these technologies and prohibit the most invasive practices.

In closing, Jasleen noted that while a great deal of public and policy attention is being directed toward speculative risks of AI, real harms are already unfolding today. Technologies like LFR are quietly reshaping interactions with the justice system and inflicting unaddressed damage. She acknowledged that the EU AI Act is not perfect but at least reflects a willingness to engage with these issues. She concluded by arguing that the UK needs to demonstrate similar resolve.

Questions:

Lord Bach - Chair

Lord Bach thanked the panel for their contributions and opened the floor to audience questions.

Question 1 – Rebecca Wilkinson, LawWorks

Rebecca raised a point inspired by the recent LawTech UK conference. Referring to Jasleen's comments about bias, she asked whether we can realistically hope to build technology that doesn't replicate the existing biases in our legal system. Drawing on Ellen's earlier remarks, she noted that the legal system already contains significant human bias. Her central question was whether we are more comfortable accepting bias when it is human, but fearful of it when it is produced by machines?

Question 2 – Chris Minnoch, Legal Aid Practitioners Group:

Chris directed his question towards Peter Chamberlin's comments on AI as another form of information technology. His concern focused on data, specifically, where AI-generated data resides and who holds it. Citing a recent data breach at the Legal Aid Agency, he highlighted the risks posed by sensitive information being compromised. He asked whether AI introduces an additional layer of risk regarding data storage and ownership, and whether AI systems themselves generate new categories of data based on client inputs, further complicating data protection.

Question 3

The third question concerned the structural challenges faced by charities working directly with government to provide access to justice. The person who asked the question observed that while AI and law tech solutions have predominantly developed in the B2B space, the most pressing need is in the B2C space, where ethical, accuracy, and risk issues are amplified. They emphasised that many individuals seeking legal support are unaware of legal changes affecting them and often rely on non-commercial or state-charity partnerships to navigate these issues. They asked the panel to reflect on how we can better support the development of B2C technology, particularly for organisations that aren't AI experts but are required to manage its associated risks on behalf of vulnerable consumers.

Peter Chamberlin – Public Digital:

Answering the first question on whether it is possible to create unbiased technology, Peter expressed doubt that any system could ever be entirely free of bias. He emphasised the need to measure and act upon bias within systems, arguing that traditionally, we haven't done enough to evaluate how effective or fair our systems are, particularly in terms of outcomes. He stated that his belief is that the focus should shift toward continuously measuring and improving technology. Peter noted that it is not about achieving perfection but about establishing mechanisms that allow for ongoing evaluation and enhancement.

Ellen Lefley – JUSTICE:

Ellen agreed with Peter's remarks, stressing that machine bias is, at its core, human bias. She pointed out that humans create and train machines, using data that is itself shaped by human decisions and structures. Separating machine bias from human bias, she argued, can reinforce one of the most significant issues in this area: the misplaced trust in machine outputs as being neutral or objective. She stressed the importance of users maintaining a critical perspective

when interpreting AI outputs and called for robust evaluation and monitoring systems to ensure fairness and accountability.

On the intersection of B2C technology and data privacy, Ellen linked her response to the concerns raised in both the second and third questions. She observed that the privacy stakes are particularly high in the B2C context, where individuals often seek help in moments of vulnerability, lacking both resources and legal support. In these situations, they are expected to trust that their personal and sensitive data will be safeguarded—not just their own, but also that of others potentially involved in a legal dispute.

Ellen advocated for the creation of communities of shared good practice and pointed to the potential of regulatory sandboxes in this space. She also highlighted the importance of government-backed initiatives that support responsible innovation. Encouraging innovators who are committed to solving the access-to-justice challenge in the B2C space should, she said, be a central part of the policy response.

Jasleen Chaggar – Big Brother Watch:

Responding to the earlier question on bias, Jasleen stated that it is right we are terrified of computer bias. She suggested the key distinction lies in the fact that when a human makes a decision, it is often possible to ask them why they reached that conclusion and to interrogate the reasoning. This is in contrast to algorithms, where the decision-making process can be opaque. She stressed the importance of transparency in algorithmic systems, particularly within the public sector, where decisions made about citizens need to be understandable. Transparency, she argued, allows for the kind of measurement and analysis that Peter Chamberlin had referenced earlier in the discussion.

Jasleen also addressed the question of who holds the data, noting that in the public sector especially, it is important that individuals retain control over their own data and how it is used. She underlined the principle of purpose limitation as a key safeguard to ensure that when a person shares their data for one specific purpose, it should not be reused for another, unrelated one.

Peter Chamberlin – Public Digital:

Peter added a technical clarification in response to the question of data and AI systems. He began by acknowledging the incident mentioned earlier involving a data breach at the Legal Aid Agency. He explained that the framing of the question depends on the specific system in use. AI systems, he noted, are built on data and also generate new data. There are, as he pointed out, ongoing questions within the UK's creative industries around data sources and training data that he did not delve into at this stage.

To illustrate, he offered the example of AI systems deployed via commercial cloud providers such as Amazon Web Services or Microsoft Azure. He suggested that tools like ChatGPT usually offer assurances that they will not retain or use the customer's data for model training. However, Peter cautioned that there remains some ambiguity around how thoroughly these assurances are tested and how well-evidenced the claims truly are.

On the question of whether AI creates an additional layer of data, Peter affirmed that it does, but it depends on how the system is used. He explained that any information system will generate new data. For instance, in a case like CIC, systems may generate new summaries or outputs based on original source data. If the primary data includes personal or special category data, the AI-generated outputs can also constitute new data derived from that personal information. Peter concluded that, in more automated implementations, AI systems can indeed be acting on and producing further data from the original personal input data, and this will often be the case in practical applications of such models.

Lord Bach - Chair

Lord Bach thanked the audience for their questions and continued with the rest of the meeting.

Martha de la Roche – Access to Justice Foundation

Martha opened by highlighting the importance of centring user needs, particularly for marginalised communities, including people experiencing digital exclusion. She stressed that any innovation in this space must be guided by the needs of those most affected by the justice gap and warned against the risk of further excluding groups who already face significant barriers.

She spoke about the Justice and Innovation Group, describing it as a small, engaged group from across the sector that has been focusing on small AI that function as applied, pragmatic technologies rather than large-scale systems. She noted that the group has produced a JusticeTech tools list, mapping existing technologies in the space. Martha noted that the emphasis has been on AI used for service delivery, to help more people get access to justice, rather than purely for organisational efficiency.

Martha outlined that this has significant potential because, with the right tools and approaches, the sector could serve 100,000 more people a year with the same resources. She clarified that this isn't about replacing services but about improving frontline capacity, enabling services to reach more communities in need, and prioritising and preserving face-to-face support for those who most need it.

She acknowledged the real challenges, particularly around infrastructure, pointing to the systems at the Legal Aid Agency (LAA, as well as the need for upskilling, and the difficulty of scaling and sustaining promising solutions. Martha stressed the need to bring together tech partners, pro bono support, government, and frontline organisations, not just to improve delivery but to manage risk. She emphasised the importance of co-creation and design with the sector as the only credible way forward.

On resourcing, Martha stated that there has been an inadequate market response to the needs of the sector. She argued that more targeted funding and investment is required, with a focus on solutions that work across the ecosystem—not just for businesses, but for legal aid and other free legal advice providers. She noted that there is an ambition to establish a JusticeTech fund, but this will require dedicated funding to get off the ground.

Mia Leslie – Institute for the Future of Work

Mia began by highlighting the significant and cumulative impacts that the introduction of AI and emerging technologies are having on the workplace and the individuals within it. She explained that these impacts are being felt not only in the nature, conditions, and quality of work, but also in employment relationships, the broader organisation of work, and even society as a whole.

According to Mia, at an organisational level, choices around how AI is designed, developed, deployed, and governed are reshaping business models and employment structures. These choices are enabling new ways of contracting and organising work and are fundamentally shifting the foundations of workplace management. Mia noted that this shift is also evident at the individual level, where AI technologies are being used to automate, augment, transfer, or intensify specific work tasks. She pointed out that AI is not only replacing or enhancing tasks but is also increasingly being used to reconfigure core management functions.

Mia emphasised that while some of these changes are visible and explicit, many are more subtle. AI systems are often integrated into existing workplace processes, tools, or IT infrastructures in ways that can be difficult for workers to detect or understand. This, Mia warned, can mask their real effects and make it harder for individuals to know how their roles are being influenced by new technologies.

Shifting her focus to the individual level, Mia stressed that this is where most people experience access to justice in the workplace. She raised concerns about workers' ability to understand the basis for changes they experience in their roles. Mia explained that monitoring and tracking systems may produce inferences used to measure performance, determine pay and benefits, allocate work, or even hire and fire employees. While the most severe consequence may be job loss, Mia made clear that there is a broad spectrum of possible impacts on working conditions and access to work.

Mia noted that when workers are not actively involved in decisions about which systems are introduced and how they are implemented, this results in a serious information imbalance. Workers may not know what data is being collected, how it is analysed, or what it is used for. Mia warned that this poses a serious risk of unfair or discriminatory outcomes, particularly when algorithmic decisions are used to make high-impact employment determinations. Mia highlighted that without transparency or meaningful consultation, workers face major obstacles in challenging decisions made about them.

To address these issues, Mia outlined a series of proposed amendments to the Employment Rights Bill (ERB). These include stronger oversight and safeguards for high-risk ADM in the workplace, mandatory AI risk assessments, legal duties to consult workers on AI systems, a right to explanation, enforceable remedies, a right to human reconsideration of algorithmic decisions, and the creation of an AI systems register.

Mia also drew attention to the particular challenges facing workers in the gig economy. She emphasised that these workers are often entirely subject to algorithmic management, with every aspect of their work determined by platform systems. Mia explained that these individuals face even greater barriers to accessing information about how platforms collect, use, and share their data which is the same data that underpins the decisions affecting their work.

Mia cautioned that upcoming changes in the Data Use and Access Bill are likely to further disadvantage vulnerable gig workers. She warned that these changes could increase the conditions under which ADM is permitted and make it even more difficult for self-employed platform workers to legally challenge breaches of their rights. Mia added that the ERB is unlikely to address the issue of platforms exploiting bogus self-employment classifications to deny workers basic protections.

In closing, Mia underlined the critical need for ongoing systems to monitor and surface the impacts of AI in the workplace. She referenced the Good Work Algorithmic Impact Assessment as a promising approach. Mia explained that this framework is specifically designed to assess the effects of adopting new technologies at work. Importantly, it involves workers at every stage: from identifying risks and opportunities, to shaping mitigation strategies, to participating in the continuous monitoring of these systems. Mia emphasised that this kind of proactive, participatory oversight is essential if AI is to be introduced in ways that are fair, transparent, and just.

Questions:

Lord Bach – Chair:

Lord Bach reopened the floor to questions, noting that approximately fifteen minutes remained for discussion. He invited audience members to pose questions, asking that they introduce themselves before doing so.

Question 4 – Tony Guise, LawtechUK

Tony addressed the panel and thanked the speakers for their presentations. He directed attention to the Online Procedure Rule, authorised by Parliament, which is creating a new digital justice ecosystem. He noted that this system brings together many of the themes discussed during the event, such as a B2C model aimed at resolving disputes and helping individuals to understand their legal challenges and identify available resources—whether through platforms such as Advicenow, solicitors, or mediators.

Tony noted that this initiative was introduced via the Judicial Review and Courts Act 2022, and that a committee and sub-committee were fully formed by July 2024, and the rule was recently authorised through a statutory instrument to begin setting rules in relation to property and possession disputes, particularly under the Renters' Rights Bill.

He highlighted that this ecosystem will necessarily engage with questions around the use of artificial intelligence, particularly in terms of how AI might help individuals to identify their legal issues. Tony then invited the panel to respond to how they have engaged with the Online Procedure Rule.

Question 5 - Amit Kohli, Access Social Care

Amit explained that his organisation provides AI-driven assistance to help individuals resolve legal issues, while rigorously observing data protection and security procedures. He noted,

however, the limitations that arise in the face of widespread public use of generative AI tools like ChatGPT. Despite the carefully tailored assessments carried out by Access Social Care, citizens frequently turn to large-scale AI systems for advice.

Amit pointed to a wider contradiction: while his organisation takes every precaution to protect personal data, individuals often share highly sensitive information freely across public platforms such as Facebook, TikTok, and Instagram.

He expressed concern that discussions around ethical AI deployment often remain at a high level, without adequate attention paid to public education. He stressed the need for foundational education systems that empower citizens to understand their rights, use digital tools responsibly, and critically assess what information sources to trust.

Question 6 – Dr Tom Grange, Clinical Psychologist and Expert Witness

Dr Grange shared his recent experience working with expert witnesses across various legal contexts. He recounted being unsettled by a report in which AI had generated responses to expert instructions. While the result was not perfect, it was, in his view, impressively competent.

Dr Grange raised a broader concern: what happens when AI systems become demonstrably better than human professionals (such as judges, lawyers, psychologist) at carrying out their roles? He posed that, in an already constrained justice system where access is limited, such developments could undermine some key aspects of our society.

Question 7 - Ray Smith, Inside Time Newspaper

Ray highlighted that people are fundamental to our justice system—from victims and defendants to lawyers, juries, and judges. He warned that the increasing reliance on AI risks deskilling the human professionals within this ecosystem.

Focusing on probation, Ray described how probation officers once had the time and experience to assess patterns and use professional judgment in complex cases. Today, however, the pressures of the job have changed. Officers now often defer to AI-generated outputs when making decisions, particularly around recall procedures. This shift, in his view, undermines both the role and the motivation of those working in probation.

He expressed concern that AI risks making crucial decisions based on algorithms, rather than human assessment. Ray concluded by warning of the potential for misuse, including breaches of confidentiality, system hacks, and malicious actors exploiting AI capabilities. While acknowledging that AI cannot be stopped, he emphasised that its deployment must be guided by common sense and the personal integrity of the human professionals overseeing it.

Question 8:

Following Martha's earlier focus on local-level advice providers, the final question focused on the variability in local provision. The person who asked the question noted that, in some areas,

there are complete legal advice deserts where provision comes mostly from independent services, and different organisations vary greatly in their complexity and emphasis. The panel was asked how their thinking could address this highly variable local landscape.

Martha de la Roche – Access to Justice Foundation

Martha began by addressing the issue of variability. She noted that technology, more generally, does offer the opportunity to support better collaboration and coordination at the local level of advice services. She gave the example of developing a tool that could, based on the information you put into it, signpost you to your best chance service or a good kind of success service.

She cautioned that what technology can't do is create a specialism where one doesn't already exist. Particularly with AI, she stressed that we need to look closely at the data it's using and the expertise it's drawing from. Martha highlighted that the sector is facing a huge workforce crisis in free legal advice provision, especially among housing legal practitioners. There are areas across the country that are, effectively, legal advice deserts with little or no housing advice provision.

Her key point was that we can't use AI to sit on top of a non-existent specialism and expect it to replicate something that isn't already there. She emphasised the importance of continuing to invest in specialist, expert provision, because without that, there's no underlying intelligence for an AI or tech-based tool to build on or democratise. She concluded by stating that if it doesn't exist, it can't be automated or scaled.

Jasleen Chaggar – Big Brother Watch

Jasleen responded to the part of the question about how to get people to care about privacy and be more cautious about handing over their data. As a privacy rights organisation, she said, this is a challenge they face often. She noted that when people first signed up to social media services for free, they didn't realise that these platforms weren't really free, and that users were, in fact, giving up a lot more than they were receiving in return.

Over time, she said, as more scandals emerged and more became known about how our data is used, public awareness of what it actually means to give up your data has grown. Jasleen was clear that there's no simple answer to the question but that the work continues.

Ellen Lefley – JUSTICE

Ellen picked up on Jasleen's point, noting that we are now a bit more knowledgeable about the fact that tech is buggy, and that blind trust is probably in decline especially in light of the Horizon scandal. She pointed out that while Horizon wasn't artificial intelligence in the way we talk about it now, it was still complex software, and it showed us what can happen when systems are trusted too much without sufficient scrutiny.

She reflected more broadly on the nature of the discussion, saying that this is often the problem with talking about AI, that It depends entirely on the deploying context.

For example, a smart and ethical lawyer might use an LLM chatbot to double-check their understanding of a case. If it says, "Yes, you're right," then they use it — potentially saving time and money for clients. That's very different, she said, to someone in a state of panic, outside of the scope of legal aid, inputting information about a child arrangement issue into a chatbot and hoping for answers.

Ellen argued that what we need are context-specific risk management assessments, because the reality is that AI is not one single thing. The way the state uses it is very different to how it might be used to empower marginalised communities that have been excluded from access to justice.

Peter Chamberlin – Public Digital

Peter began by saying he would try to answer all the questions quickly. He said there's real value in tools like ChatGPT for refreshing your memory or checking your understanding and emphasised that if people are using these tools anyway, then these systems must be designed with expertise and care.

He reiterated that we have no option but to engage with this technology, and that this calls for leadership, clarity, and purpose in the public space.

In response to Dr Grange's question, he said it's an important question for legislators and public policy makers. The nature of work is at risk of radically changing under people's feet, in ways they didn't expect, and this needs to be taken seriously.

Coming back to the issue of local-level advice, Peter summarised that this and the other questions really point to a broader need for leadership. He was clear that this is not a technology problem but a social problem, a political policy problem, and a regulatory problem.

Mia Leslie – Institute for the Future of Work

Mia responded to the question about AI capabilities, and the question regarding the (dis)incentives for developing skills.

She referred back to the start of her presentation, saying that while the development of technology may feel outside of the control of those in the room, many decisions are still within our hands.

When it comes to roles and the workplace, she said, there are different ways the exact same technology can be implemented, and the outcomes for workers can be very different. She gave the example of high-discretion versus low-discretion augmentation.

High-discretion augmentation is where a person is still making all the high-level decisions, using the technology simply to cross-reference or offer suggestions. Low-discretion augmentation is the more concerning model, where decision-making shifts away from someone who has trained and qualified to do the job, and instead rests with an AI system.

She stressed the importance of having a detailed and nuanced conversation, and not just focusing on what the technology can do, but about how we want or don't want to use it.

Lord Bach – Chair

Lord Bach thanked everyone for their engagement with the event and the speakers for giving up their time. Lord Bach then reiterated the announcement made at the beginning of the event, inviting everyone to attend the upcoming meeting with former Lord Chancellor, David Gauke, on the launch of the Independent Sentencing Review. Lord Bach thanked LAPG and Rohini Jana for putting the event together.