CONFIDENTIAL COMPLIANCE DOCUMENT

AI Policy — Essentials

Prepared exclusively for Brightwell Accountancy

UK Jurisdiction · 08 August 2026

CONTENTS

What is in this pack

Each document below is also attached to your delivery email as its own PDF and editable Word file, so you can circulate or sign any one of them on its own.

Document	What it is for
AI Acceptable Use Policy	Your formal AI policy. Sign it, circulate it to every member of staff, and review it annually or whenever your AI tool usage changes.
Plain-English Employee Summary	A one-page summary of the AI policy that staff will actually read. Print it, pin it up, hand it out at induction.

PART 1

Your AI Compliance Profile

Brightwell Accountancy is a UK-based accountancy and bookkeeping firm with 11–50 employees, handling payroll and self-assessment data for approximately 400 clients. Staff are actively using ChatGPT, Microsoft Copilot, and Otter.ai for drafting correspondence, meeting summaries, and research — in many cases without organisational approval or oversight. This means personal data, which is likely to include sensitive financial and payroll information belonging to identifiable individuals, is being processed via third-party AI systems without a documented lawful basis, data processing agreements, or staff governance controls. The firm currently has no AI policy in place, placing it in a materially non-compliant position under UK GDPR and exposing it to significant regulatory risk from the ICO, whose 2024/25 enforcement strategy explicitly prioritises AI use and unlawful data processing.

Maximum Fine Exposure

Under UK GDPR Article 83(5), the most serious infringements — including processing personal data without a lawful basis, failing to implement adequate technical and organisational measures, and breaching international transfer rules — attract a maximum penalty of £17,500,000 or 4% of global annual turnover, whichever is higher. The statutory maximum in absolute terms applicable to Brightwell is therefore £17,500,000. The turnover-based limb (4% of global annual turnover) cannot be computed here because Brightwell has not disclosed its annual turnover figures; if 4% of Brightwell's actual global turnover exceeds £17,500,000, that higher figure would apply, but this is unlikely for a firm of 11–50 employees and should be confirmed against Brightwell's accounts. For lesser infringements under Article 83(4) — such as failures in record-keeping (Article 30) or failure to execute Data Processing Agreements (Article 28) — the ceiling is £8,750,000 or 2% of global annual turnover, whichever is higher. In practical terms, Brightwell faces concurrent exposure across multiple Article 83(5) heads simultaneously: processing client payroll and self-assessment data via unapproved AI tools without documented lawful bases, without Data Processing Agreements with vendors, and without UK GDPR-compliant international transfer mechanisms for US-based tools (ChatGPT/OpenAI and Otter.ai). Each failure is a discrete infringement. The ICO has discretion to treat related failures as a single infringement or as separate matters, and its 2024/25 enforcement strategy confirms that AI-related data protection failures are an active priority. The firm's exposure is therefore material and multi-limbed, with a worst-case statutory ceiling of £17,500,000 for the most serious breach — a figure that dwarfs the cost of implementing a compliant AI governance programme.

AI Tools Inventory

Tool	Use	Risk Level	Notes
ChatGPT	Drafting client correspondence	**High**	Staff are using ChatGPT on their own initiative without approval or a governing policy. Client correspondence for an accountancy firm routinely involves names, addresses, tax reference numbers, income figures, and payroll data — all personal data under UK GDPR. Inputting such data into ChatGPT (operated by OpenAI, a US-based third party) without a Data Processing Agreement, a UK GDPR-compliant international transfer mechanism, and staff training creates a serious compliance gap. There is no evidence of a lawful basis being documented for this processing. The risk of data being used to train

Tool	Use	Risk Level	Notes
			OpenAI models unless enterprise controls are in place adds further exposure. Hallucination risk in client-facing correspondence also creates potential professional liability.
Microsoft Copilot	Research and analysis	**Limited**	Microsoft Copilot, when accessed via Microsoft 365 commercial licensing with a Data Processing Agreement in place, offers materially stronger data protection controls than consumer-grade AI tools — including commitments not to use customer data to train foundational models. However, risk remains 'Limited' rather than 'Minimal' because Brightwell has not confirmed whether the appropriate commercial licensing tier is active, whether staff are using the consumer or enterprise version, and whether a Data Processing Agreement has been formally executed. If staff are using a personal or free-tier Copilot account, risk escalates to High. Immediate verification of licensing status and contractual documentation is required.
Otter.ai	Summarising meetings	**High**	Otter.ai transcribes and summarises audio recordings of meetings, which in the context of an accountancy firm will frequently capture personal data — including client names, financial details, payroll figures, and self-assessment information — spoken aloud. Otter.ai is a US-based service. Processing personal data via Otter.ai requires a valid Data Processing Agreement and a UK GDPR-compliant international data transfer mechanism (such as the UK International Data Transfer Agreement). Meeting participants — including clients — must be informed that AI transcription is in use, as covert recording of conversations containing personal data is likely to breach the transparency requirements of UK GDPR Articles 13 and 14. No such controls appear to be in place. The risk of transcription data being retained on Otter.ai servers and used for service improvement is a further concern requiring contractual verification.

Applicable Regulations

UK GDPR

The primary data protection framework governing Brightwell's processing of client personal data, including payroll and self-assessment information. Key obligations include: identifying a lawful basis for each AI-enabled processing activity (Article 6); ensuring transparency with data subjects about how their data is processed, including via AI tools (Articles 13–14); maintaining Records of Processing Activities (Article 30); executing Data Processing Agreements with all AI tool vendors acting as processors (Article 28); and implementing appropriate technical and organisational measures to ensure data security (Article 32). Where AI tools are based outside the UK, international transfer rules apply and a transfer mechanism such as a UK International Data Transfer Agreement must be in place (Articles 44–49). The ICO's AI Auditing Framework operationalises these obligations specifically in the context of AI systems and should be treated as authoritative guidance for Brightwell's governance programme.

Key date: In force. The ICO's 2024/25 enforcement strategy, launched 5 June 2025, explicitly prioritises AI and unlawful data processing as enforcement priorities. · Max penalty: £17,500,000 or 4% of global annual turnover, whichever is higher (for the most serious infringements under Article 83(5)–(6)). Lesser infringements attract up to £8,750,000 or 2% of global annual turnover.

UK GDPR Article 22 — Automated Decision-Making

Article 22 restricts solely automated decisions that produce legal or similarly significant effects on individuals without human involvement. Brightwell has confirmed it does not currently use AI for decisions about people, which reduces immediate Article 22 exposure. However, as AI use expands — for example, into automated payroll calculations or client risk assessments — Article 22 obligations will engage. The Data (Use and Access) Act 2025 (see below) has modified and expanded the circumstances in which automated decision-making is permissible, so Brightwell must monitor this area as its AI use matures. The ICO has specifically flagged automated decision-making in employment contexts as an area of 'real risk of harm' warranting increased scrutiny.

Key date: In force. Modified by the Data (Use and Access) Act 2025, which took full effect on 5 February 2026. · Max penalty: Enforcement falls under the UK GDPR penalty regime: up to £17,500,000 or 4% of global annual turnover for the most serious breaches.

Data (Use and Access) Act 2025

This Act took full effect on 5 February 2026 and introduces reforms to UK data protection law with relevance to Brightwell's AI activities. It expands the circumstances in which solely automated decisions can be made about individuals, potentially liberalising some AI use cases that would previously have required explicit human review under Article 22. It also clarifies that 'scientific research' can encompass commercial research and technological development, and permits broader consent for such processing — relevant if Brightwell or its AI vendors rely on research-based lawful bases. While broadly AI-friendly in orientation, the Act retains essential data subject safeguards. For Brightwell, the practical implication is that its AI governance documentation must be updated to reflect the post-February 2026 legal position, and any vendor data processing agreements should be reviewed for alignment with the Act's provisions.

Key date: Took full effect 5 February 2026. · Max penalty: Enforcement and penalties operate through the UK GDPR framework as modified by this Act. The maximum penalty position under UK GDPR therefore applies: up to £17,500,000 or 4% of global annual turnover.

ICO AI Auditing Framework

The ICO's AI Auditing Framework sets out how the ICO assesses whether organisations using AI systems comply with data protection law. It is not a standalone legal instrument but represents the ICO's authoritative interpretation of how UK GDPR obligations apply to AI — and it is the lens through which any ICO investigation of Brightwell's AI practices would be conducted. The Framework expects organisations to: maintain an inventory of AI tools in use; assess and document the data protection risks of each tool; implement Data Protection Impact Assessments (DPIAs) where processing is likely to result in high risk (which the unsanctioned use of generative AI with client financial data almost certainly is); ensure human oversight is meaningful; and govern staff use of AI through documented policies and training. Brightwell's current position — no policy, no oversight, unsanctioned tool use, no DPIAs — represents a significant gap against every dimension of this Framework.

Key date: Active and applied by the ICO in current enforcement activity. The ICO's 2024/25 strategy launched 5 June 2025 confirms AI auditing as a live enforcement priority. · Max penalty: Not a source of independent penalties — enforcement is under UK GDPR. Non-compliance with the Framework's expectations is a material factor in ICO investigations and penalty calculations.

UK Equality Act 2010

The Equality Act 2010 is relevant to Brightwell's AI use to the extent that AI tools used in any employment-related context — including HR correspondence, staff communications, or any future people-related decisions — must not produce outputs that discriminate against individuals on the basis of protected characteristics (including age, sex, race, disability, religion, and others). The ICO has specifically flagged AI use in employment and recruitment as an area of heightened scrutiny. While Brightwell has confirmed it does not currently use AI for decisions about people, any AI-generated content touching on employment matters must be reviewed for discriminatory outputs before use. As AI use expands, a documented equality impact assessment of AI tools in scope may be prudent.

Key date: In force. No AI-specific commencement date; applies continuously to all employment-related activities. · Max penalty: Employment Tribunal claims are uncapped for discrimination. There is no statutory maximum for Equality Act liability in employment cases; awards are determined by loss suffered and, in some cases, injury to feelings.

AI Acceptable Use Policy

Your formal AI policy. Sign it, circulate it to every member of staff, and review it annually or whenever your AI tool usage changes.

Brightwell Accountancy — AI Acceptable Use Policy

Document reference: BAC-AI-AUP-001
Version: 1.0
Issue date: 01 August 2026
Next mandatory review: 01 August 2027
Policy owner: [Name and job title of designated Policy Owner to be inserted on adoption]
Approved by: [Name of Director / Senior Partner to be inserted on adoption]

1. Purpose

Brightwell Accountancy handles payroll data, self-assessment records, and related personal and financial information for approximately 400 clients. Members of staff have begun using artificial intelligence (AI) tools on their own initiative. This policy exists to bring that activity under formal governance before it gives rise to a data breach, a regulatory enforcement action, or a loss of client trust.

The policy sets out which AI tools are approved, what may and may not be entered into those tools, how outputs must be verified, when human judgement is mandatory, what records must be kept, and what happens when the policy is broken. It is a live compliance document. Every member of staff must read it, sign the acknowledgement at the end, and act in accordance with it from the date of signature.

2. Scope

This policy applies to every individual who works for or on behalf of Brightwell Accountancy, including permanent employees, part-time staff, fixed-term contractors, agency workers, and any placement or work-experience student. It applies regardless of seniority, job function, or whether the person works on site, from home, or at a client's premises.

The policy governs the use of AI tools in all activities carried out for or in connection with Brightwell Accountancy's business, including drafting client correspondence, summarising meetings, conducting research and analysis, preparing payroll documentation, generating self-assessment workings, and any other task where a member of staff uses an AI tool to assist with work output.

The policy covers all devices used to access AI tools: firm-owned laptops, mobile devices, desktop workstations, and any personal device used for work purposes. If a member of staff accesses an AI tool on their personal phone or laptop for a work task, this policy applies in full.

3. Approved Tools and the Process for Requesting a New Tool

As at the issue date of this policy, the following AI tools are approved for use at Brightwell Accountancy, subject to the conditions set out in this document:

ChatGPT (OpenAI) — approved for drafting client correspondence templates, general research and analysis tasks, and producing first-draft written content where no personal data or confidential client material is involved.

Microsoft Copilot — approved for use within the Microsoft 365 environment for drafting correspondence, summarising documents, and research support, subject to the data-handling controls set out in Section 5 below.

Otter.ai — approved for transcribing and summarising internal team meetings only. Otter.ai must not be used to record or transcribe any call, meeting, or conversation in which a client is present, unless the client has given explicit, documented, prior consent and been informed of the tool's identity and the purpose of the recording.

No other AI tool, application, browser extension, plugin, or AI-assisted feature within existing software is approved for work use. This includes, but is not limited to, Grammarly's AI features, Google Gemini, Anthropic Claude, Meta AI, and any AI assistant embedded in a tool not listed above. The discovery that a member of staff has used an unapproved tool will be treated as a breach of this policy.

If a member of staff believes there is a legitimate business case for adding a new AI tool, they must submit a written request to the Policy Owner before using the tool. The request must identify the tool, its provider, the data it would process, the purpose of use, and the provider's data processing terms. The Policy Owner will carry out or commission a data protection impact assessment (DPIA) under UK GDPR Article 35 before approving any tool that would process personal data. No tool will be approved until the DPIA is complete, the data processing agreement is in place, and the Policy Owner has given written approval. Approval of a new tool will be followed by an updated version of this policy being issued to all staff.

4. What Must Never Be Entered into an AI Tool

The following categories of information must never be entered into any AI tool under any circumstances, regardless of whether the tool is on the approved list. This prohibition is absolute and has no exceptions unless an exception is expressly stated in writing by the Policy Owner.

Personal data of clients: This includes names, addresses, National Insurance numbers, Unique Taxpayer References, earnings figures, tax computations, bank account details, payroll records, dates of birth, and any other information that identifies or could identify a natural person. Brightwell Accountancy processes this data as a data controller under UK GDPR, and transmitting it into a third-party AI tool constitutes a disclosure to a third-party processor. No such disclosure may occur without a compliant data processing agreement and, where relevant, a DPIA. Until those safeguards are formally confirmed in writing by the Policy Owner for a specific tool, no client personal data may be entered.

Personal data of staff: Salary information, disciplinary records, performance assessments, health information, and any other information about employees or workers must not be entered into any AI tool.

Confidential client material: All information provided by clients in the course of Brightwell Accountancy's engagement is confidential. This includes business financial records, management accounts, correspondence, commercial arrangements, and any information shared under an expectation of confidentiality, whether or not there is a formal confidentiality clause in place. The professional duty of confidentiality owed by accountants applies independently of data protection law.

Credentials and access data: Passwords, API keys, authentication tokens, account numbers, and any other credentials must never be entered into an AI tool.

Information subject to a non-disclosure agreement or legal privilege: Any material covered by a non-disclosure agreement to which Brightwell Accountancy is party, or any information that is or may be subject to legal professional privilege, must not be entered into an AI tool.

The most common error staff are likely to make is copying a real client's details into a prompt to make the AI's output more specific. This is not permitted. Where context is genuinely needed, staff must either use anonymised or entirely fictitious placeholder data, or request a tool configuration that prevents data leaving the firm's controlled environment — and that configuration must be verified and approved by the Policy Owner before use.

5. Verification of AI Output and Responsibility for Errors

No AI-generated output may be used externally, sent to a client, submitted to HMRC or any other body, or incorporated into a file without first being reviewed and verified by a qualified member of staff. This requirement is not a formality. The courts of England and Wales have made clear, in cases including R (Ayinde) v Haringey LBC and Choksi v IPS Law LLP, that professionals who rely on AI-generated content without checking it face serious professional and legal consequences, including wasted costs orders and regulatory referrals. More recently, in July 2026, a criminal trial was halted after AI-generated documents were submitted without proper scrutiny. These cases all concern different professional contexts, but the lesson applies with equal force to accountancy.

Generative AI tools can produce figures, calculations, citations, legislative references, and factual statements that appear authoritative but are wholly or partially incorrect. This is commonly described as 'hallucination.' At Brightwell Accountancy, every piece of AI-assisted work must be checked against primary sources before it leaves the firm. For tax and payroll work, 'primary source' means the relevant HMRC guidance, legislation, or client-supplied data. For correspondence, it means checking that every factual assertion is accurate and that no client information has been misrepresented or fabricated.

Responsibility for errors in AI-assisted output rests with the member of staff who approved and sent the output. The fact that an AI tool produced an error is not a mitigating factor in client complaints, professional indemnity claims, or regulatory proceedings. The Policy Owner is responsible for ensuring that verification standards are communicated and maintained, but the individual reviewer bears personal responsibility for the work they sign off.

6. Where Human Decision-Making Is Mandatory

Brightwell Accountancy has confirmed that it does not use AI to make automated decisions about people within the meaning of UK GDPR Article 22, and the Data (Use and Access) Act (which took full effect on 5 February 2026 and which modified the UK's automated decision-making framework) does not change the firm's obligations in this respect because no such automated decisions are currently made.

Nevertheless, this policy identifies the following categories of decision as ones where human judgement is mandatory and AI output may only be used as background research or a drafting aid. A human decision-maker must form and record their own independent view before the decision is implemented:

Any advice given to a client on their tax liability, tax planning, or payroll position. AI may be used to draft the correspondence or summarise background, but the substantive advice must be reviewed and approved by a qualified accountant.

Any decision to report a suspicion of fraud or money laundering. These decisions engage professional obligations and potentially criminal liability. They must be made by a designated senior individual, not assisted or delegated to an AI output.

Any decision affecting a member of staff's employment, remuneration, performance assessment, or disciplinary position. These decisions engage the UK Equality Act 2010 and must be made by a human being following a fair process.

Any communication to HMRC, Companies House, or another regulatory body. AI may assist in drafting, but the submission must be reviewed, verified, and approved by a qualified member of staff before despatch.

If the firm's use of AI changes in future such that it involves automated decisions that produce legal or similarly significant effects on clients or staff, a fresh DPIA will be required under UK GDPR Article 35 before that use commences, and this policy will be updated accordingly.

7. Disclosure to Clients

Brightwell Accountancy's clients are entitled to know when AI tools have played a material role in the work done on their behalf. This matters not only as a matter of professional integrity but because clients have a reasonable expectation, grounded in their contract and in data protection law, that they understand how their information is being handled.

The following disclosure standard applies. Where an AI tool has been used to draft correspondence that is sent to a client in substantively the form the tool produced, the covering or footer note on that correspondence must state: 'Elements of this communication were drafted with AI assistance and have been reviewed and approved by [name or role of approving accountant].' Where AI has been used only for background research or internal summarisation that informed the work but is not directly reflected in client-facing output, no disclosure is required on each individual document, but the firm's engagement terms and privacy notice must inform clients that AI tools are used in this way.

Brightwell Accountancy will update its client engagement letter and privacy notice to reflect the use of approved AI tools within 30 days of this policy being adopted. The privacy notice must comply with UK GDPR Articles 13 and 14 and must describe the categories of processing that involve AI tools, the identity of the tool providers, and the safeguards in place.

Otter.ai may not be used to transcribe or summarise any meeting in which a client participates unless the client has been told, before the meeting begins, that the tool will be used, the identity of the tool provider, and the purpose of the recording. The client's agreement must be recorded in writing, either by email confirmation before the meeting or by a note signed at the start of the meeting.

8. Record Keeping and Retention

Brightwell Accountancy must be able to demonstrate to the ICO, to a professional indemnity insurer, or to a client that AI tools were used appropriately. This requires records.

For each piece of client-facing work where an AI tool was used, the member of staff must retain: the substantive prompt or instruction given to the AI (excluding any personal data that should not have been entered in the first place); a copy of the AI-generated output before human editing; a record of the substantive changes made during review; and the name of the member of staff who reviewed and approved the final output. These records must be saved to the client's file in the firm's practice management system, not stored only on a personal device or in the AI tool's own history.

For Otter.ai transcripts and summaries of internal meetings, records must be stored in a designated folder in the firm's document management system. Any transcript that contains a discussion of client matters must be treated as a confidential document and handled accordingly.

Records relating to AI-assisted work must be retained for the same period as the underlying client file to which they relate. For payroll and self-assessment work, HMRC recommends retaining records for at least six years from the end of the relevant tax year. Brightwell Accountancy's existing file retention schedule applies, and AI-related records form part of the file.

Staff must not rely on AI tools' own history functions as a substitute for saving records to the firm's systems. Chat histories in ChatGPT or Microsoft Copilot do not constitute a firm record. The firm's record is what is saved to its own systems.

9. Breach of Policy: Reporting and Disciplinary Position

A breach of this policy includes, but is not limited to: using an unapproved AI tool for work purposes; entering personal data, confidential client material, or credentials into an AI tool; failing to verify AI output before sending it externally; failing to disclose AI use to a client where this policy requires disclosure; and failing to keep the records required by Section 8.

Any member of staff who discovers or suspects a breach must report it to the Policy Owner immediately, and in any event within 24 hours of becoming aware of it. If the breach involves personal data being entered into an AI tool or otherwise disclosed without authorisation, the Policy Owner must assess whether it constitutes a personal data breach within the meaning of UK GDPR Article 4(12). Where there is a reasonable likelihood that the breach poses a risk to the rights and freedoms of the individuals whose data was involved, the ICO must be notified within 72 hours of the firm becoming aware of it, in accordance with UK GDPR Article 33. Where the risk to individuals is high, those individuals must also be notified directly under UK GDPR Article 34. The maximum fine for a serious breach of UK GDPR is £17.5 million or 4% of global annual turnover, whichever is higher. The turnover-based limb cannot be calculated here without Brightwell Accountancy's financial figures, but the statutory cap of £17.5 million is the minimum upper exposure for a serious breach.

Breaches of this policy will be investigated under Brightwell Accountancy's disciplinary procedure. Depending on the nature and severity of the breach, disciplinary action may range from a formal written warning to summary dismissal for gross misconduct. Entering confidential client data or personal data into an unapproved AI tool, or doing so knowingly in breach of this policy, will ordinarily be treated as gross misconduct. The firm reserves the right to report breaches to relevant professional bodies and, where legally required, to affected clients and the ICO.

10. Review Cycle and Policy Ownership

This policy must be reviewed in full no later than 01 August 2027. It must also be reviewed immediately, and updated before continued use of AI tools, in any of the following circumstances: a new AI tool is adopted or an approved tool is discontinued; a significant change is made to the way an approved tool processes data; new ICO guidance materially affecting the firm's obligations is published; there is a legislative or regulatory change that affects this policy's compliance status; or a data breach or near-miss connected to AI use occurs.

The Policy Owner holds day-to-day responsibility for maintaining this policy, monitoring compliance, approving new tools, managing breach reports, and ensuring that new staff receive a copy of this policy as part of their induction. The Policy Owner must also monitor ICO publications, including the ICO's AI Auditing Framework and any updates to it, and bring material developments to the attention of the firm's senior leadership.

This policy was approved by the firm's senior leadership on the date shown in the header. Any amendment to this policy requires the approval of a director or senior partner and must be reissued to all staff with a covering note explaining the nature of the change.

11. Staff Acknowledgement

Every member of staff must sign and return this section to the Policy Owner. Signed copies must be retained on the individual's personnel file for the duration of their employment and for six years thereafter.

I confirm that I have read, understood, and will comply with Brightwell Accountancy's AI Acceptable Use Policy (version 1.0, dated 01 August 2026). I understand that breach of this policy may result in disciplinary action up to and including dismissal.

Full name: .....

Job title: .....

Signature: .....

Date: .....

Policy Owner signature: .....

Policy Owner name and title: .....

Date of adoption: .....

Plain-English Employee Summary

A one-page summary of the AI policy that staff will actually read. Print it, pin it up, hand it out at induction.

AI at Brightwell Accountancy — What Every Member of Staff Needs to Know

This summary tells you what you can and cannot do when using AI tools at work. It applies to everyone at Brightwell Accountancy, whatever your role. Our approved AI tools are ChatGPT, Microsoft Copilot, and Otter.ai. If you want to use a different tool, you must get approval from your manager before you start — do not simply try it out on your own initiative.

What You May Do With AI at Work

You may use ChatGPT and Microsoft Copilot to help you draft letters, emails, and other written communications to clients, provided you review and edit the output before it goes anywhere. You may use these tools to support research and analysis tasks, such as summarising guidance notes or helping you structure a report. You may use Otter.ai to transcribe or summarise internal meetings where all participants have been told in advance that AI transcription is being used and have agreed to it.

All AI use must have a legitimate business purpose, and you must always apply your own professional judgement to whatever the tool produces. AI tools can and do make mistakes, including stating things that are simply untrue.

What You Must Never Put Into an AI Tool

Never enter any information that could identify a client or relate to a specific client's affairs. This means no names, addresses, National Insurance numbers, tax reference numbers, payroll figures, salary details, bank details, self-assessment data, or any other personal or financial information belonging to our clients or their employees. We hold this kind of data for around 400 clients, and entering it into an external AI tool would almost certainly breach our duties under data protection law and our professional obligations to those clients.

Never enter sensitive information about Brightwell's own staff, such as pay, performance, or personal circumstances. Never enter commercially sensitive business information such as client lists, fee arrangements, or internal financial data. The rule is simple: if it belongs to a client or could identify anyone, it stays out of the AI tool.

Check Before You Send — The One Rule That Prevents Most Incidents

Before any AI-assisted work leaves your hands — whether that is an email, a letter, a summary, or a report — read it in full yourself. Check that every fact, figure, and piece of advice is accurate. AI tools sometimes produce confident-sounding content that is wrong, outdated, or entirely made up. You are responsible for everything you send out under your name or on behalf of Brightwell. A quick check before you send is the single most effective way to avoid a serious mistake.

Who to Ask, and How to Report a Mistake Without Being Blamed

If you are unsure whether a particular task or AI tool is permitted, speak to your line manager or the firm's data protection lead before you proceed. It is always better to ask first.

If you think you have already made a mistake — for example, you have accidentally entered client data into an AI tool, or sent out something that contained an AI error — tell your manager straight away. Brightwell operates a no-blame reporting culture for honest mistakes that are reported promptly. Early reporting allows us to contain any problem before it becomes serious. Sitting on a mistake, on the other hand, makes things significantly worse for

everyone, including you. We need to know quickly so that we can protect our clients and, where required, notify the relevant authorities within the legal timeframes that apply to us.