

Data Concerns Undermine AI Applications for Climate

Emilie Vrain^{a,*} and Charlie Wilson^{a,b}

^a Environmental Change Institute, University of Oxford, Oxford, UK.

^b International Institute for Applied Systems Analysis (IIASA), Laxenburg, Vienna, Austria.

*corresponding author: emilie.vrain@eci.ox.ac.uk

Abstract

The global climate crisis demands innovative solutions, and artificial intelligence (AI), including large language models (LLMs), offers significant potential to improve efficiency and enable lower-carbon lifestyles. However, public concerns around data privacy may limit the use of AI for climate-beneficial applications. Drawing on three empirical studies with a nationally representative UK sample (N=2078), we show that heightened awareness of AI's data collection practices reduces willingness to use such technologies - especially when perceived user benefits are low. We also find that data protection behaviours and familiarity with AI can mitigate data privacy concerns. These behavioural dynamics call for interdisciplinary governance frameworks that integrate human factors alongside technical and environmental dimensions of AI deployment. Our results illustrate how everyday perceptions and practices can support or constrain digital transitions for sustainability.

Introduction

Recent discussions on artificial intelligence (AI) and climate change emphasise AI's dual potential for substantial positive and negative impacts. On one hand, AI has been identified as a powerful tool in reducing greenhouse gas emissions¹⁻³. For example, recent reviews find that AI-enabled smart manufacturing can reduce energy consumption, waste and carbon emissions by 30-50% and intelligent transportation systems can reduce CO₂ emissions by approximately 60%⁴. On the other hand, the surge in AI capabilities and applications, particularly through large language models (LLMs) and other high-computation systems, has led to soaring energy demands in data centres worldwide⁵⁻⁷. This rise in computational infrastructure poses a growing challenge as tech companies become significant energy consumers, sparking concerns over the sector's expanding carbon footprint⁸.

Beyond these direct impacts lie less obvious but increasingly influential dimensions of AI's climate impacts: the ways in which AI reshapes consumer engagement, and aggregate consumption patterns. AI applications in areas like micro-target marketing and agentic systems could indirectly influence societal energy demands, potentially leading to increased consumption^{9,10}. At the same time, AI enables numerous consumer-oriented digital applications that have the potential to reduce carbon emissions by helping individuals change daily behaviours. We refer to these as 'climate-beneficial AI applications'. Examples include: smart energy management tools providing real-time feedback, helping users see how and when they're consuming the most energy; peer-to-peer (P2P) platforms enabling collaborative consumption through usership business models or more efficient resource use, extending product lifespan; and food waste reduction apps that redistribute surplus food¹¹⁻¹³. The extent to which these applications contribute to meaningful carbon reductions depends on how they are adopted and used⁹. Their impact may be positive by replacing high-carbon behaviours, or negative if they lead to increased usage and energy consumption through rebound effects^{14,15}.

Despite the growing landscape of climate-beneficial AI applications, a critical knowledge gap remains: how public concerns about AI, particularly around personal data use, affect the use and impact of these technologies. Prior research has examined drivers of sustainable technology use (e.g.¹⁶) and data privacy concerns in digital tools¹⁷⁻¹⁹ but not their intersection in the context of climate-beneficial AI applications. This oversight is increasingly consequential as AI becomes more salient in public discourse and the demand for climate action accelerates. This study addresses the gap by investigating how data privacy concerns, amplified by perceptions of AI risk, influence use of climate-beneficial AI applications, and what factors mitigate or exacerbate these effects.

The urgency of this investigation is further highlighted by the peak in public attention to AI during 2023-24, following the launch of ChatGPT, which our study directly exploits, reflecting on changes in perceptions over the year. With generative AI (GenAI) tools like ChatGPT becoming "the fastest application to reach 100 million users"²⁰, AI has surged in prominence with considerable media attention, raising the salience of its usage across sectors²¹⁻²³. Public surveys report heightened

perceptions of AI risk, distrust, and concerns over opaque data use^{24–28}, amplified by high-profile data misuse scandals (e.g., Cambridge Analytica’s misuse of Facebook data²⁹ and Clearview AI’s image-scraping³⁰). AI governance has progressed including through the EU’s AI Act (Regulation (EU) 2024/1689), however, risk mitigation focuses on system safety, fundamental rights and trustworthiness. There remain notable gaps not only in the need to align AI systems with sustainability^{31–33} but also to account for how public trust, privacy concerns and perceptions of AI risk affect societal uptake of tools that could support climate goals.

As AI technologies become more prominent in society, we argue that public awareness of data practices increases in tandem—potentially creating new barriers to the usage of climate-beneficial AI applications. Are data concerns undermining AI’s contribution to tackling climate change? Although we reason that data privacy concerns play a key role, it is overly simplistic to view them as the sole determinant of technology use. A range of other psychological factors – including perceived ease of use, perceived functionality, and environmental motivation – are known to influence both usage and usage propensity^{34,35}. We explore how salience of AI’s data collection practices interacts with these other drivers of technology acceptance.

While AI applications are more visible in domains like targeted advertising – where personalisation is overt and often discussed³⁶ – their use in shared mobility or food waste reduction platforms tends to be less explicitly communicated such that users may be unaware of their data-driven functionalities. According to the privacy paradox theory, individuals express concerns about data privacy yet continue to engage in behaviours that expose their personal data^{37,38}. Two prominent explanations for this paradox are: 1) individuals weigh the perceived benefits of using a technology against its abstract privacy risks (so-called privacy calculus) and 2) decisions are made in the absence of clear awareness about how data are collected, used, or shared³⁹. We argue that concerns about data privacy may deter some users, privacy calculus may lead others to use protective behaviours that mitigate perceived risks, potentially increasing trust and permitting use of AI applications.

In this paper, we offer a novel, empirically grounded investigation into how the use of climate-beneficial AI applications is affected by data privacy concerns, and, what factors mitigate or exacerbate these effects. We do so through three studies using a nationally representative UK sample (N=2,078). Fig. 1 provides an overview of the research design. Study 1 and 2 examine four case study climate-beneficial AI applications across domains of daily life: 1) retail – P2P retail platform; 2) mobility – bike share scheme platform; 3) food – food waste reduction app; 4) home - smart thermostat. Study 1 uses a vignette experiment to test how varying levels of AI salience and perceived data collection affect usage propensity. The impact of perceived relative advantage is also explored. Study 2 investigates how perceived AI risks influence data privacy concerns, and in turn how concerns impact the usage of the four applications. It also examines how changes in privacy concerns over the past year relate to application usage. Additionally, we study whether the recent surge of information and awareness of

GenAI – often termed the "GenAI hype"^{40–42} – influences perceived AI risks, as well as whether behaviours around data protection positively mediate privacy concern's impact on application use. Finally, Study 3 tests the generalisability of our findings across a diverse set of 15 climate-beneficial AI applications.

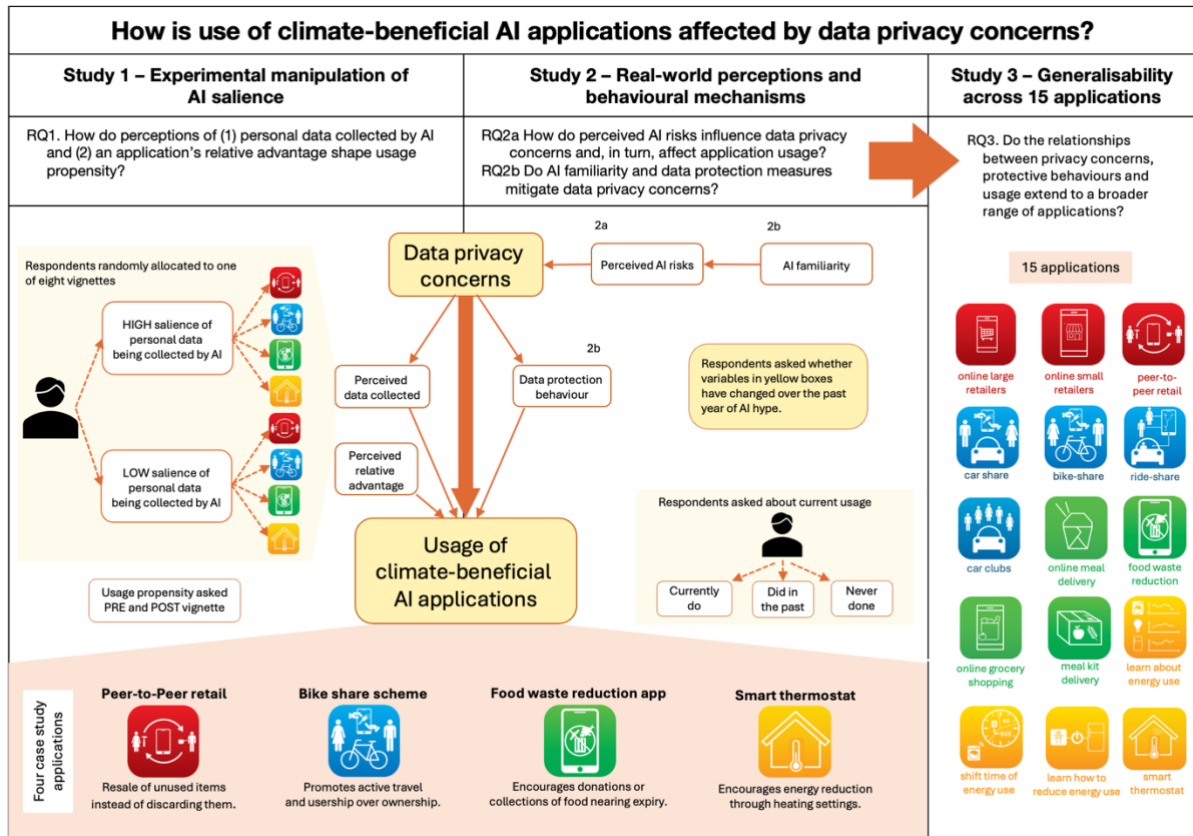


Fig. 1. Summary of the research design framework. Study 1: Respondents were randomly assigned to one of eight vignettes describing a scenario of using a climate-beneficial AI application, emphasising either 1) high salience or 2) low salience of AI collecting personal data for functionality. Respondents were asked about relative advantage, perceived data collection and usage propensity. Study 2: Respondents were asked about their familiarity of AI (including exposure to information about AI and its importance, ability to recognise AI, seeking information about AI, and use of GenAI), perceived AI risks, current data privacy concerns, data protection behaviours, application usage frequency, usage propensity and whether some of these have changed over the previous year of AI hype. Study 3: Respondents were asked about their application usage frequency for an additional 11 climate-beneficial AI applications across daily life domains.

Results

Study 1 Impacts of AI data collection practices and relative advantage

The salience of AI collecting personal data significantly reduced the *usage propensity* for climate-beneficial AI applications (mean reduction -0.260, 95% CI [-0.350 to -0.170], $p < 0.001$, $r = -0.161$ small effect). This effect was particularly pronounced for P2P retail platforms and smart thermostats, shifting respondents' usage propensity from 'likely' to 'unlikely'. Even low salience of AI collecting data reduced propensity for P2P retail platform usage, suggesting inherent concerns regarding data practices in such applications. No significant effects were observed for food waste reduction apps,

whereas salience increased *usage propensity* for bike share schemes; however, respondents remained ‘unlikely’ to use them (Supplementary Table 1).

To assess whether our experimental design effectively manipulated perceptions of AI collecting personal data [*perceived data collected*], we examined differences in responses to our vignettes. Participants exposed to a high AI salience vignette were more likely to perceive that personal data was being collected compared to those exposed to a low AI salience vignette ($U = 606167.50$, $z = 5.169$, $p < 0.001$, $r = 0.113$ small effect). Differences in perceptions were significant for P2P retail platforms ($U = 35082.50$, $z = 2.946$, $p = 0.003$, $r = 0.132$ small effect) and bike share schemes ($U = 40931.00$, $z = 4.475$, $p < 0.001$, $r = 0.196$ small effect) but not for other applications (Supplementary Table 2). While our vignette-based experimental design effectively introduced variation in perceptions of AI data collection practices, the resulting effects differed across applications. Because our main interest lies in the role of perceived data collection, regardless of vignette assignment, we shift our analysis from between-vignette comparisons to participants’ responses on *perceived data collection*. This enables us to better capture the underlying causal and mediating mechanisms in our ex-post framework.

Structural equation modelling (SEM) results indicate that the association between *data privacy concerns* and *usage propensity* is fully mediated by *perceived data collected*, ultimately reducing *usage propensity* across all four climate-beneficial AI applications (Table 1, with model fit indices provided in Supplementary Table 3). *Data privacy concerns* have a positive effect on *perceived data collected*, and *perceived data collected* has a negative effect on *usage propensity*. These findings suggest that concerns about data privacy shape individuals’ willingness to use AI applications by influencing their perceptions of AI data collection practices. Thus, perceptions of such practices potentially act as a psychological barrier, amplifying the impact of privacy concerns on usage decisions.

We next demonstrate that *perceived relative advantage* acts as a counterbalancing influence on users’ concerns over data privacy (e.g., for P2P retail $\lambda = 0.63$, $p < 0.001$; Fig.2). This is consistent with privacy calculus theory which characterises how digital users trade off benefits with privacy risks. We find *perceived relative advantage* has the strongest effect on *usage propensity* for P2P retail platforms and smart thermostats (Supplementary Fig.1).

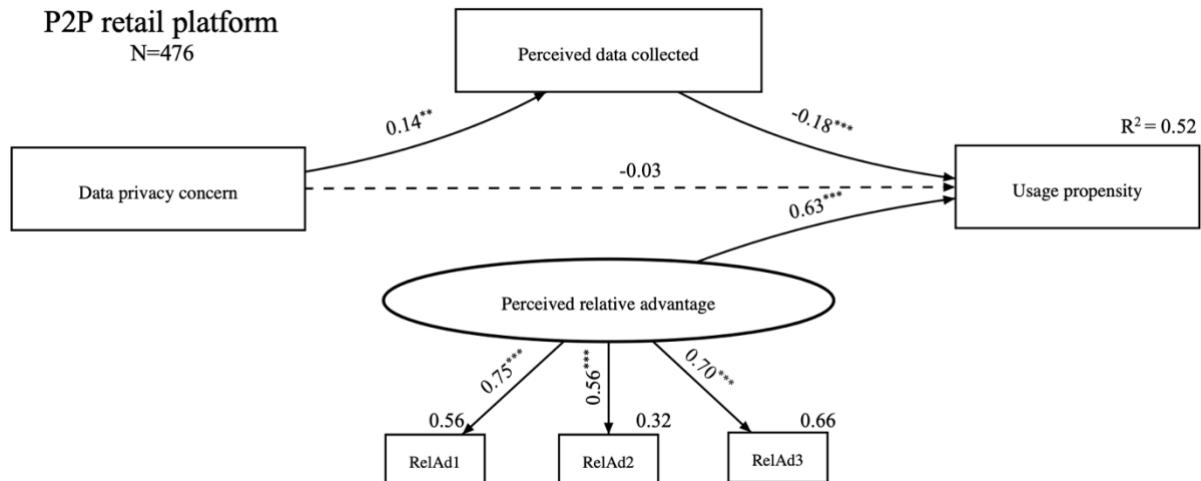


Fig. 2. Structural equation model showing the direct and indirect effects of data privacy concern, perceived data collected and relative advantage on usage propensity of a P2P retail platform (measured post-vignette). The latent variable ‘perceived relative advantage’ consists of: RelAd1 ‘better than alternatives’, RelAd2 ‘easy to use’ and RelAd3 ‘useful’. Numbers adjacent to the arrows are standardised path coefficients and indicative of the effect size of the relationship. Dashed lines indicate nonsignificant relationships. *P<0.05, **P< 0.01 and ***P< 0.001. R² represents the proportion of variance explained by the relations in the path model.

Table 1. Study 1’s test for mediation using a Bootstrap analysis with a 95% Confidence Interval. Unstandardised coefficients reported. Bootstrap Sample = 5000 with replacement.

Study 1 – Path analysis	Data privacy concerns → <i>Perceived data collected</i> → Usage propensity								Perceived relative advantage → Usage propensity	
	n	Direct effect	P	Indirect effect	95% CI		P	Mediation conclusion	Direct effect	P
		β	(2-tailed)	a x b = β	Low	High	(2-tailed)		β	(2-tailed)
P2P retail platform	476	-0.005	n.s.	-0.005	-0.010	-0.001	0.001	full	1.255	<0.001
Bike share scheme	513	-0.010	n.s.	-0.004	-0.010	-0.001	0.002	full	0.748	<0.001
Food waste reduction	512	-0.003	n.s.	-0.009	-0.016	-0.004	<0.001	full	1.086	<0.001
Smart thermostat	519	0.010	n.s.	-0.003	-0.007	0.000	0.021	full	1.014	<0.001

a: Effect of the independent variable on the mediator.

b: Effect of the mediator on the dependent variable.

We have shown that AI salience amplifies data privacy concerns that undermine usage propensity of climate-beneficial AI applications, particularly if their relative advantage is weak. However, this was through an experimental manipulation of AI data collection salience. In our second study, we test the validity of our findings given current levels of AI awareness particularly over the past year during the GenAI hype.

Study 2 Impacts of perceived AI risks, familiarity and data protection behaviours

Existing perceptions of AI are strongly biased towards AI-related risks: 36% of Study 2 participants think AI poses greater risks than benefits – double the proportion who view its benefits as outweighing its risks (18%). (Personal characteristics of respondents across three risk-benefit subgroups are summarised in Supplementary Table 5). The most frequently cited perceived negative impact of AI was

'*use of personal data without consent*,' reported by 72% of participants. These concerns are consistent with population-level trends²⁴. *Data privacy concerns* increased over the past year, while *usage* of the four climate-beneficial AI applications decreased (Supplementary Table 4). Together these findings highlight potential widespread issues about AI's implications for data privacy and trust.

Perceived AI risks were positively associated with self-reported *data privacy concerns* ($r_s = 0.163$, $p < 0.001$) and *changes in concerns* over the past year ($r_s = 0.201$, $p < 0.001$) (Supplementary Table 6a). However, although high privacy concerns are expected to deter application usage, we find no significant difference in levels of concern between users and non-users of climate-beneficial AI applications (Supplementary Table 7). We also find significant positive correlation between *changes in data privacy concerns* and *changes in application usage*. These findings on real-world behaviour are consistent with the privacy paradox that explains why users report concerns yet continue using applications. Can risk mitigation measures help further counter the undermining effect of AI risks and data privacy concerns on use?

First, AI familiarity through *exposure to AI-related information* and the *perceived importance of information source* (e.g., media, social interactions) were associated with lower *perceived AI risks*, though effect sizes were weak (Cramer's $V < 0.2$ for exposure and Spearman's rho values ≈ 0 for importance - Supplementary Table 8 and Table 6c). Stronger associations emerged for other aspects of AI familiarity, namely proactive interactions, where *actively seeking AI-related information*, *recognising AI use*, or *using ChatGPT* correlated with lower *perceived AI risks* ($r_s = -0.211, 0.219, 0.317$, $p < 0.001$, Supplementary Table 6d). Results suggest that such familiarity may mitigate uncertainty around AI risks (or conversely, lack of familiarity breeds mistrust).

Second, *data protection* actions to protect personal data online (e.g., removing location tracking, creating a difficult-to-guess/strong password) fully mediated the relationship between *data privacy concerns* and *usage frequency* for P2P retail platforms, food waste reduction apps and smart thermostats. For bike share schemes, partial mediation was found as *data privacy concerns* were significantly negatively associated with *usage frequency* ($\lambda = -0.010$, $p = 0.007$) (Supplementary Table 9 and 10).

We show that proactive engagement with AI, especially usage of GenAI, may help reduce perceived AI risks, and that data protection behaviours help mitigate the potential negative impact that data privacy concerns have on the usage of four climate-beneficial AI applications. Do these findings hold for a more diverse set of AI applications? In Study 3 we expand our investigation to an additional 11 climate-beneficial AI applications and explore data privacy concerns and the mediating role of data protection behaviours on usage.

Study 3 Impacts on a broad range of climate-beneficial AI applications

Broadening the investigation to a wider range of climate-beneficial AI applications, trends show a decline in usage over the past year across home, food, mobility and retail domains (Fig. 3). SEMs revealed that *data privacy concerns* do not uniformly deter usage; instead, their effect is largely mediated through *data protection* behaviours, which in many cases are associated with higher *usage frequency*. For most applications, *data privacy concerns* were fully mediated, suggesting that rather than disengaging, concerned individuals adopt mitigating strategies of data protection that enable continued use of AI-enabled applications. Privacy concerns can motivate individuals to implement protective strategies, allowing them to navigate AI-related risks without forgoing its benefits—a novel finding to inform the privacy paradox.

Notably, for bike share schemes (reported in Study 2) and meal kits, *data privacy concerns* retained a significant negative direct association with *usage frequency* ($\lambda = -0.009$, $p = 0.048$), even after accounting for *data protection* behaviours.

In contrast, informational energy applications – particularly those aimed at helping users reduce energy consumption i.e., ‘learning how to reduce energy use’, exhibited a small but significant positive direct association between *data privacy concerns* on *usage frequency* ($\lambda = 0.025$, $p < 0.001$). This finding diverges from the broader trend, suggesting that in contexts where AI is perceived as non-invasive and empowering for users, the benefits (e.g., gaining energy-saving knowledge) outweigh the minimal privacy risks.

Finally, it is important to note that model fit was stronger for retail and home energy applications, suggesting greater clarity or salience of data-driven functionalities in these domains (Supplementary Table 11).

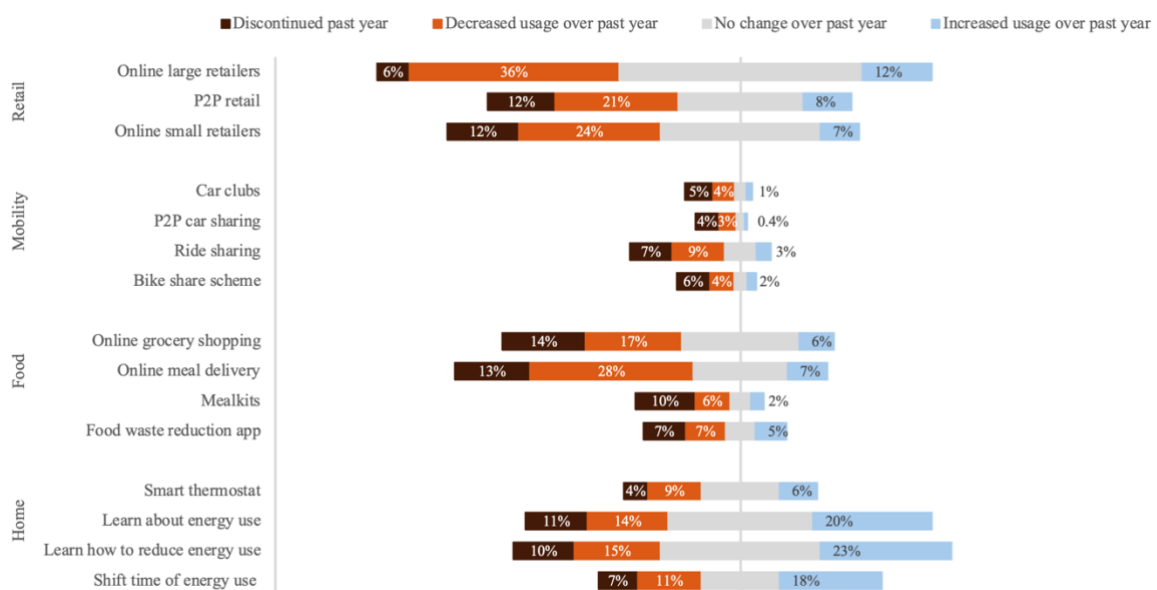


Fig. 3. Participants' change in usage of 15 climate beneficial AI applications compared to a year ago (n=2078).

Discussion

In this paper we present new empirical evidence on how data privacy concerns shape people's usage of climate-beneficial AI applications. Through three studies with a nationally representative UK sample, we substantiate the privacy paradox theory but significantly extend its explanatory value by uncovering more complex causal pathways in individuals' risk calculus. Our findings show that although data privacy concerns are widespread and increasing, they do not uniformly reduce the use of climate-beneficial AI. Crucially, privacy concerns do not act as direct deterrents, but rather through perceptions that AI applications extract and utilise personal data.

Our experimental results from Study 1 show that when AI involvement is made explicit, willingness to use climate-beneficial AI applications declines, particularly when their perceived relative advantage is low. This effect is most pronounced in areas such as online retail and smart energy systems. Two explanations for continued usage despite privacy concerns is that many AI-enabled applications are not yet perceived as posing significant privacy risks, or that the AI component of their functionality is not salient to users. Our results align with the privacy calculus model, according to which individuals weigh perceived benefits against perceived privacy risks^{39,43,44}.

Building on these findings, Studies 2 and 3 further show that individuals with strong privacy concerns are more engaged with climate-beneficial AI applications if they implement data protective behaviours. This mediation effect suggests that rather than avoiding AI applications, users actively manage their privacy concerns by mitigating risks. In addition to protective behaviours, Study 2 found that familiarity with AI through proactive engagement such as seeking information or using GenAI significantly reduces perceived AI risks. This highlights the importance of reducing 'fear of the unknown' through user engagement with AI, suggesting that increasing familiarity can act as a strategy for mitigating perceived risks and concerns. These findings offer a behaviourally grounded explanation of the privacy paradox to explain why concerns do not always translate into disengagement. Our findings add new empirical weight to this theory and provide evidence for its relevance in the climate-AI nexus.

Our findings also introduce a behavioural dimension that complements existing AI risk frameworks, such as the MIT AI Risk Repository⁴⁵, which primarily catalogue technical, operational and societal risks. Although such frameworks aim to identify and inform risk mitigation strategies, our results suggest that end users themselves also play a role in risk mitigation. Specifically, we find that data privacy concerns can shape usage patterns through behavioural responses like protective strategies and familiarity with AI, which can positively influence the uptake of climate-beneficial AI applications. As AI becomes more embedded in climate-related infrastructures and services, understanding these user responses will be essential for promoting widespread adoption.

However, an important insight is that data privacy concerns can still directly impact usage, even when protective behaviours are implemented. For example, residual negative effects of privacy concerns on the use of bike-share schemes and meal kits may reflect perceived limits to individual control over data, such as opaque data practices, mandatory data sharing or distrust of how services manage personal information. For example, Chinese bike-sharing services such as Ofo and Mobike have faced criticism for collecting extensive geolocation and behavioural data without clear user consent, prompting public scrutiny and regulatory responses^{46,47}. Similarly, the UK Data Regulation Authority fined HelloFresh for sending millions of unsolicited commercial emails, drawing attention to the misuse of personal data in the meal kit sector^{48,49}. These cases contribute to heightened concerns, reinforcing the need for clear communication and greater responsibility from application providers regarding their data practices.

These behavioural and perceptual dynamics occur within a broader system of technological governance. As global frameworks are increasingly advocating for AI to support sustainability goals^{31,32}, our findings suggest that without sustained attention to data privacy, progress towards these goals may be compromised. Responsible design and policy interventions are necessary to ensure that data privacy concerns do not become a structural barrier to the adoption of climate-beneficial AI applications.

Building on these insights, our findings highlight two key avenues for further inquiry. First, following Bossert and Loh⁵⁰, we argue for holistic sustainability assessments of AI, including its social acceptability, data practices and demands, and perceived intrusiveness. Addressing such dimensions would help identify trade-offs and conflicting interests. Second, future research should directly measure perceived data collection across applications, to better triangulate practices with perceptions. These two lines of enquiry would clarify not just whether AI should be used in each context but whether it can be used in ways that people trust and are willing to sustain.

To mitigate AI-related privacy concerns, we recommend increasing transparency and enhancing user control over data collection practices. The need for clear consent mechanisms and better transparency in data usage are critical for building trust, as emphasised in studies on AI^{33,44}. Based on our findings, we propose three practical strategies for AI providers and users to mitigate the impacts of data privacy concerns on climate-beneficial AI applications:

- 1) Utilise the privacy paradox** by emphasising the relative advantage of the applications, communicating clearly both climate (pro-social) and personal (functional) benefits, to outweigh privacy concerns.
- 2) Support digital privacy literacy** by equipping users with skills to protect themselves, empowering safe engagement without opting out entirely.
- 3) Strengthen familiarity with AI** by both increasing exposure to AI-related information and encouraging greater use of GenAI, which our data show is linked to reduced perceived risks.

These strategies are increasingly pertinent in light of recent evidence on the widespread undisclosed data tracking by 100 leading Android apps⁵¹. Without transparency of data practices, AI-enabled services risk eroding public trust⁵² – even when they serve climate goals. To scale climate-beneficial AI equitably and sustainably, our findings support calls for developers to incorporate "privacy by design" principles, ensuring responsible data collection and usage. As results from Study 1 demonstrated, perceived relative advantage of an application is important in user decision making. Effectively communicating tangible user benefits and offering strong privacy as a unique selling proposition (USP) could help unlock wider public uptake of climate-beneficial AI applications. Within a diverse ecosystem of AI application providers, those with a strong privacy (USP) can differentiate themselves and build user trust. As discussed in recent literature on LLMs³³, our results suggest the importance of: AI developers implementing rigorous frameworks to protect user privacy and ensure fairness; and policymakers mandating transparency and user control as a default, not optional features. Our results support calls by the Global Partnership on AI for stronger governance to ensure privacy, equity and environmental alignment⁵³.

Broader implications emerge from our research. While attention to existential or otherwise extreme risks of AI is growing^{54–57}, our findings advocate for incorporating a more scientifically grounded and broader range of social considerations into AI application development. Such considerations are critical for aligning AI systems with climate mitigation goals and promoting low-carbon lifestyles.

Conclusion

Our research demonstrates that AI data collection practices and individuals' data privacy concerns shape the usage of climate-beneficial AI applications, but privacy concerns do not uniformly deter usage. Instead, perceived data harvesting acts as a key psychological barrier to usage propensity, but data protection behaviours and perceived relative advantage mediate the negative effects of privacy concerns. As AI becomes more prominent in society and ingrained in daily life, so too will public awareness of the personal data AI applications collect and analyse — potentially creating new barriers to the usage of sustainable digital products and services. The explosive rise of GenAI appears to have increased perceived AI risks, but exposure to AI-related information and proactive engagement mitigates adverse consequences. Users need to navigate privacy risks through protective strategies rather than avoiding AI applications. This emphasises the importance of designing AI-driven solutions that balance transparency, security, and usability to maximize their beneficial climate impact.

Methods

National online survey

We conducted three studies to investigate how data privacy concerns influence the use of climate-beneficial AI applications, as well as the factors shaping this relationship. Data were collected through a nationally representative survey of UK adults ($n = 2,078$), for which participants received monetary compensation. The survey instrument was administered by a market research company (Qualtrics) in April 2024. The UK was chosen as the reference country due to the wide availability of climate-beneficial AI applications and the policy imperative to reduce per capita emissions in line with legally binding domestic climate targets⁵⁸. The final sample passed quality checks and screening. The median survey duration was 11.9 minutes. Sample demographics are available in Supplementary Table 13.

The survey instrument comprised six blocks of questions followed by a vignette-based experimental survey block (Table 2). Questions used either single or multi-item scales based on precedents in the literature (with slight modifications to fit the current research context) and newly developed items. Many questions consisted of statements with level of agreement captured by a 5-point Likert scale. The full survey instrument provided in Data Availability also outlines which questions came from precedent literature. Further details of the sampling method and data quality checks are provided in Supplementary Information – Methods.

Table 2. Survey blocks and example questions

Block	Topic	Example questions
1	Socio demographics and household characteristics	What is the highest level of education you have completed?
2	Application usage status and change in usage (x15 applications)	How often do you do the following [domain]-related activities? Compared to a year ago, how would you describe the frequency you do the following activities?
3	Application usage intention (1 application x 4 domains)	Within the next year, how likely are you to use [application]?
4	Technology attitudes	To what extent do you agree or disagree with the following statement...? Often it is easier to do things without using digital technologies.
5	Data privacy concerns and change in concern	When using the internet in daily life, how concerned are you that... Your internet usage information (including details of items you searched or purchased) is shared with websites or companies which you don't use? Has your level of concern changed compared to a year ago for the following..?
6	AI exposure and perceptions	How important have these sources of information been in shaping your opinion of AI? To what extent do you agree or disagree that AI will benefit you?
Vignette – participants randomly allocated to one block (application domain – salience of AI collecting personal data)		
7	Retail – High	Based on the information provided in the story, how likely are you to use a [application] like the one described?
8	Retail – Low	
9	Mobility – High	To what extent do you agree or disagree with the following statements about [application] like the one described?
10	Mobility – Low	
11	Food - High	Using them would be better than other available options
12	Food - Low	I think they would be easy to use
13	Home - High	I think they would be useful
14	Home - Low	When using them, they would collect personal data

Footnote: [domain] and [application] are specific to survey variants. Four domains are: retail, mobility, food, home. The 15 innovations include P2P retail platform, bike share scheme, anti-food waste app (for full list see Fig. 3).

Initial survey blocks captured demographic and contextual characteristics to inform inclusion criteria (above the age of 18) and quota fulfilment (nationally representative sample on age and gender), followed by self-reported use of a diverse set of 15 applications largely enabled by AI (e.g., learning algorithms, cloud-based services, or other forms of data-dependent service-provision) in four domains of daily life: retail, mobility, food and home energy (Fig. 3). These applications are referred to as ‘climate-beneficial’ throughout given available evidence showing the potential for carbon emission reduction^{11–13}. We use the term ‘climate beneficial’ but acknowledge that their net climate impact depends on user behaviour, system effects and potential rebound dynamics. This label reflects their functional potential to support low-carbon behaviours rather than their assumed effectiveness. Logic branching allowed follow-ups regarding usage frequency to be customised according to current, past or non-use. For example, a participant who indicated current use of an application was then asked, ‘how often do you do [activity]?’. Subsequent blocks assessed usage propensity for four case study applications (one per domain), general attitudes towards technology⁵⁹, data privacy concerns^{60,61}, data protection behaviours^{62,63} and perceptions of and attitudes to AI²⁴. Several items on AI information sources were adapted from³⁴, and additional questions included, for example to assess AI familiarity e.g., usage of GenAI.

To examine the potential effects of the recent surge in AI awareness following the release of ChatGPT 3.5, we also included retrospective items measuring changes in attitudes and behaviours over the previous year of 2023 (the first full year of the GenAI hype). To capture change in application usage, we asked ‘Compared to a year ago, how would you describe the frequency you do the following activities?’.

Following the main survey, we implemented a vignette experiment to test how awareness of AI data practices affects usage propensity. Each respondent was randomly assigned to one of eight vignettes describing the use of one of four climate-beneficial AI applications, with either an increased or reduced emphasis on personal data being collected. The vignettes respected the principles of realism, clarity, simplicity and internal consistency^{64,65}. A pilot survey (n = 56) was used with two additional exploratory questions to identify applications with varying perceptions of data collection and associated concerns, informing vignette application selection (Supplementary Table 14). Quota sampling was used to approximate 250 respondents per vignette to ensure sufficient statistical power. Final sub-sample sizes ranged from 245 to 275. To elicit accurate estimates of usage propensity, respondents were asked post-vignette whether they would use the application described, rather than assessing what others should do⁶⁴.

Additional post-vignette questions were informed by prior research to investigate mechanisms of the privacy paradox by assessing respondents’ perceptions of personal data collection during application use, alongside perceived relative advantage (operationalised as a latent construct using three items:

better than alternatives, easy to use, and useful; see Fig. 3, Supplementary Fig. 1). Vignettes and respondent sub-sample characteristics are available in Supplementary Tables 15 and 16.

Data preparation

We conducted data reduction and preparation steps for all three studies. Reliability testing demonstrated acceptable internal consistency for key scales: *perceived relative advantage* ($\alpha = 0.784$), *data privacy concerns* ($\alpha = 0.878$), and technology attitudes – referred to as *technophilia* ($\alpha = 0.705$). *Data protection behaviour* items showed marginal internal consistency ($\alpha = 0.699$). *Perceived relative advantage* items were combined into a latent variable for SEM, while *data privacy concerns* and *technophilia* scores were summed following established approaches^{60–62}. Notably, scale items for *technophilia* were recoded where necessary to ensure consistent directional meaning⁶². *Data protection behaviours* were summed to create a scale.

For each AI application, respondents who stated, ‘used in the past but not now’ or ‘never used’, were recategorised as ‘never’ for *usage frequency*. Ordinal variables representing change over the past year were constructed for *data privacy concerns* and *application usage*. For *data privacy concerns*, change scores were created by summing responses across six items (*privacy_compare 7.2_1* to *7.2_6*), yielding a scale ‘*change data privacy concerns*’ from – 6 (large decrease) to +6 (large increase). The variable ‘*change application usage*’ was derived by combining responses from current users’ reported changes in frequency over the past year (*compare_[application]*), and past users’ discontinuation timelines (*lastused_[application]*), i.e., capturing past users who discontinued during the past year.

To maintain analytical clarity, ‘don’t know’ responses were recategorised based on variable type:

- For current *application usage*, and *data protection behaviours*, ‘don’t know’ responses were recoded as ‘never done’, based on the assumption participants unable to recall engaging in a behaviour or with an application likely had not done so recently or meaningfully.
- For *technophilia*, ‘don’t know’ responses were recategorised as ‘neither agree nor disagree’ preserving the ordinal scale structure and avoid unnecessary missingness.
- For *perceived AI risks*, *usage propensity* and *usage frequency*, ‘don’t know’ responses were labelled as missing data and regression imputation was used to deal with missing data and to facilitate bootstrapping in the SEMs.

Data analysis

We ran statistical analyses separately for each study using SPSS (version 30.0.0.0) and Amos Graphics (version 29.0). For Study 1, to verify participants were randomly assigned across the eight vignette groups and assess baseline differences across socio-demographic variables and technology attitudes, Kruskal-Wallis tests were used for continuous variables e.g., *technophilia*, and chi square tests for nominal variables e.g., *education level*. No significant differences were detected (Supplementary Table 16).

Wilcoxon signed-rank tests assessed changes in application *usage propensity* (measured pre and post vignette exposure). Mann-Whitney U tests compared differences in *usage propensity* changes, as well as differences in *perceived data collection* between high and low AI salience conditions.

Causal pathways between *data privacy concerns* and application *usage propensity* were examined through SEMs which incorporated *perceived relative advantage* and *perceived data collection*. The overarching model was specified based on established theory (privacy paradox) and applied to the four case study climate-beneficial AI applications. Both measurement and structural components were included, with a latent construct derived from three observed indicators of *perceived relative advantage*. The scale was fixed by constraining the variance to 1. SEMs were estimated using the covariance matrix, applying the maximum likelihood (ML) procedure with bootstrapping to account for non-normality (Shapiro-Wilk $p > 0.05$).

Model fit indices show acceptable to excellent fit across application models (Supplementary Table 3). The specific application models for food waste reduction app (RMSEA = 0.059, CFI = 0.978) and smart thermostat (RMSEA = 0.069, CFI = 0.974) demonstrated excellent fit. Fit was weaker for the P2P retail platform and bike-share scheme (RMSEA > 0.09) though still within acceptable thresholds suggesting models fit the data very well⁶⁶.

For Study 2, descriptive statistics were produced for independent variables (Supplementary Table 4), and for participant subgroups based on their *perceived AI risk* (Supplementary Table 5). Spearman's rank correlations were used to explore associations between variables (Supplementary Table 6); monotonicity was confirmed through preliminary visual checks. Mann-Whitney U tests assessed group differences in *data privacy concerns* between current users and non-users of each of the four case study applications.

SEM was conducted per application to investigate the mediation of *data protection behaviours* between *data privacy concerns* and application *usage frequency*. ML with bootstrapping based solely on observed variables was used. However, SEM model fit indices indicated relatively poor fit: RMSEA values ranged from 0.144 (bike-share scheme), CFI values from 0.862 to 0.614, and TLI values from 0.724 to 0.228 (Supplementary Table 9).

The models' RMSEA values indicate a reasonable fit for two of the applications (P2P retail platform, 0.036 and food waste reduction app, 0.079), but the low CFI and TLI scores across the four models (Supplementary Table 9) suggest that the specifications did not fully capture the underlying data structures. The results should therefore be interpreted as indicative.

For Study 3, SEM statistical analyses used in Study 2 were applied to an additional 11 climate-beneficial AI applications investigated in the survey, providing insights on a total of 15 AI applications. In the main text we report all relevant results for answering our research questions listed in Fig. 1. In the

Supplementary Information, we report additional model results and assessments as part of robustness checks and sensitivity testing.

Data Availability

The dataset generated by the survey research during the current study will be deposited in a formal data repository but is temporarily available during the review process for anonymous download at: <https://www.dropbox.com/scl/fo/ffurwa7cs4f5onuke5ife/AGabYYE1-YAh1ofHT05T9Lo?rlkey=gfnaddoy5z8dmshef9vqoewo0&st=54olp68e&dl=0>

Ethics Declaration

This research was approved by the [Name of Ethics Committee] at [Institution] [Approval Ref] (details removed for review process). All participants gave informed consent before participating in the online survey used in the studies. No personally identifying data were collected or retained.

References

1. Stern, N. et al. Green and intelligent: the role of AI in the climate transition. *npj Clim. Action* **4**, 19 (2025).
2. Rolnick, D. et al. Tackling Climate Change with Machine Learning. *ACM Comput Surv* **55**, (2023).
3. Kaack, L. H. et al. Aligning artificial intelligence with climate change mitigation. *Nat Clim Chang* **12**, 518–527 (2022).
4. Chen, L. et al. Artificial intelligence-based solutions for climate change: a review. *Environmental Chemistry Letters* **21**, 2525–2557 (2023).
5. Sfondrini, N. GenAI's Carbon Footprint: A New Challenge for Corporations. *Forbes* <https://www.forbes.com/councils/forbestechcouncil/2024/03/28/genais-carbon-footprint-a-new-challenge-for-corporations/> (2024).
6. Zewe, A. Explained: Generative AI's environmental impact. *MIT News* <https://news.mit.edu/2025/explained-generative-ai-environmental-impact-0117> (2025).
7. Vartziotis, T. et al. Carbon Footprint Evaluation of Code Generation through LLM as a Service. (2025).
8. International Energy Agency. *Energy and AI*. <https://www.iea.org/reports/energy-and-ai> (2025).
9. Luccioni, A. S., Strubell, E. & Crawford, K. From Efficiency Gains to Rebound Effects: The Problem of Jevons' Paradox in AI's Polarized Environmental Debate. (2025).
10. Ertel, W. & Bonenberger, C. Rebound Effects Caused by Artificial Intelligence and Automation in Private Life and Industry. *Sustainability* **17**, (2025).
11. Wilson, C., et al. Potential climate benefits of digital consumer innovations. *Annu Rev Environ Resour* **45**, 13.1-13.32 (2020).
12. Wilson, C. et al. Evidence synthesis of indirect impacts of digitalisation on energy and emissions. in *ICT4S 2024: ICT for Sustainability* (Stockholm, Sweden, 2024).
13. Pettifor, H., Agnew, M., Wilson, C. & Niamir, L. Disentangling the carbon emissions impact of digital consumer innovations. *J Clean Prod* **485**, (2024).
14. Horner, N. C., Shehabi, A. & Azevedo, I. L. Known unknowns: indirect energy effects of information and communication technology. *Environmental Research Letters* **11**, 103001 (2016).

15. Bremer, C. et al. How Viable are Energy Savings in Smart Homes? A Call to Embrace Rebound Effects in Sustainable HCI. *ACM Journal on Computing and Sustainable Societies* **1**, 1–24 (2023).
16. Neves, C., Oliveira, T. & Santini, F. Sustainable technologies adoption research: A weight and meta-analysis. *Renewable and Sustainable Energy Reviews* **165**, (2022).
17. Maseeh, H. I. et al. Exploring the privacy concerns of smartphone app users: a qualitative approach. *Marketing Intelligence and Planning* **41**, 945–969 (2023).
18. Kelly, L., Kerr, G. & Drennan, J. Privacy concerns on social networking sites: a longitudinal study. *Journal of Marketing Management* **33**, 1465–1489 (2017).
19. Jang, C. & Sung, W. J. Beyond the privacy paradox: The moderating effect of online privacy concerns on online service use behavior. *Telematics and Informatics* **65**, (2021).
20. Paris, M. ChatGPT Hits 100 Million Users, Google Invests in AI Bot and CatGPT Goes Viral. *Forbes* <https://www.forbes.com/sites/martineparis/2023/02/03/chatgpt-hits-100-million-microsoft-unleashes-ai-bots-and-catgpt-goes-viral/> (2023).
21. Forsdick, S. The charts that defined 2023 – from the rise of AI to empty offices. *Raconteur* <https://www.raconteur.net/economy-trends/business-stories-2023-dat> (2023).
22. Deering, I. The big debate: is generative AI over-hyped? *Raconteur* <https://www.raconteur.net/technology/big-debate-generative-ai-hype> (2024).
23. McKinsey. *The State of AI in 2023: Generative AI's Breakout Year*. <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai-in-2023-generative-ais-breakout-year> (2023).
24. ONS. *Public Awareness, Opinions and Expectations about Artificial Intelligence*. <https://www.ons.gov.uk/businessindustryandtrade/itandinternetindustry/articles/publicawarenessopinionsandexpectationsaboutartificialintelligence/julytooctober2023> (2023).
25. Deloitte. *Europeans Are Optimistic about generative AI but There Is More to do to Close the Trust Gap*. <https://www2.deloitte.com/us/en/insights/topics/digital-transformation/trust-in-generative-ai-in-europe.html> (2024).
26. European Commission. *The Digital Decade. Eurobarometer Survey*. <https://europa.eu/eurobarometer/surveys/detail/3174> (2024).
27. Ipsos. *A 31-Country Global Advisor Survey How People across the World Feel about Artificial Intelligence and Expect It Will Impact Their Life*. (2023).
28. UKGOV. *Public Attitudes to Data and AI: Tracker Survey (Wave 4) Report*. <https://www.gov.uk/government/publications/public-attitudes-> (2024).
29. BBC. Facebook fined £500,000 for Cambridge Analytica scandal. *BBC Technology* (2018).
30. Hern, A. TechScape: Clearview AI was fined £7.5m for brazenly harvesting your data – does it care? *The Guardian* (2022).
31. Gaffney, O. et al. The Earth alignment principle for artificial intelligence. *Nature Sustainability* (2025).
32. UN. *Governing AI for Humanity*. <https://www.un.org/en/ai-advisory-body> (2024).
33. Khowaja, S. A., et al. ChatGPT Needs SPADE (Sustainability, PrivAcy, Digital divide, and Ethics) Evaluation: A Review. *Cognitive Computation* **16**, 2528–2550 (2024).
34. Wilson, C., Vrain, E. & Andrews, B. Consumer Uptake of Digital Low-Carbon Innovations. in *ICT4S* (2022).
35. Pettifor, H. et al. Are low-carbon innovations appealing? A typology of functional, symbolic, private and public attributes. *Energy Res Soc Sci* **64**, 101422 (2020).

36. Ada Lovelace Institute and The Alan Turing Institute. *How Do People Feel about AI? A Nationally Representative Survey of Public Attitudes to Artificial Intelligence in Britain*. <https://adalovelaceinstitute.org/report/public-attitudes-ai> (2023).
37. Acquisti, A., Brandimarte, L. & Loewenstein, G. Privacy and Human Behavior in the Information Age. *Science* **347**, 509 (2015).
38. Norberg, P. A., Horne, D. R. & Horne, D. A. The privacy paradox: Personal information disclosure intentions versus behaviors. *Journal of Consumer Affairs* **41**, 100–126 (2007).
39. Barth, S. & de Jong, M. D. T. The privacy paradox – Investigating discrepancies between expressed privacy concerns and actual online behavior – A systematic literature review. *Telematics and Informatics* **34**, 1038–1058 (2017).
40. Ljubljana, M. How Overhyped is ChatGPT in 2023? The Truth about its Hype Explained. *Shape Labs* <https://shape-labs.com/articles/how-overhyped-is-chatGPT> (2023).
41. Kovanovic, V. Generative AI hype is ending – and now the technology might actually become useful. (2024).
42. Anirudh VK. How ChatGPT Broke the AI Hype Cycle. *AIM* <https://analyticsindiamag.com/ai-origins-evolution/how-chatgpt-broke-the-ai-hype-cycle/> (2024).
43. Gerber, N., Gerber, P. & Volkamer, M. Explaining the privacy paradox: A systematic review of literature investigating privacy attitude and behavior. *Computers and Security* **77**, 226–261 (2018).
44. Saliju, L. & Deboi, V. Unraveling the Paradox: Balancing Personalization and Privacy in AI-Driven Technologies. (Jonkoping University, 2023).
45. Graham, J. et al. The AI Risk Repository: A Comprehensive Meta-Review, Database, and Taxonomy of Risks from Artificial Intelligence. (2024) doi:10.48550/arXiv.2408.12622.
46. Cui, J. & Lim, C. K. The Privacy Concerns’ Influences on Bike-Sharing Consumers’ Behavior. *Journal of System and Management Sciences* **12**, 252–266 (2022).
47. Bliss, L. Are Dockless Bikes a Cybersecurity Threat? *Bloomberg* <https://www.bloomberg.com/news/articles/2018-02-15/dockless-bikesharing-and-china-s-cybersecurity-laws> (2018).
48. Morrison, O. Beware consumer backlash about data sharing, food apps told. *Food Navigator Europe* https://www.foodnavigator.com/Article/2020/08/24/Beware-consumer-backlash-about-data-sharing-food-apps-told/?utm_source=chatgpt.com (2020).
49. Demsey, J. HelloFresh fined over millions of spam texts and emails. *BBC* <https://www.bbc.co.uk/news/business-67949521> (2024).
50. Bossert, L. N. & Loh, W. Why the carbon footprint of generative large language models alone will not help us assess their sustainability. *Nature Machine Intelligence* **7**, 164–165 (2025).
51. Tang, F. & Østvold, B. M. User Interaction Data in Apps: Comparing Policy Claims to Implementations. in *18th IFIP Summer School on Privacy and Identity Management (IFIPSC 2023)*. doi:10.48550/arXiv.2312.02710.
52. Steedman, R., Kennedy, H. & Jones, R. Complex ecologies of trust in data practices and data-driven systems. *Inf Commun Soc* **23**, 817–832 (2020).
53. GPAI. *Climate Change and AI: Recommendations for Government Action*. (2021).
54. Dafoe, A. & Russell, S. Yes, we are worried about the existential risk of artificial intelligence. *MIT Technology Review* <https://www.technologyreview.com/2016/11/02/156285/yes-we-are-worried-about-the-existential-risk-of-artificialintelligence/> (2016).
55. Center for AI Safety. Statement on AI risk. <https://www.safe.ai/work/statement-on-ai-risk> (2023).

56. Gyevnár, B. & Kasirzadeh, A. AI safety for everyone. *Nature Machine Intelligence* **7**, 531-542 (2025).
57. Lazar, S. & Nelson, A. AI safety on whose terms? *Science* **381**, 138 (2023).
58. Climate Change Committee. *The Seventh Carbon Budget Advice for the UK Government*. www.theccc.org.uk/publications (2025).
59. Blank, G., Dutton, W. H. & Lefkowitz, J. *Perceived Threats to Privacy Online: The Internet in Britain. Oxford Internet Survey 2019*. (2019).
60. Potoglou, D., Dunkerley, F., Patil, S. & Robinson, N. Public preferences for internet surveillance, data retention and privacy enhancing services: Evidence from a pan-European study. *Comput Human Behav* **75**, 811–825 (2017).
61. Dinev & Hart. Privacy Concerns and Levels of Information Exchange: An Empirical Investigation of Intended e-Services Use. *e-Service Journal* **4**, 25 (2006).
62. Blank, G. *2019 OxIS Questionnaire*. <https://oxis.oii.ox.ac.uk/research/questionnaires/> (2019).
63. Park, Y. J. Offline Status, Online Status: Reproduction of Social Categories in Personal Information Skill and Knowledge. *Soc Sci Comput Rev* **31**, 680–702 (2013).
64. Finch, J. The Vignette Technique in Survey Research. *Sociology* **21**, 105–14 (1987).
65. Barter, C. & Renold, E. *Social Research Update: The Use of Vignettes in Qualitative Research*. <https://sru.soc.surrey.ac.uk/SRU25.html> (1999).
66. Byrne, B. M. *Structural Equation Modelling With AMOS*. (Routledge, 2013).