Government Information Quarterly 40 (2023) 101876

Contents lists available at ScienceDirect

Government Information Quarterly

journal homepage: www.elsevier.com/locate/govinf

Citizens' acceptance of artificial intelligence in public services: Evidence from a conjoint experiment about processing permit applications

Laszlo Horvath [a,*], Oliver James [b], Susan Banducci [b], Ana Beduschi [c]

[a] School of Social Sciences, Birkbeck, University of London, Malet St, London WC1E 7HX, United Kingdom
[b] Department of Social and Political Sciences, Philosophy, and Anthropology, University of Exeter, Northcote House, The Queens Drive, Exeter Devon EX4 4QJ United Kingdom
[c] Law School, University of Exeter, Northcote House, The Queens Drive, Exeter Devon EX4 4QJ United Kingdom

ARTICLE INFO

Keywords:
Artificial intelligence
Public service
Human involvement
Conjoint experiment

ABSTRACT

Citizens' acceptance of artificial intelligence (AI) in public service delivery is important for its legitimate and effective use by government. Human involvement in AI systems has been suggested as a way to boost citizens' acceptance and perceptions of these systems' fairness. However, there is little empirical evidence to assess these claims. To address this gap, we conducted a pre-registered conjoint experiment in the UK regarding acceptance of AI in processing public permits: for immigration visas and parking permits. We hypothesise that greater human involvement boosts acceptance of AI in decision-making and associated perceptions of its fairness. We further hypothesise that greater human involvement mitigates the negative impact of certain AI features, such as inaccuracy, high cost, or data sharing. From our study, we find that more human involvement tends to increase acceptance, and that perceptions of fairness were less influenced. Yet, when substantial human discretion was introduced in parking permit scenarios, respondents preferred more limited human input. We found little evidence that human involvement moderates the impact of AI's unfavourable attributes. System-level factors such as high accuracy, the presence of an appeals system, increased transparency, reduced cost, non-sharing of data, and the absence of private company involvement all boost both acceptance and perceived procedural fairness. We find limited evidence that individual characteristics affect these results. The findings show how the design of AI systems can increase its acceptability to citizens for use in public services.

1. Introduction

The use of artificial intelligence (AI) in the delivery of public services is growing. Although the specific forms of AI technology vary, they have in common the use of computer algorithms as sets of rules to enable the automated processing of data for routine bureaucratic decision-making. This technology is seen by policy makers as a way to increase the effectiveness of services, reduce costs, and, consequently, to improve efficiency (Engin & Treleaven, 2019; European Commission, 2020; Janssen & Kuk, 2016; Veale & Brass, 2019). However, in democratic politics and policy-making, the use of AI goes beyond technical concerns and impacts state accountability. Therefore, it is necessary to move beyond the simplistic policy rhetoric of using AI for "the public good" to consider AI's acceptance as part of politico-social relations in the design of systems (Taddeo & Floridi, 2018). Studies have found that most respondents to public opinion surveys express concerns about AI across a range of different domains (Smith, 2018). The degree of acceptance depends on the specific form of AI and the context for its use (Kitchin, 2017; Smith, 2018). Despite these findings, there is relatively little research directly about how features of the design of AI systems affect citizens' acceptance of its use in the context of public services, a gap in knowledge our research aims to address.

We contribute to understanding citizens' acceptance of AI in public services which should help inform the future design of such systems in democracies. We define acceptance of AI as citizens' agreement that the use of AI in decision-making is appropriate for that service. Such acceptance is important for activities where public authorities provide services using the state's powers to coerce citizens or grant privileges to individuals or groups (Janssen & Kuk, 2016; Levy, 2018; Taylor, 2016). Activities of this kind include using public regulations to issue permits that give individuals a particular status or allow them to undertake activities (Busch & Henriksen, 2018). We focus on AI use in systems for

* Corresponding author at: 10 Gower Street, WC1E 6DP London, UK.
E-mail address: L.Horvath@bbk.ac.uk (L. Horvath).

https://doi.org/10.1016/j.giq.2023.101876
Received 3 June 2021; Received in revised form 2 October 2023; Accepted 20 October 2023
Available online 28 October 2023
0740-624X/

issuing public permits to individuals within two contexts: the processing of national visas for immigration and the allocation of local car parking permits. Citizens' acceptance of AI in these contexts matters because, if they do not have confidence in the procedures, they may not use the services, abide by the decisions (Aoki, 2020), or recognise the legitimacy of the decisions made by such processes (Tyler, 2006).

The acceptance of AI across a range of contexts has been linked to several factors, especially perceptions of the procedural fairness of systems utilizing AI (for a review, see Starke, Baleis, Keller, & Marcinkowski, 2022). Procedural fairness emphasises the appropriateness of the procedures considering the criteria used and their application to the individuals affected, rather than the outcomes of decisions themselves. Procedural fairness has previously been identified as a key aspect of the acceptance of laws and public regulations (Tyler, 2006) and is important in public services because it influences how citizens perceive the legitimacy of systems and whether they comply with judgements reached by them. Incorporating AI into how public service decisions are made entails important changes to procedures. Understanding citizens' perceptions of fairness is a key part of developing a 'society-in-the-loop' approach to the use of AI, to examine the relationship between societal values and the values embedded in the design of such systems (Rahwan, 2018; Starke et al., 2022). For these reasons, we supplement our focus on citizens' acceptance by additionally examining citizens' perceptions of the fairness of procedures in permit processing systems incorporating AI.

The literature on the use of AI has increasingly found evidence about the effect of different features of systems on acceptance and perceptions of fairness. The role of human involvement has been identified as being of particular importance (Jones, 2017; Meijer & Grimmelikhuijsen, 2020). A review of the general literature on fairness and AI found that for contexts similar to ours, where contextual information influences decisions, people tended to prefer systems with human involvement compared to automated systems that exclude them (Starke et al., 2022). Maintaining human involvement is potentially important to public authorities' use of AI because it can alleviate concerns about full automation ensuring the systems appear procedurally fair to citizens. In this study, we incorporate this insight to assess whether involving humans in the systems increases acceptance and perceptions of fairness in the use of AI for processing permits.

In Section 2, we draw on theories about AI and human involvement in decision-making systems to generate hypotheses about citizens' views about AI's acceptability and procedural fairness in the context of processing decisions about issuing public permits. The literature suggests that maintaining human involvement makes the use of AI more acceptable to individuals. This leads us to the first set of hypotheses that citizens' concerns about AI will make its use less preferred, and make it seem less fair, the greater the proportion of AI (and the less human involvement) in the decision process (Lee, 2018; Smith, 2018). Developing this perspective further, we propose a second set of hypotheses: that more human involvement can boost support for, and increase the perceived fairness of, systems that in other respects have undesirable features. Specifically, we hypothesise that human involvement can reduce the negative effects of AI systems that are more costly, less accurate, have limited appeal procedures in the case of error, and involve private sector-developed AI technology. We also draw on the literatures on technology acceptance and knowledge about AI to move beyond the notion that all citizens share the same views about AI system features. Consequently, we examine individual-level moderators of the acceptance of systems flowing from differences in individuals' tendency to adopt technology, and differences in their knowledge about AI.

In Section 3, we introduce a pre-registered, conjoint experiment with UK residents as participants. In the experiment, we offered respondents choices about which form of a hybrid (that is, human plus AI) permit processing system to adopt. These choices allowed us to assess the effects of a range of features of AI on acceptability and perceived fairness simultaneously. To achieve a breadth of relevance for the findings, we presented three permit processing scenarios: processing a national visa application, processing a local car parking permit application where human administrators have little discretion, and processing the same applications but human administrators have more discretion.

Section 4 sets out our findings. In two out of the three cases of permit processing, respondents preferred systems with greater human involvement. Yet, when substantial human discretion was introduced in the second of the parking permit scenarios, they preferred more limited human input. We found little evidence that individual-level covariates moderate these effects, with the exception that respondents' preferences for human involvement in visa decisions changed depending on their propensity to adopt technology in other domains. In Section 5, we highlight the conclusions of the research and the contribution of the findings to wider debates about AI system design and technology acceptance in the public domain. We also discuss the limitations of the study and how these findings open directions for future research.

2. Artificial intelligence and public services

AI offers transformative potential for citizen-state interactions. Defined broadly, AI allows computational tools to perform tasks that otherwise would require human intelligence (Taddeo & Floridi, 2018). These technologies are differentiated from traditional decision-making systems because they learn from existing data using algorithms, rather than relying solely on predefined rules. Such technology is now prevalent in governance at the macro-, meso-, and "street-levels," both automating and augmenting administrative tasks (Engin & Treleaven, 2019; Janssen & Kuk, 2016; Veale & Brass, 2019; and European Commission, 2020; p. 2.). According to a recent report, governments increasingly view AI technologies as tools to that can "maximise service delivery and target early intervention as a way of saving resources" (Centre for Data Ethics and Innovation, 2020; p. 74).

We focus on the use of AI where the state uses its public authority for issuing permits to citizens, specifically in the cases of visas and parking permits. When public authorities process permit applications they share key characteristics with the broader operation of legal systems, as the permits are issued based on specific legal rules and standards (Tyler, 2006). In this context, AI enables different degrees of automation in the decision-making process to decide if specific individuals will be granted permit based on complex data and rules (for example, document verification based on image recognition, or pattern recognition for fraud). In what follows, we consider a theoretical framework about citizens' acceptance and perceived procedural fairness of AI in public services for permit applications. We highlight the role of human involvement in influencing acceptance and perceptions of AI and the impact of additional individual citizens' characteristics that potentially moderate the relationship between human involvement and AI acceptance.

2.1. Acceptance of AI in public services

Our focus is on the acceptance of AI by citizens for use in the public service of issuing permits. Acceptance is citizens' agreement that the use of AI is appropriate for the service. More generally, the acceptability of AI has been found to be sensitive to the specific domain of its use (Smith, 2018). Chohan and Hu (2020) survey recent developments in digital government services and argue for the need for research that recognises the specific 'public' sector context in which services are provided. Our approach draws on the Technology Acceptance Model, a framework used to explain citizens' support of digital technologies across different domains of application (Davis, 1986). Accordingly, users' perceptions that a piece of technology is "useful," and/or "easy to use" would lead to support for its use.

Digital technologies are transforming the way that users interact with a wide range of organisations providing services. In the context of publicly provided services such as parking permits and visas, a body of literature has started to clarify the additional contextual and individual-level determinants, such as the quality of services, transparency, and knowledge, that can lead to citizens' acceptance of digital government and Internet of Things (IoT) technologies. To date, much of the current literature is split into either examining the implementation of digital government initiatives compared to stated policy intentions, or work that examines the attitudes of citizens or other groups. Researchers have pointed out a gap with relatively little research that crosses this divide to evaluate how citizens assess these digital initiatives (Gil-Garcia & Flores-Zúñiga, 2020).

The need to focus on citizens' acceptance of procedures is important because public sector applications of AI typically rely on processing personal data. Citizens' views about the acceptability of such technology engages their potential concerns about its use—especially about the privacy of their data and how it will be used (e.g., Anthopoulos, Reddick, Giannakidou, & Mavridis, 2016; Horvath et al., 2022). We focus on citizens' views about the appropriateness of using AI in systems for issuing public permits—with the empirical contexts being permits for visas to allow immigration and permits for local parking. Such acceptance should be of concern for the state. On the one hand, for the state to deliver its services effectively, citizens need to accept its authority to either coerce or grant privileges in this manner (Taylor, 2016). On the other hand, politicians, and populists in particular, may mobilise sentiment about the loss of control to elites, including big tech companies involved in providing AI technology, in order to score political points (Janssen & Kuk, 2016; Levy, 2018).

Because of the public context, the perceived procedural fairness of decision processes using AI is particularly important for citizens' acceptance. Procedural fairness is based on the perception that the criteria used, and their application to affected individuals, are appropriate, rather than on the outcomes of the decisions themselves. Fairness perceptions in general have previously been found to be important for the public acceptability of AI (Smith, 2018). However, procedural fairness is especially important for issuing public permits because it involves the exercise of state powers that in, a democratic society, need to have public legitimacy. Such procedural fairness has previously been identified as a key aspect of the acceptance of laws and public regulations more generally (Tyler, 2006), although the issue has received little previous attention in the context we explore. However, reflecting these concerns, algorithmic decision-making is becoming increasingly evaluated on fairness grounds by public servants and contractors providing AI (Veale, Van Kleek, & Binns, 2018). For these reasons, we focus the outcomes of AI in systems for both citizens' acceptability and perceived procedural fairness of the decision-making systems.

2.2. Human involvement and AI in public services

A growing body of research about AI points to people seeing the involvement of humans in decision making as being required to protect the human dignity of those subject to regulations (Dietvorst, Simmons, & Massey, 2015; Jones, 2017; Meijer & Grimmelikhuijsen, 2020). An important aspect of dignity when rules are being applied is the fairness of procedures for applicants subject to the decision-making procedures. For this reason, we assess the potential for the degree of human involvement in the process of decision making in public services as a way to make these systems more acceptable and perceived to be procedurally fairer to citizens.

Experimental evidence has found a degree of 'algorithm aversion' when people are offered a choice between human or algorithmic processing in decision making. These findings are consistent with people seeing humans as a requirement for such decision making (Dietvorst et al., 2015). Having a 'human in the loop' is expected to increase the acceptability of using AI in hybrid decision making, particularly in sensitive issues or unique circumstances (Centre for Data Ethics and Innovation, 2020; p. 6.). On this basis, we hypothesise that human involvement in the decision process, which entails less reliance on automation, increases its acceptability to citizens. By implication, reducing the proportion of human involvement by administrators decreases acceptance of the system. This leads us to the first hypothesis:

H1. Increasing the proportion of human administrative involvement in the decision process increases acceptance and perceptions of fairness.

The impact of human involvement in procedures further extends to influencing other factors. A body of previous work suggests several characteristics of AI can affect its acceptability. First, a lack of information about how the systems work can reduce acceptance and some research has found that this can be mitigated by providing the public with more transparent information about the operation of AI systems (Brauneis & Goodman, 2018). Second, having a public rather than a private company as the source of the AI technology can boost acceptance. This is partly due to concerns about private profit jeopardising the public interest and also because of the lower costs associated with administering the overall process (Johnson & Verdicchio, 2017; Meijer & Grimmelikhuijsen, 2020; Taddeo & Floridi, 2018). Third, within the public sector, the accuracy of decisions (or the proportion of accurate predictions made by algorithms) is seen as a significant benefit. Accuracy is likely to increase support for their use (Meijer & Wessels, 2019). Fourth, the availability of appeals processes influences acceptability, with AI decision outcomes being subject to potential appeals increasing trust (Harrison & Luna-Reyes, 2020, p. 12) and, by extension, increasing overall acceptance of such systems.

We examine how human involvement conditions the impact of these other four AI characteristics. The baseline expectation is that negative system features, including a lack of accuracy, reduced transparency, limited appeals, and privately developed systems, will reduce acceptance, and lower perceived fairness. However, we suggest that having more human involvement in issuing permits would mitigate the presence of these "undesirable" system characteristics that would otherwise reduce acceptance:

H2. Increasing human involvement can further enhance the acceptance and perceived fairness of systems that are less accurate, more costly, less transparent, have limited appeal options, or are privately developed rather than publicly.

2.3. Moderators of the need for human involvement: Knowledge about AI and tendency to adopt technology

Besides characteristics of the technology itself, several studies have worked on identifying individual-level characteristics that can explain technology acceptance across different domains. Digital skills and literacy have been found useful to explain adoption of new technologies (Yu, Lin, & Liao, 2017). In the public domain, Gil-Garcia and Flores-Zúñiga (2020) developed a model linking digital technology implementation by governments and adoption by the users of government services.

Expanding on this, we focus on 'AI literacy.' This competency covers the ability of recognising instances of AI and "distinguish[ing] between technological artefacts that use and do not use AI" (Long & Magerko, 2020; p. 4). Those with higher AI literacy are more knowledgeable about AI technology and its uses. Those who are more AI literate are expected to need less human involvement in the system in order to accept its decisions. Thus, we consider AI literacy as a moderator of the influence of human involvement on acceptability and perceived fairness of systems using AI:

H3. Being more knowledgeable about AI reduces the negative impact of less human involvement on acceptance of the system and perceived fairness.

Whilst there are potentially many relevant factors influencing the need for humans in hybrid systems, an additional variable recently argued to be particularly relevant to the context of technology acceptance is individuals' "previous use" of technology. The proliferation of digital technology in everyday life, including in the public domain, results in a variation in how familiar individuals are with digital technology prior to encountering our specific case of the use of AI. A similar point is made in Guppy et al. (2022) about a "digital disconnect" prior the Covid-19 pandemic, in which the authors hypothesised (but found no evidence) that only teachers who were previously adopters of digital learning technologies could successfully transition to online teaching during the pandemic. Here we formulate the following hypothesis about technology adoption as a second moderator of both acceptance and perceived fairness of systems using AI that have different degrees of human involvement:

H4. Those with higher tendency to adopt technology over traditional methods relying on in-person contact will be less sensitive to human involvement in their acceptance of the system and perceived fairness.

We have summarized our research hypotheses in Fig. 1 showing the direct effects of human involvement in decision making and how we hypothesise that this will be moderated by both knowledge and technology adoption.

2.4. Public permit processing and administrative discretion

Our theoretical framework, which includes human involvement in AI systems, is potentially applicable across several public service domains. However, our primary focus is on permits as a public service, an under-researched aspect of state activity when it comes to the impacts of AI. The use of AI for permit decision processing affects a core activity of the state informed by public regulations (Busch, André & Zinner, 2018). These are typically complex applications of rules to particular cases and offer the potential for either more manual approaches using administrators to make decisions, or for more automation including through the use of AI. The use of more rule-based systems builds on long running attempts to automate public administration systems (Margetts, 1999).

We focus our empirical work on a set of different kinds of public permit processing in order to make the findings more general across contexts within this public service domain. In the research design we present below, we create different experimental scenarios that cover different kinds of permit decision processing (the contexts of local parking permits and national government visa processing). These different cases enable us to ask our research participants to consider AI implemented in systems that they might themselves use on a daily basis (local parking), and in systems that are mostly about others such as non-citizens (national visa processing).

We use two forms of parking permit systems because permit processing can vary in the extent to which administrative discretion is present—the degree of autonomy available for bureaucrats when they implement policy (Lipsky, 1980). Considering the effect of digital technology on "frontline discretion," Buffat (2015) argues that both limited and amplified forms of bureaucratic discretion are possible developments in digital government. Looking at AI technology in particular, Criado, Valero, and Villodre (2020) demonstrate a case when AI positively impacted discretionary power, as well as note that involving administrators who can apply discretion can increase algorithmic transparency. The extent of substantive human discretion, as opposed to "artificial discretion" (Young, Bullock, & Lecy, 2019) leads us to ask whether citizens support more human involvement in different kinds of permit processing. For this reason, we include two examples of parking permit processing, differing in whether the public official can potentially exercise a larger or smaller amount of discretion in issuing the permits.

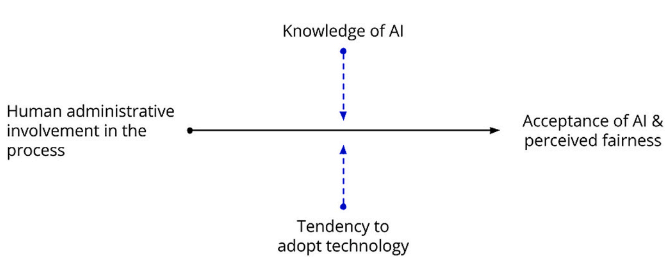

Fig. 1. Research model and summary of hypotheses.

3. Methods

We administered a series of three choice (conjoint) experiment tasks to a sample of citizens in the UK. The experiment was embedded in a broader survey that also asked questions about other issues, specifically the UK 2019 General Election and general socio-political attitudes. We pre-registered the research design and our hypotheses on AsPredicted, Report Number #33170 [report attached for peer review]. Participants provided informed consent to participate in the survey. Our research design received ethical approval from the Institutional Ethics Committee.

3.1. Data and subjects

Our study uses an online panel of $N = 2143$ provided by Dynata, a recognised research panel vendor who directed their respondents to our survey, hosted on our institutional platform.[1] All respondents were part of the experiment, however, the responses are subject to missingness – an issue that we discuss at the end of Section 3.2. The respondent characteristics mirror the UK's demographic makeup as of its 2011 census, also presented in more detail in Appendix A. The surveys were completed between 19th and 23rd December 2019. In addition to the measures presented in the following sections, we also included data quality checks, following the recommendations in Berinsky et al. (2013) which we report in Appendix A, with robustness checks using these data quality measures in Appendix C.

3.2. Conjoint experiment

We administered three conjoint experiment tasks offering respondents choices about which form of hybrid permit processing system to adopt. The three tasks consisted of five trials each, presenting participants with a series of five pairwise comparisons of permit processing systems, displayed in a table format. The task for respondents was to select which one of the two systems they preferred. All systems incorporated a form of AI and in all tasks the systems were described as processing "thousands of applications each year and the decision-making process can be assisted by artificial intelligence". The description further set out "artificial intelligence in the form of a computer algorithm—that is, a set of rules to enable the automated processing of data for the application, making it a simple."

The three tasks differed in the domain in which AI was implemented (as discussed in Section 2.4). All participants completed Task 1 where they were presented with visa permit systems. While individual citizens are not directly subject to visa decisions within their own country, they have a vested interest due to its general importance for a political community, and the implications of who is allowed entry. In the scenario, we told respondents that a central government department, the Home Office, would use AI when deciding on visa permits for work in Britain for extended periods of time (longer than six months).

In the second case, individual citizens can be directly affected by the process of applying for a residential parking permit. Here, we examine high and low discretion contexts for issuing the permit. The high

[1] There are a number of online panels available to researchers. The advantages and disadvantages of these platforms for conducting experiments have been reviewed (Weinberg, Freese, & McElhattan, 2014). We have opted for Dynata which has been recommended for online studies (Reviewed in https://brl.mit.edu/researcher/online-studies/).

discretion case allows administrators more autonomy in decision making, while the low discretion context reduces the scope of human agency of the officials involved in the process. In Task 2A, the low discretion case, we said permits were given under a policy based on location, giving local governments "limited discretion in whether to issue the permit." In Task 2B, the high discretion cases, permits were given based on need and caring responsibilities, a policy giving local governments "a large amount of discretion in whether to issue the permit." We used a split sample, meaning participants were randomly divided between the two parking permit tasks, as they were very similar to each other. The alternative design of making all respondents do both parking permit tasks in a sequence may have risked experimenter demand effects or other biases when respondents are asked to complete similar tasks in sequence.

Acceptance of AI is indicated by stated preferences for use of AI across the six attributes, independently randomised and presented in pairs of alternative systems. As well as the extent of human involvement in the processing, we include the attributes set out Section 2.2 (above) where we discussed factors that have been found to influence the acceptability of use of AI. These are: the rate of accuracy of the AI process, having an appeals process for people affected by a decision, the transparency of the system, and costs of the processing system. The fifth attribute differs across the visa and parking tasks. In the parking permit context, we included data sharing beyond the immediate application process because it directly relates to the respondent's personal data. For the visa processing task, which involves personal data of unknown applicants, we instead listed the organisations involved in the AI system. We varied these organisations to include the government, a private firm or a university. We show the full list of attributes along with the attribute levels in Table 1, and give an example of what participants viewed when responding to the task about visa processing in Fig. 2.

We note that our process of randomisation across these attributes was successful given an even distribution of what was displayed across the sample, and we found no link between the attribute level displayed and participant demographics (age, gender, and education) either. When conducting the experiment, 185 or 8.6% of our respondents encountered a display error on one of the above attributes during the study. Data for these respondents is missing at random which we confirm by noting the successful randomisation as outlined above, as well as by regressing demographics on missingness. We found no significant patterns, although these respondents were on average 2.27 years younger than those who did not encounter these errors, $t(178.01) = 1.57$, $p = 0.12$.

Table 1
Overview of system attributes and attribute levels displayed to participants.

Attributes & attribute levels			
Involvement of Home Office (visa task) or local government (parking tasks) administrators*			
An administrator processes the application with 100% of the decision made by the computer algorithm	An administrator processes the application with 75% of the decision made by the computer algorithm	An administrator processes the application with 50% of the decision made by the computer algorithm	An administrator processes the application with 25% of the decision made by the computer algorithm
Accuracy rate of the computer algorithm			
80% of decisions by the computer algorithm are correct	90% of decisions by the computer algorithm are correct	99% of decisions by the computer algorithm are correct	
Computer algorithm developed by (visa task only)			
Home office	Private company	University researchers	
Cost to permit applicant for each application			
£10	£50 (parking task only)	£100	£1000 (visa task only)
Information about how the computer algorithm works			
Available upon request	Not available upon request		
Use of your personal data (parking tasks only)			
Personal data used only for decision and not retained		Personal data retained and shared with other government departments	

* We reversed the direction of this variable in the analysis to reflect the amount of human involvement thus the variable ranges "None" to "75%".

3.3. Dependent variable – AI acceptance & perceived fairness

Following standard practice in conjoint experiments, we use the binary choice (selected vs not selected) made across two competing profiles as the dependent variable (in our case, decision processes involving AI) to estimate acceptance of the system. An advantage of the forced-choice setup is that participants are invited to consider a series of trade-offs that forces them to reveal preferences across a series of tasks, making it impossible to support alternative systems equally. In addition, we asked respondents to indicate the perceived fairness of each system as a process to issue public permits on a 7-point scale under each profile. The additional question takes away the constraint of choice and allows us to compare the results with this attitudinal measure.

3.4. Analytical framework

The conjoint experiment produced data on the choices that respondents made in the survey across the different permit processing profiles. We used this choice data in statistical models to infer the importance that respondents attach to the attributes of the systems by estimating Average Marginal Component Effects (AMCEs). In this way, the AMCE estimates allow us statistically to test our hypotheses about the effect of changing levels of the attributes, and about moderators of these relationships (Hainmueller, Hopkins, & Yamamoto, 2014). The model coefficients predict the increase in the population probability that the system with a given attribute level is chosen over the reference category (pp. 10–12) including component interactions which we illustrate more in detail in Section 4.2 about these interactions.

3.5. Moderators

Our key moderators affecting the relationship between human involvement and acceptance of systems are individuals' tendency to adopt technology and their knowledge about AI. We define high tendency to adopt technology as a willingness to opt for technological solutions, while those with low tendency to adopt technology will tend to opt for more traditional methods, such as face to face contact. In the survey, prior to the conjoint experiment, we ask about tech adoption across three domains: parking permit application, reporting local issues such as potholes, and passport or visa applications. Either of the three were randomly asked for each participant, to be able to tap into tech adoption on a range of domains. The specific question after describing the domain was "Provided there is a smartphone app available to assist you with …, how likely is it that you would take advantage of the technology?" scaled 1, "full contact with administrator in person or via phone" to 7, "prefer to do process entirely on app." Tendency to adopt technology measured this way has a moderate negative correlation with age in our sample, $r(1976) = -0.25$, $p < 0.01$.

The second moderator, knowledge about AI, taps into respondents' ability to recognise processes as artificial intelligence. We adopted the Elements of AI questionnaire from the popular online MOOC (University of Helsinki & Reaktor, 2021) presenting six use cases to respondents, asking them to decide whether they were instances of AI or not. A non-AI example was "Spreadsheet that calculates sums and other pre-defined functions on given data," and there were five AI examples such as "A music recommendation system such as Spotify that suggests music based on the user's listening behaviour." We provide the list of quiz questions in Appendix A. For each individual, we derived an AI quiz score (factor

Please read the descriptions of Home Office decision processes below. Then indicate which of the two processes you would personally prefer to see used to decide whether a work visa should, or should not, be issued to an applicant.

	Process 1	**Process 2**
Cost to government for each application	Costs government £10 each to process	Costs government £100 each to process
Opportunity for applicant to appeal decision	Appeal of decision is possible	Appeal of decision is possible
Computer algorithm developed by	Home Office	University researchers
Involvement of the Home Office administrator and/or computer algorithm	An administrator processes the application with 100% of the decision made by the computer algorithm	An administrator processes the application with 100% of the decision made by the computer algorithm
Accuracy rate of the computer algorithm	99% of decisions are correct	80% of decisions are correct
Information about how the computer algorithm works	Details of the algorithm are available upon request	Details of the algorithm are not made publicly available

○ **Prefer decision process 1**

○ **Prefer decision process 2**

Fig. 2. An example trial (visa task).

scores) under an Item Response Theory framework. To our knowledge, the distribution of the Elements of AI quiz has not yet been described in the general population. Therefore, to explore construct validity, we check instead the following cognitive correlates: The latent AI quiz score correlates positively with education, $r(1865) = 0.21$, $p < 0.01$; with performance on the attention checks in the survey, $r(2035) = 0.18$, $p < 0.01$; but not with response time on survey, $r(2141) = -0.01$, $p < 0.86$. Across the two moderators, tendency to adopt technology and knowledge about AI, we find only a weak positive correlation, $r(1996) = 0.07$, $p < 0.01$.

While the main results regarding acceptance of systems as formulated in H1 and H2 are presented in the AMCE framework, subgroup preferences relating to tech adopters and those who are knowledgeable about AI are estimated with the Marginal Means method to avoid bias stemming from the choice of reference category within each subgroup, following the recommendation in Leeper, Hobolt, and Tilley (2020). Within this framework, we derive mean support to each system feature averaged across all other features.

4. Results

The four sections below relate to the four hypotheses presented in Section 2, assessing whether the results across three different conjoint tasks, representing three different contexts of AI-assisted public service provision, provide evidence in support of each of the hypotheses.

4.1. Average Marginal Component Effects (AMCE)

In this section, we focus on whether AI-assisted systems with greater human involvement are preferred, an attribute we randomised independently from the other attributes that would also potentially affect preference. We visualise the results for acceptance of the system as a dependent variable. We have repeated the analysis with the perceived fairness of systems and report these separately (see Appendix B for complete results of the models). In addition, we used the attention check questions (Appendix A) to perform a series of robustness analyses applying different thresholds on data quality, obtaining consistent results but, as expected, decreasing the precision of the estimates as the sample size drops (see Appendix C).

We found partial support for H1 in that increased human involvement increases acceptance for certain tasks: For two out of the three tasks, increasing human involvement to 75% translated into increased acceptance of the system. However, in the scenario where parking permits were given out based on a policy giving administrators a large amount of discretion in deciding who gets the permit, respondents preferred to cap human involvement at 25% (which we discuss further in Section 5). By contrast, when the policy gave local administrators only limited discretion, respondents reacted similarly to our baseline hypothesis and preferred systems with the highest (75%) human involvement. We show these effects in the context of all conjoint attributes in Fig. 3 and Table B1 in Appendix B. We note that the impacts associated with human involvement are small relative to the more influential effects of cost and the accuracy of the algorithm. Additionally, the magnitude of effects is even smaller for the fairness response scale. While we found a small but significant increase in perceived fairness with 75% human involvement in the visa task, but were unable to reject the null hypothesis when it came to the parking tasks, (see Table B2).

In terms of the other system attributes, we find clear support for our original expectations. Increased accuracy, having an appeals system, increasing transparency, lower cost, non-sharing of data and not having a private company involved in delivery all increase acceptance as shown in Fig. 3 below. These patterns are identical for perceived fairness of systems.

4.2. Average Component Interaction Effects (ACIE)

In this section, we discuss H2: whether undesirable system attributes, for example costly or less accurate systems, gain more support if embedded within systems with greater human involvement. To test this, we estimated component interaction effects. For consistency with the results presented earlier, we kept the same baseline category for each attribute. Our hypothesis would be supported if either:

(a) with increasing human involvement, the negative effects such as high cost move in the positive direction thus becoming “less rejected” over their respective reference category, or.
(b) with increasing human involvement, the positive effects such as high accuracy move in the negative direction thus no longer preferred as strongly over their respective reference category.

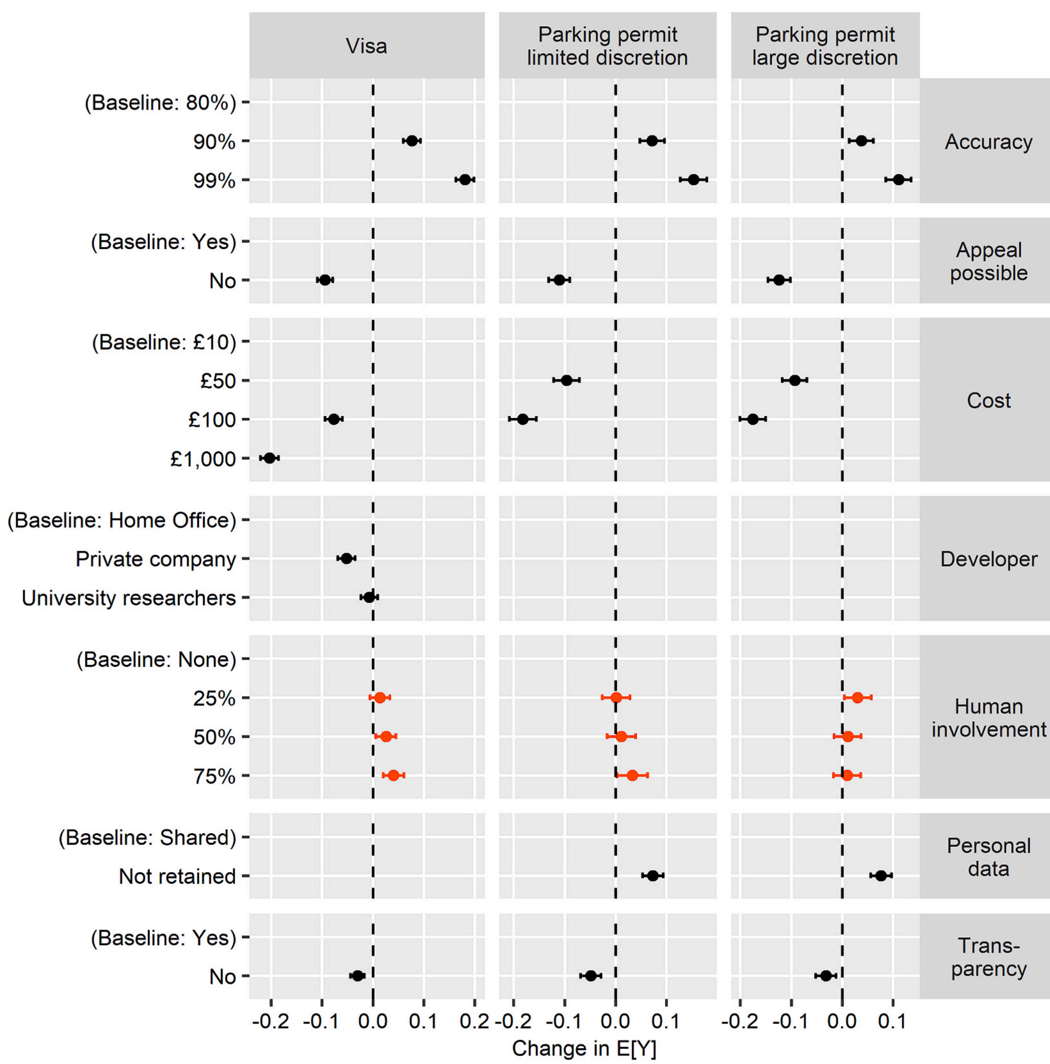

Fig. 3. Average Marginal Component Effects: Predicted increase in the population probability that the permit system with given attribute level is chosen (supported) over the reference category (Hainmueller, Hopkings & Yamamoto, 2014, pp. 10–12), across three tasks.

As system acceptance is measured in a binary choice task in this instance, (a) and (b) depend on which particular attribute level was selected as the reference category and therefore reflect the same process. Fig. 4 below shows these effects.

We found little evidence that human involvement moderated the impact of other attributes on acceptance in a systematic way. Notably, we found that costly parking permit systems, in which the local government had no discretion in deciding when to approve applications, were less rejected when displayed alongside the high level of human involvement. The result suggests that at least in this particular scenario, citizens accepted the higher costs associated with greater human involvement in the processing of permits.

4.3. Marginal means by tendency to adopt technology

In this section, we turn to the breakdown of results by respondent-varying characteristics to assess evidence for H3 and H4, that being more knowledgeable about AI and having a higher tendency to adopt technology, respectively, reduce the negative impact of less human involvement on acceptance and fairness. The characteristics were measured in the survey prior to administering the conjoint experiment. Respondents who scored high on tendency to adopt technology preferred making advantage of smartphone apps over traditional methods in a range of domains (see Section 3.5). We show the estimated Marginal Means as mean support for each system feature averaged across all other features (Leeper et al., 2020, p. 210) by tendency to adopt technology in Fig. 5 below.

Across two tasks, that the analysis shows that the choices made by the group with a high propensity to adopt technology were not influenced by the extent of human involvement (feature support near 50% across all levels of human involvement). By contrast, those with low tendency to adopt technology seem to be supporting systems with greater human involvement. We tested for differences in marginal means which is significant for visa systems ($F(4847.1,4) = 2.35$, $p = 0.04$) but not for either parking tasks either separately or pooled together ($F(4784.1,4) = 0.54$, $p = 0.70$). For parking systems with high human discretion, however, we have observed that greater human involvement is penalised regardless of citizens' tendency to adopt technology. As for the rest of the attributes, the effect of an individual's propensity to adopt technology seems negligible.

4.4. Marginal means by knowledge about AI

We asked if respondents' knowledge about artificial intelligence (ability to correctly identify examples of AI) moderated preferences regarding the extent of human involvement in hybrid systems. We show the marginal means of acceptance moderated by AI knowledge in Fig. 6 below.

The moderating impact of AI knowledge on preference for systems with greater human involvement is insignificant ($F(4787.1, 4) = 1.06$, p

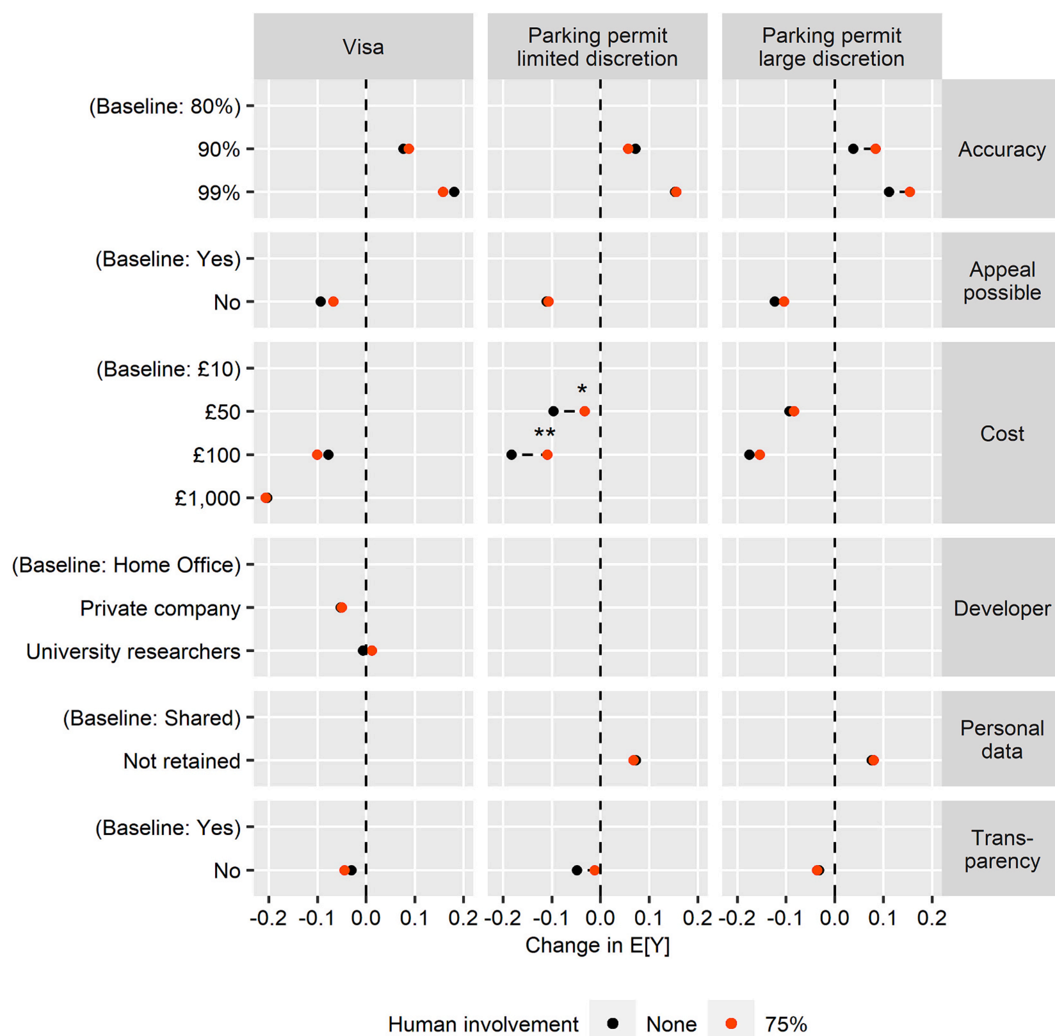

Fig. 4. Support for system attributes conditional on human involvement (Component interactions).

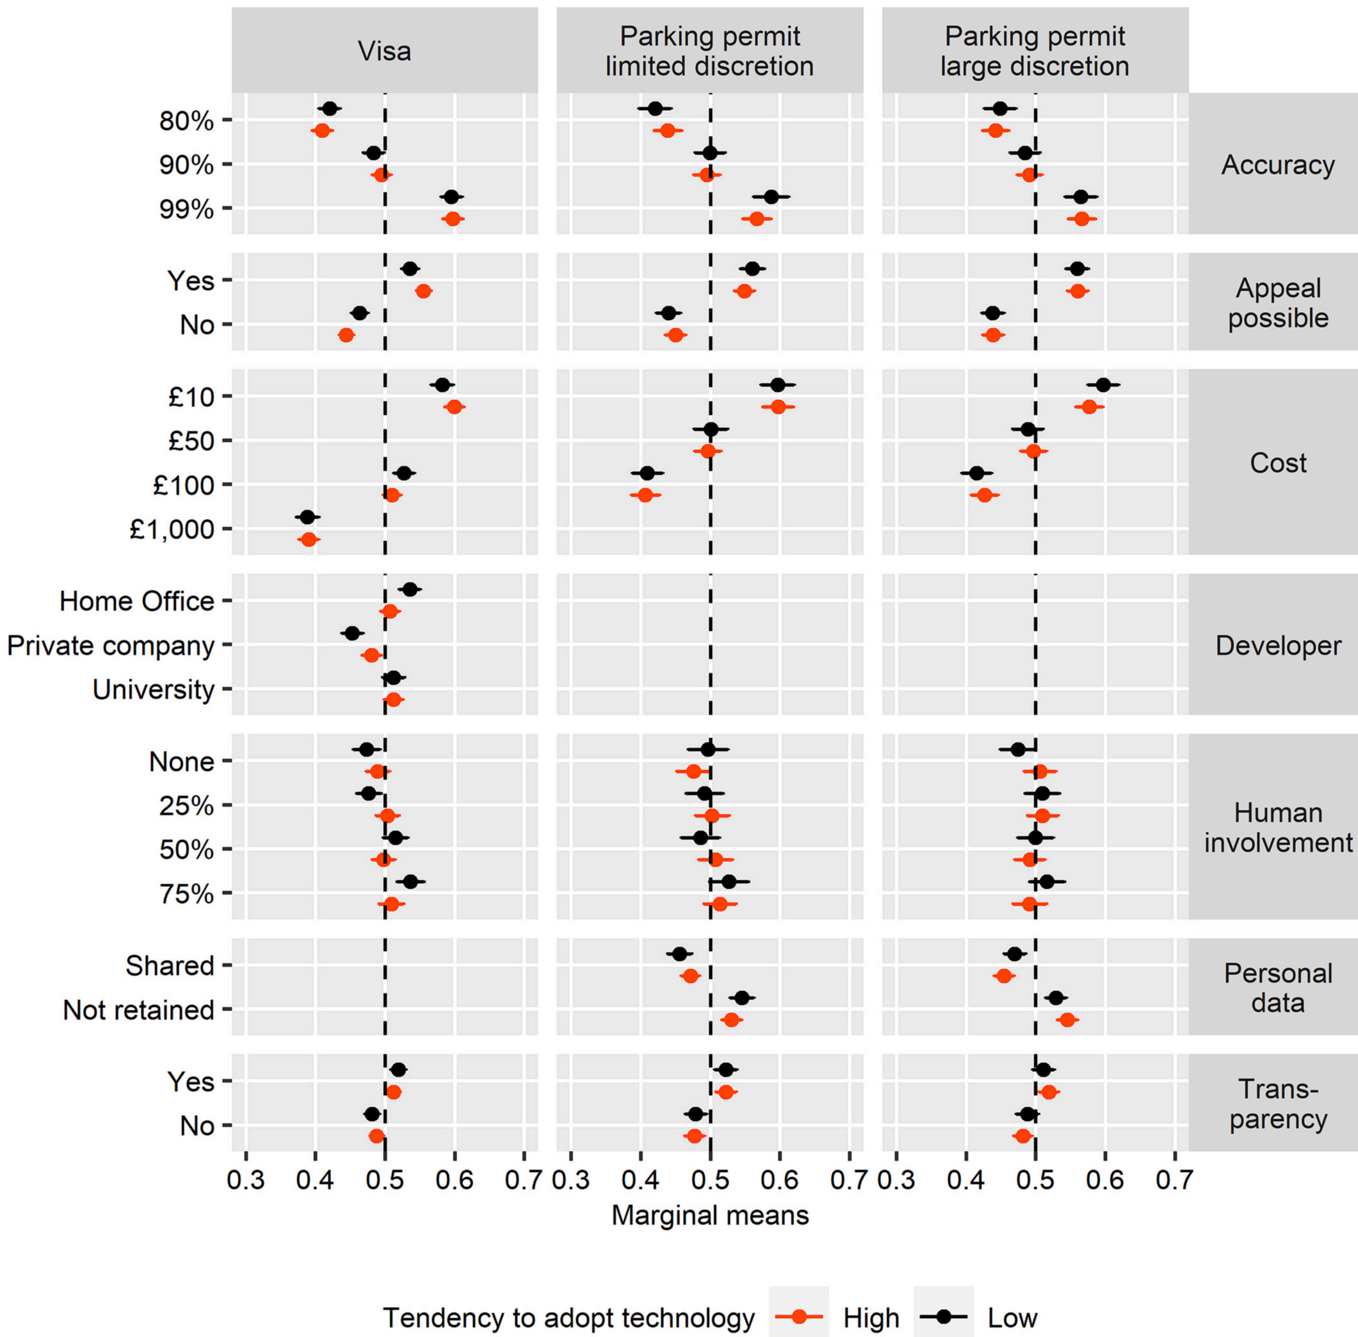

Fig. 5. Mean support for system attributes conditional on tendency to adopt technology, using Leeper, Hobolt and Tilley's method (2020).

= 0.37 in the visa task and $F(4850.9, 4) = 0.12$, $p = 0.97$ across the two parking tasks). Respondents with low knowledge about AI seemed to be opposing systems without human involvement even more than those with high AI knowledge. Nevertheless, the overall trend and pattern of effects remained consistent.

5. Discussion

If public authorities are to adopt AI assisted decision making, public acceptance will be an important aspect of its successful implementation. Our study provides empirical evidence about an important application of AI in the public sector, clarifying the factors leading to acceptance and perceived fairness. The limited support for the role of human involvement and lack of evidence that individual characterises such as AI literacy influence acceptance contrasts with the strong support for the other hypotheses about the influence of the characteristics of AI. These findings are consistent with research on AI in other domains and are reassuring that the null findings for other hypotheses are not the result of a failure of the research design in general.

In terms of the magnitude of effect sizes, our survey respondents appeared to be more strongly influenced by the costs and accuracy of the technology than by concerns about "humans in the loop," transparency, or even data sharing. This suggests citizens may perceive legitimacy more profoundly in terms of the system's efficiency and its ability to deliver accurate (and cost-effective) results. Future research needs to investigate this perspective further. As AI systems become more prevalent and integrated into citizen-state interactions, people might assess the trade-offs between efficiency and human feedback, including

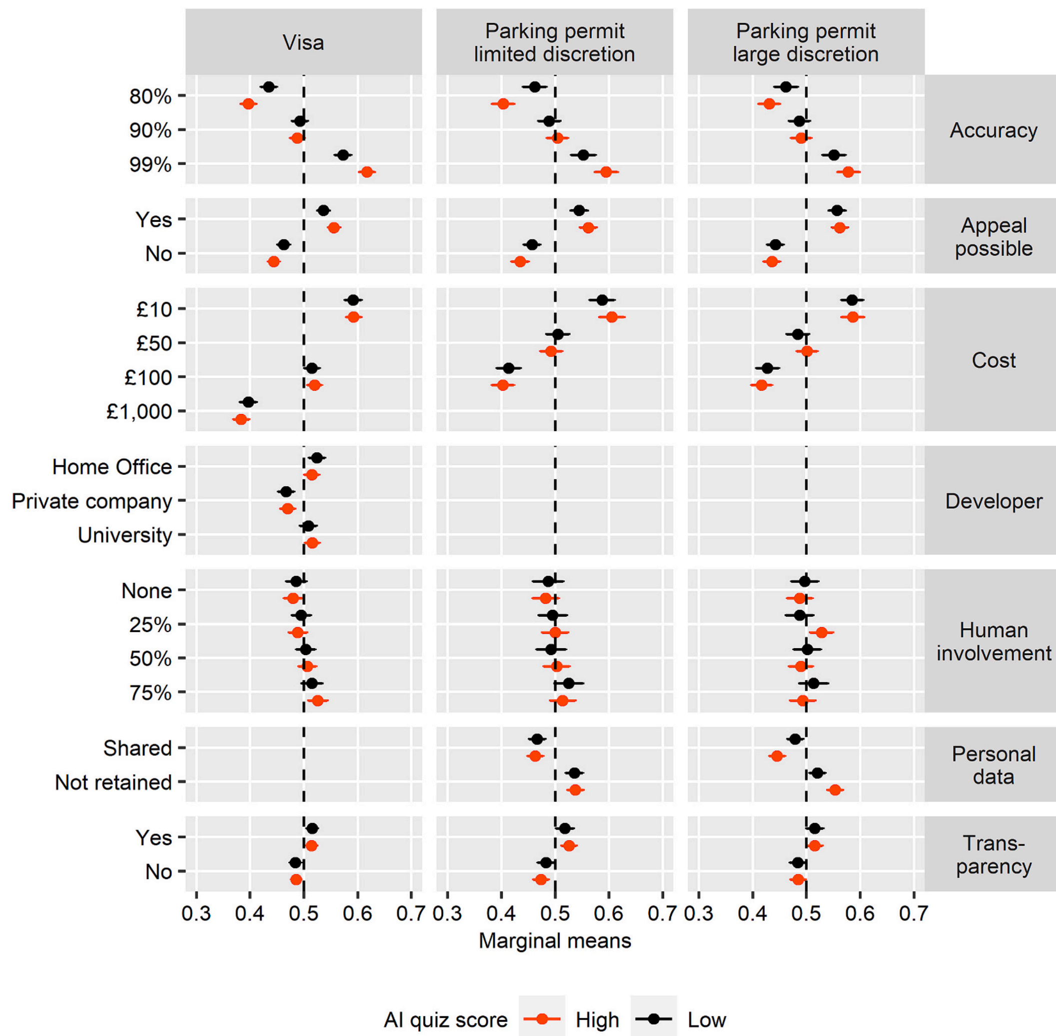

Fig. 6. Acceptance of system attributes conditional on knowledge about AI, using Leeper, Hobolt and Tilley's method (2020).

collaborative forms of decision-making, differently.

6. Conclusions

The findings of this study contribute to existing debates in three main ways: 1) how human versus AI involvement in public service provision shapes its acceptance by citizens; 2) how technology can shape the relationship between citizens and states; and 3) the key mechanisms, both in terms of the experiences of individuals and features of AI, that underpin its acceptance.

6.1. Human vs. AI involvement in public service provision

First, our research provides a framework for understanding citizens' perceptions of AI in public services, especially in administrative bureaucracy where discretion plays a role. We found that, in most scenarios, respondents did prefer processes with more human involvement, although these effects were relatively small compared to accuracy and cost considerations. Yet in specific contexts, such as local government parking permits based on demonstrable need, respondents showed a tendency to cap human involvement, favouring the algorithm. The nuances of public trust in different sectors of administration, from benefits allocation to parental support, may be key determinants here. We also suggest that future research explores this issue to understand if respondents with certain demographics, potentially related to presence or absence "demonstrable need", are more likely to prefer the algorithm over human discretion or the other way around.

6.2. Technology's impact on citizen-state interaction

Second, our study contributes to the existing literature about citizen-state interactions which is extensive but has so far not addressed AI in sufficient depth (for an overview see Jakobsen, James, Moynihan, & Nabatchi, 2019). This is a crucial insight, especially because the role of technology in this interaction is changing rapidly (Lindgren, Madsen, Hofmann, & Melin, 2019). One area of concern is data privacy. Our results suggest resistance to the accumulation and sharing of citizens' data—but we also show, in the context of other system-level characteristics, that accuracy seems to be more influential than data privacy.

6.3. Mechanisms underpinning acceptance

Third, our results contribute to the theoretical frameworks of technology acceptance in digital government and AI. By testing empirically vital mechanisms and examining the intricate relationships between individual characteristics, including literacy about AI, and features of the AI systems, we have highlighted new variables in this domain. However, we only tested two mechanisms against a controlled set of AI choices, which might not capture the full range of possible reactions. We suggest that a future observational study using survey questions could include consideration of more complex underlying mechanisms and interactions. In addition, future research could assess more individual level characteristics including age, gender, personality traits and the extent to which the use of the technology may help or hinder the state to achieve goals that are particularly important to individuals.

6.4. Limitations and assumptions

Our study's focus on the government's role in granting permits limits its applicability to public service contexts beyond this domain. In cases such as education, social care, or police interventions, the balance between humans and machine involvement might be different. Nevertheless, many routine interactions with the government involve permit applications similar to the kinds we examined such that the findings are of broad relevance. We also note that our experimental setup allow us to produce findings using "complete" information about the AI systems in a tabulated format. In real-world scenarios, citizens may neither have access to such comprehensive information nor actively seek it out. In particular, we suggest that further studies should investigate not only citizens' perceptions but also the effects of varying official communications about AI systems to citizens.

CRediT author statement

L Horvath: Conceptualization, Methodology, Software, Formal analysis, Data curation, Writing.
O James: Conceptualization, Methodology, Writing.
S Banducci: Conceptualization, Methodology, Writing, Resources, Supervision, Funding acquisition.
A Beduschi: Conceptualization, Writing.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A. Demographics, attention checks, and AI quiz

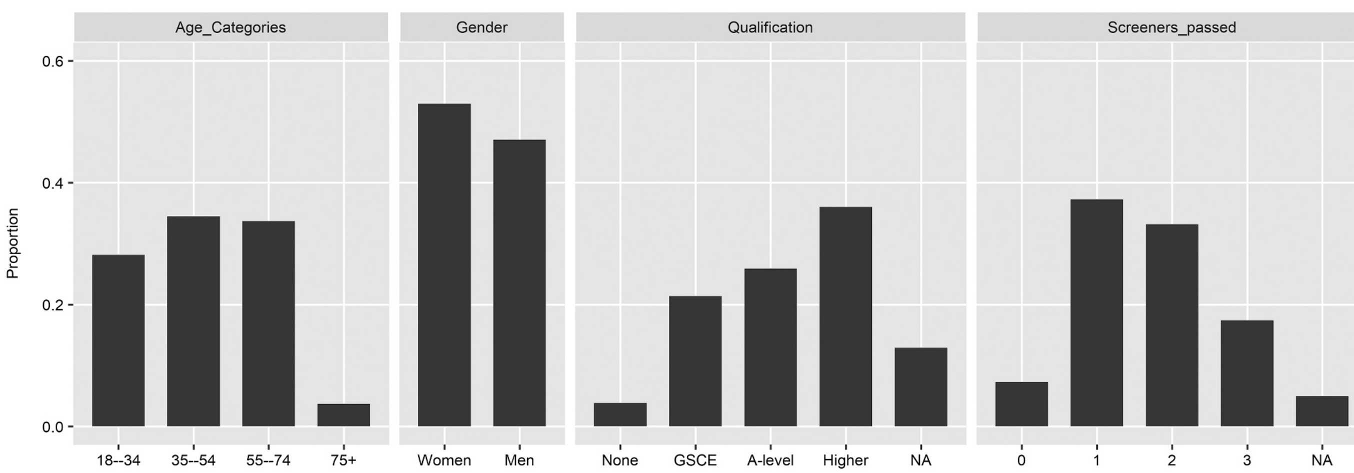

Fig. A1. Respondent demographics and attention checks.

Fig. A2. Time spent on survey by attention checks passed.

Table A1
Attention checks (adapted from Berinsky, Margolis, & Sances, 2014), for robustness checks using these measures and respondent N passing different data quality thresholds, see Appendix C.

Item ID	Survey question	Response options
Screener 1	How much do you agree or disagree with the following statements? (1 to 5 + DK)	Feminists want women to have equal power to men (1) For most women, equality means seeking special favors, such as hiring policies that favour them over men. (2) Most women interpret innocent remarks or acts as being sexist. (6) **Please click the "neither agree nor disagree"response (9)**
Screener 2	How much do you agree or disagree with the following statements? (1 to 5 + DK)	Women should be cherished and protected by men. (3) Many women have a quality of purity that few men possess. (7) A good woman should be set on a pedestal by her man. (8) **Two is greater than one. (9)**
Screener 3	When a big news story breaks people often go online to get up-to-the-minute details on what is going on. We want to know which websites people trust to get this information. We also want to know if people are paying attention to the question. **Please ignore the question and select bbc.co.uk and Google News as your two answers.**	The Telegraph (1) bbc.co.uk (2) Google News (3) CNN.com (4) Guardian (5) Sky News (6) Another site, namely (8)

Table A2
AI quiz questions (adapted from University of Helsinki & Rekator, 2021).

Which of the following are examples of artificial intelligence?	Response options (correct in bold)
Spreadsheet that calculates sums and other pre-defined functions on given data (1)	Yes, **No**, DK
Predicting the stock market by fitting a curve to past data about stock prices (2)	**Yes**, No, DK
A GPS navigation system for finding the fastest route (3)	**Yes**, No, DK
A music recommendation system such as Spotify that suggests music based on the user's listening behaviour (4)	**Yes**, No, DK
Big data storage solutions that can store huge amounts of data (such as images or video) and stream them to many users at the same time (5)	Yes, **No**, DK
Photo editing features such as brightness and contrast in applications such as Photoshop (6)	Yes, **No**, DK

Appendix B. Model results

Table B1
Approval of systems (AMCE, predicted increase in the population probability that the system with given attribute level is chosen over the reference category, see Hainmueller et al., 2014, pp. 10–12).

Attribute	Reference	Attribute level	Visa	Parking limited discretion	Parking large discretion
Accuracy	80%	90%	0.08*** (0.01)	0.07*** (0.01)	0.04*** (0.01)
		99%	0.18*** (0.01)	0.15*** (0.01)	0.11*** (0.01)
Developer	Gov't	Private company	−0.05*** (0.01)		
		University researchers	−0.01 (0.01)		
Cost	£10	£50		−0.10*** (0.01)	−0.09*** (0.01)
		£100	−0.08*** (0.01)	−0.18*** (0.01)	−0.18*** (0.01)
		£1000	−0.20*** (0.01)		
Transparency	Yes	No	−0.03*** (0.01)	−0.05*** (0.01)	−0.03** (0.01)
Human involvement	None	25%	0.01 (0.01)	0.01 (0.01)	0.04* (0.01)
		50%	0.03* (0.01)	0.01 (0.01)	0.01 (0.01)
		75%	0.04*** (0.01)	0.03* (0.02)	0.01 (0.01)
Appeal possible?	Yes	No	−0.09*** (0.01)	−0.11*** (0.01)	−0.12*** (0.01)
Personal data shared	Yes	No		0.07*** (0.01)	0.08*** (0.01)
N trials			5	5	5
N respondents			1958	925	996

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table B2
Fairness of systems (AMCE, predicted increase in the population probability that the system with given attribute level is perceived fairer over the reference category, see Hainmueller, 2013, pp. 10–12).

Attribute	Reference	Attribute level	Visa	Parking limited discretion	Parking large discretion
Accuracy	80%	90%	0.05*** (0.01)	0.02* (0.01)	0.03*** (0.01)
		99%	0.11*** (0.01)	0.07*** (0.01)	0.07*** (0.01)
Developer	Gov't	Private company	−0.03*** (0.01)		
		University researchers	0.00 (0.01)		
Cost	£10	£50		−0.03** (0.01)	0.00 (0.01)
		£100	−0.02*** (0.01)	−0.06*** (0.01)	−0.06*** (0.01)
		£1000	−0.06*** (0.01)		
Transparency	Yes	No	−0.03*** (0.01)	−0.04*** (0.01)	−0.01 (0.01)
Human involvement	None	25%	−0.01 (0.01)	−0.01 (0.01)	0.00 (0.01)
		50%	0.01 (0.01)	−0.01 (0.01)	0.01 (0.01)
		75%	0.02* (0.01)	0.00 (0.01)	0.02 (0.01)
Appeal possible?	Yes	No	−0.06*** (0.01)	−0.06*** (0.01)	−0.07*** (0.01)
Personal data shared	Yes	No		0.04*** (0.01)	0.03*** (0.01)
N trials			5	5	5
N respondents			1958	925	996

* p < 0.1, ** p < 0.05, *** p < 0.01

Table B3
Component interactions (adding on to main effects as seen in Table B1).

Attribute	Reference	Attribute level	Visa	Parking limited discretion	Parking large discretion
Accuracy x Human inv.	80%	90% at 25% involvement	0.05* (0.03)	0.00 (0.03)	0.05 (0.03)
		99% at 25% involvement	0.04 (0.02)	−0.02 (0.03)	0.06* (0.03)
		90% at 50% involvement	−0.01 (0.02)	−0.03 (0.03)	0.01 (0.03)
		99% at 50% involvement	0.02 (0.02)	−0.03 (0.04)	0.05 (0.03)
		90% at 75% involvement	0.01 (0.02)	−0.01 (0.03)	0.05 (0.03)
		99% at 75% involvement	−0.02 (0.02)	0.00(0.03)	0.04 (0.03)
Developer x Human inv.	Gov't	Private company x 25% involvement	0.02 (0.02)		
		University researchers x 25% involvement	0.01 (0.02)		
		Private company x 50% involvement	0.00 (0.02)		
		University researchers x 50% involvement	0.02 (0.02)		
		Private company x 75% involvement	0.00 (0.02)		
		University researchers x 75% involvement	0.02 (0.02)		
Cost x Human inv.	£10	£50 × 25% involvement		0.01(0.04)	−0.01 (0.03)
		£100 × 25% involvement	0.00 (0.02)	0.04(0.03)	0.00 (0.03)
		£1000 × 25% involvement	0.02 (0.02)		
		£50 × 50% involvement		−0.02 (0.04)	0.05 (0.03)
		£100 × 50% involvement	0.03 (0.02)	0.03 (0.03)	0.03 (0.03)
		£1000 × 50% involvement	0.02 (0.02)		
		£50 × 75% involvement		0.06* (0.04)	0.01 (0.03)
		£100 × 75% involvement	−0.02 (0.02)	0.07** (0.03)	0.02 (0.03)
		£1000 × 75% involvement	0.00 (0.02)		
Transparency x Human inv.	Yes	No x 25% involvement	0.00 (0.02)	−0.03 (0.03)	−0.03 (0.03)
		No x 50% involvement	−0.01 (0.02)	−0.01 (0.03)	−0.02 (0.03)
		No x 75% involvement	−0.01 (0.02)	−0.04 (0.03)	0.00 (0.03)
Appeal x Human inv.	Yes	No x 25% involvement	0.03 (0.02)	0.02 (0.03)	−0.02 (0.03)
		No x 50% involvement	0.04* (0.02)	0.01 (0.03)	0.00 (0.03)
		No x 75% involvement	0.03 (0.02)	0.00 (0.03)	0.02 (0.03)
Personal data shared	Yes	No x 25% involvement		0.02 (0.03)	0.02 (0.03)
		No x 50% involvement		0.00 (0.03)	−0.03 (0.03)
		No x 75% involvement		−0.01 (0.03)	0.00 (0.03)
N trials			5	5	5
N respondents			1958	925	996

* p < 0.1, ** p < 0.05, *** p < 0.01.

Table B4.1
F-test for model comparison: Simple model (preference by human involvement, four attribute levels) against moderated model (preference by human involvement * technology adoption).

Task	Model	Residual Dev	Degrees of freedom	F	p
Visa	Simple model	4849.5			
	Moderated model	4847.1	4	2.40	0.04
Parking *limited discretion*	Simple model	2302.1			
	Moderated model	2300.8	4	1.21	0.30
Parking *large discretion*	Simple model	2482.0			
	Moderated model	2481.0	4	1.72	0.14

Table B4.2
Marginal means of system approval by human involvement by levels of technology adoption. Interpretable as proportion of binary choices in favour of given attribute level.

Technology adoption	Human involvement	Mean approval: Visa	Mean approval: Parking *limited discretion*	Mean approval: Parking *large discretion*
Low	None	0.47	0.50	0.47
	25%	0.48	0.49	0.51
	50%	0.51	0.49	0.50
	75%	0.54	0.53	0.52
High	None	0.60	0.48	0.51
	25%	0.49	0.50	0.51
	50%	0.50	0.51	0.49
	75%	0.50	0.51	0.49

Table B5.1
F-test for model comparison: Simple model (preference by human involvement, four attribute levels) against moderated model (preference by human involvement * knowledge about AI).

Task	Model	Residual Dev	Degrees of freedom	F	p
Visa	Simple model	4849.5			
	Moderated model	4849.2	4	0.25	0.90
Parking *limited discretion*	Simple model	2302.1			
	Moderated model	2301.9	4	0.16	0.95
Parking *large discretion*	Simple model	2482.0			
	Moderated model	2480.6	4	1.48	0.21

Table B5.2
Marginal means of system approval by human involvement by levels of knowledge about AI. Interpretable as proportion of binary choices in favour of given attribute level.

Technology adoption	Human involvement	Mean approval: Visa	Mean approval: Parking *limited discretion*	Mean approval: Parking *large discretion*
Low	None	0.49	0.49	0.50
	25%	0.50	0.50	0.49
	50%	0.50	0.49	0.50
	75%	0.51	0.53	0.51
High	None	0.48	0.50	0.53
	25%	0.49	0.50	0.49
	50%	0.51	0.51	0.49
	75%	0.53	0.46	0.45

Appendix C. Robustness checks

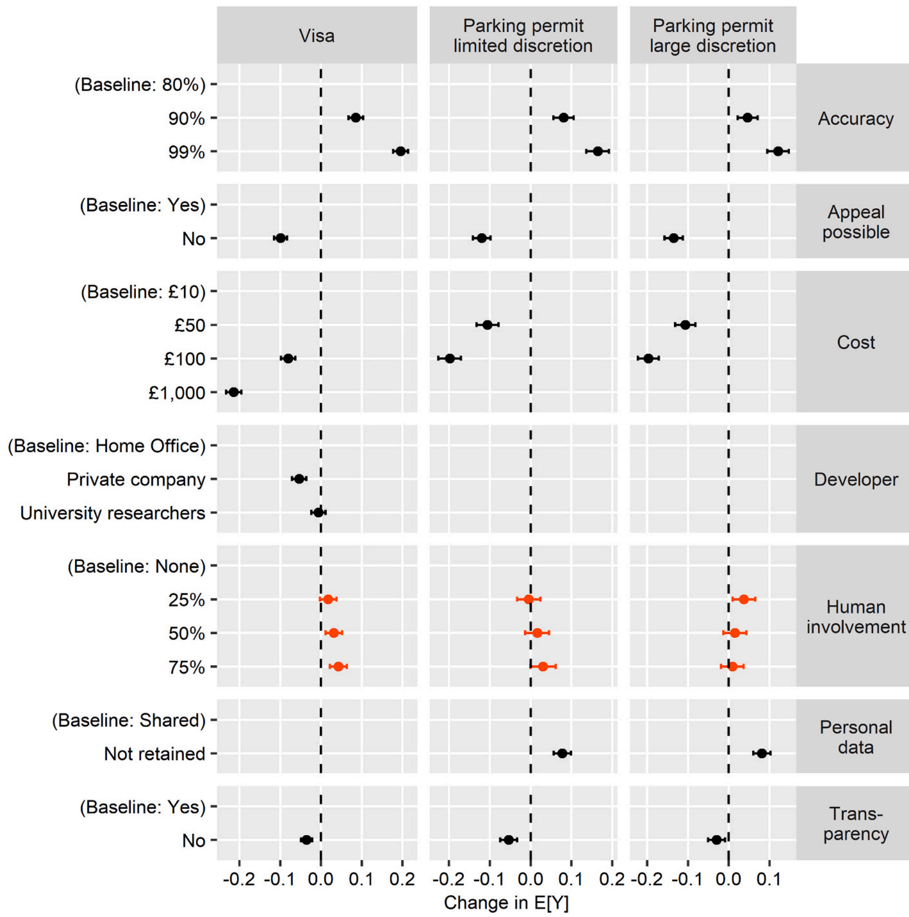

Fig. C1. AMCEs with low bar on data quality (passed one screener, respondent $N = 1987$).

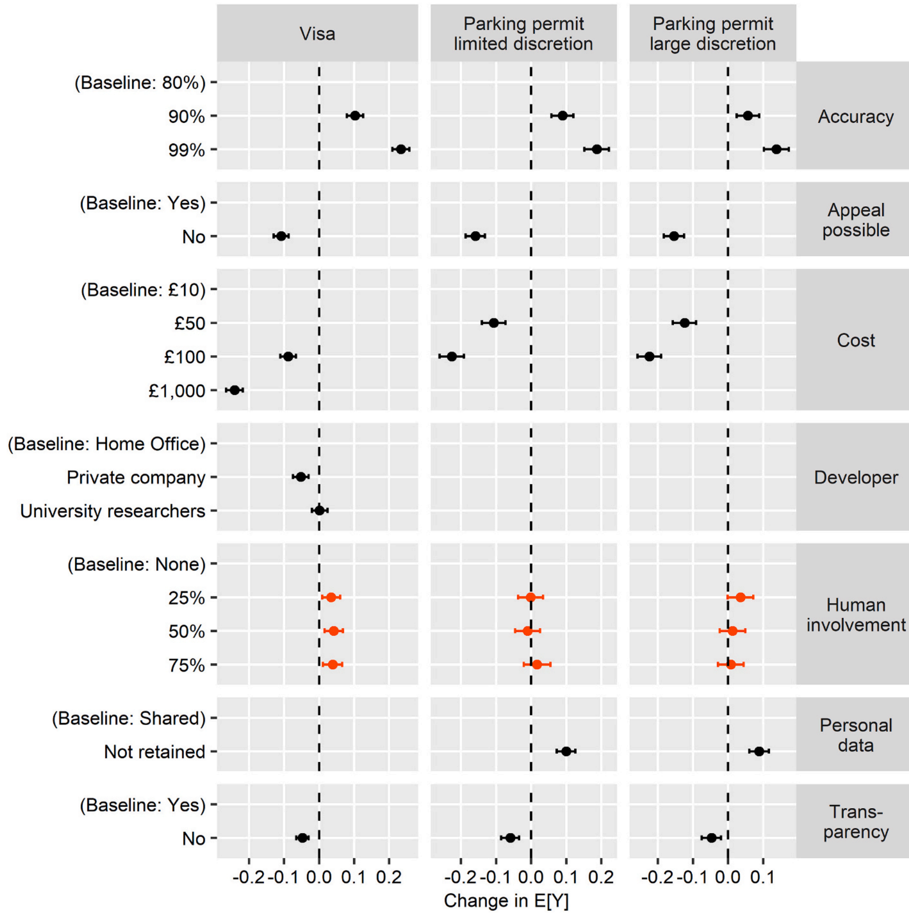

Fig. C2. AMCEs with mid bar on data quality (passed two screeners, respondent $N = 1189$).

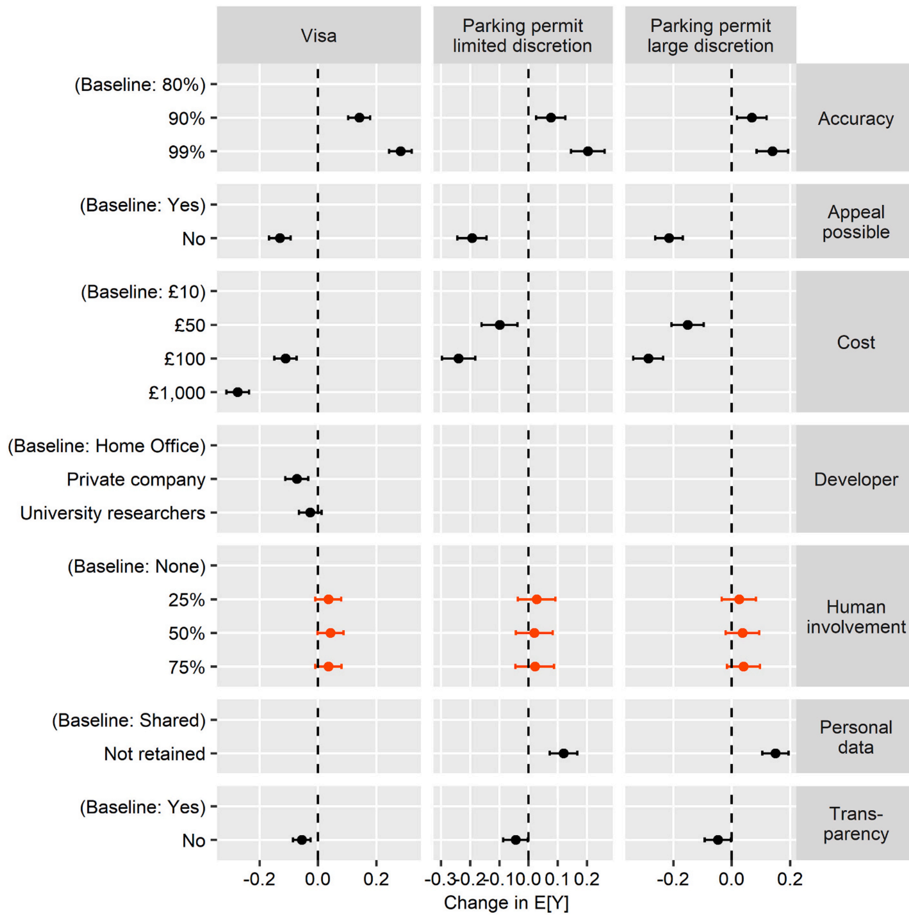

Fig. C3. AMCEs with high bar on data quality (passed three screeners, respondent $N = 479$).

Appendix D. Supplementary data

Supplementary data to this article can be found online at https://doi.org/10.1016/j.giq.2023.101876.

References

Anthopoulos, L., Reddick, C. G., Giannakidou, I., & Mavridis, N. (2016). Why e-government projects fail? An analysis of the Healthcare.gov website. *Government Information Quarterly, 33*(1), 161–173.

Aoki, N. (2020). An experimental study of public trust in AI chatbots in the public sector. *Government Information Quarterly, 37*(4), Article 101490.

Berinsky, A. J., Margolis, M. F., & Sances, M. W. (2014). Separating the shirkers from the workers? Making sure respondents pay attention on self-administered surveys. *American Journal of Political Science, 58*(3), 739–753.

Brauneis, R., & Goodman, E. P. (2018). Algorithmic transparency for the Smart City. *Yale Journal of Law and Technology, 20*, 103–176.

Buffat, A. (2015). Street-level bureaucracy and e-government. *Public Management Review, 17*(1), 149–161.

Busch, P. A., & Henriksen, H. Z. (2018). Digital discretion: A systematic literature review of ICT and street-level discretion. *Information Polity, 23*(1), 3–28.

Chohan, S. R., & Hu, G. (2020). Success factors influencing citizens' adoption of IoT service orchestration for public value creation in smart government. *IEEE Access, 8*, 208427–208448.

Criado, J. I., Valero, J., & Villodre, J. (2020). Algorithmic transparency and bureaucratic discretion: The case of SALER early warning system. *Information Polity, 25*(4), 449–470.

Dietvorst, B., Simmons, J. P., & Massey, C. (2015). Algorithm aversion: People erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General, 144*(1), 114–126. https://doi.org/10.1037/xge0000033

Engin, Z., & Treleaven, P. (2019). Algorithmic government: Automating public services and supporting civil servants in using data science technologies. *The Computer Journal, 62*(3), 448–460.

European Commission. (2020). *White paper on artificial intelligence - a European approach to excellence and trust. 2020.* COM, 65 final.

Gil-Garcia, J. R., & Flores-Zúñiga, M.Á. (2020). Towards a comprehensive understanding of digital government success: Integrating implementation and adoption factors. *Government Information Quarterly, 37*(4), Article 101518.

Guppy, N., Boud, D., Heap, T., Verpoorten, D., Matzat, U., Tai, J., & Bartolic, S. (2022). Teaching and learning under COVID-19 public health edicts: The role of household lockdowns and prior technology usage. *Higher Education, 84*(3), 487–504.

Hainmueller, J., Hopkins, D. J., & Yamamoto, T. (2014). Causal inference in conjoint analysis: Understanding multidimensional choices via stated preference experiments. *Political Analysis, 22*(1), 1–30.

Harrison, T. M., & Luna-Reyes, L. F. (2020). Cultivating trustworthy artificial intelligence in digital government. *Social Science Computer Review, 0894439320980122*.

Horvath, L., Banducci, S., Blamire, J., Degnen, C., James, O., Jones, A., & Tyler, K. (2022). Adoption and continued use of mobile contact tracing technology: Multilevel explanations from a three-wave panel survey and linked data. *BMJ Open, 12*(1), Article e053327.

Jakobsen, M., James, O., Moynihan, D., & Nabatchi, T. (2019). JPART virtual issue on citizen-state interactions in public administration research. *Journal of Public Administration Research and Theory, 29*(4), e8–e15.

Janssen, M., & Kuk, G. (2016). The challenges and limits of big data algorithms in technocratic governance. *Government Information Quarterly, 33*(3), 371–377.

Johnson, D. G., & Verdicchio, M. (2017). AI anxiety. *Journal of the Association for Information Science and Technology, 68*(9), 2267–2270.

Jones, M. L. (2017). The right to a human in the loop: Political constructions of computer automation and personhood. *Social Studies of Science, 47*(2), 216–239.

Kitchin, R. (2017). Thinking critically about and researching algorithms. *Information, Communication & Society, 20*(1), 14–29.

Lee, M. K. (2018). Understanding perception of algorithmic decisions: Fairness, trust, and emotion in response to algorithmic management. *Big Data & Society, 5*(1). https://doi.org/10.1177/2053951718756684

Leeper, T. J., Hobolt, S. B., & Tilley, J. (2020). Measuring subgroup preferences in conjoint experiments. *Political Analysis, 28*(2), 207–221.

Levy, F. (2018). Computers and populism: Artificial intelligence, jobs, and politics in the near term. *Oxford Review of Economic Policy, 34*(3), 393–417.
Lindgren, I., Madsen, C.Ø., Hofmann, S., & Melin, U. (2019). Close encounters of the digital kind: A research agenda for the digitalization of public services. *Government Information Quarterly, 36*(3), 427–436.
Lipsky, M. (1980). *Street-level bureaucracy: Dilemmas of the individual in public service*. Russell Sage Foundation.
Long, D., & Magerko, B. (2020). What is AI literacy? Competencies and design considerations. In *Proceedings of the 2020 CHI conference on human factors in computing systems* (pp. 10–16).
Margetts, H. (1999). *Information Technology in Government: Britain and America*. Routledge.
Meijer, A., & Grimmelikhuijsen, S. (2020). Responsible and accountable algorithmization: How to generate citizen trust in governmental usage of algorithms. In *The algorithmic society* (pp. 53–66). Routledge.
Meijer, A., & Wessels, M. (2019). Predictive policing: Review of benefits and drawbacks. *International Journal of Public Administration, 42*(12), 1031–1039.
Rahwan, I. (2018). Society-in-the-loop: Programming the algorithmic social contract. *Ethics and Information Technology, 20*(1), 5–14.
Smith, A. (2018). *Public attitudes toward computer algorithms*. Pew Research Center. https://www.pewresearch.org/internet/2018/11/16/public-attitudes-toward-computer-algorithms/.
Starke, C., Baleis, J., Keller, B., & Marcinkowski, F. (2022). Fairness perceptions of algorithmic decision-making: A systematic review of the empirical literature. *Big Data & Society, 9*(2), 20539517221115189.
Taddeo, M., & Floridi, L. (2018). How AI can be a force for good. *Science., 361*(6404), 751–752.
Taylor, M. Z. (2016). *The politics of innovation: Why some countries are better than others at science and technology*. Oxford University Press.
Tyler, T. R. (2006). *Why people obey the law*. Princeton University Press: Princeton.
University of Helsinki, & Reaktor. (2021). A free online introduction to artificial intelligence for non-experts. Global. Retrieved 24 May 2021, from https://www.elementsofai.com/.
Veale, M., & Brass, I. (2019). Administration by algorithm? Public management meets public sector machine learning. In K. Yeung, & M. Lodge (Eds.), *Algorithmic regulation*. Oxford University Press.
Veale, M., Van Kleek, M., & Binns, R. (2018). Fairness and accountability design needs for algorithmic support in high-stakes public sector decision-making. In *Proceedings of the 2018 CHI conference on human factors in computing systems* (pp. 1–14).
Weinberg, J. D., Freese, J., & McElhattan, D. (2014). Comparing data characteristics and results of an online factorial survey between a population-based and a crowdsource-recruited sample. *Sociological Science, 1*, 292–310.
Young, M. M., Bullock, J. B., & Lecy, J. D. (2019). Artificial discretion as a tool of governance: A framework for understanding the impact of artificial intelligence on public administration. *Perspectives on Public Management and Governance, 2*(4), 301–313.
Yu, T. K., Lin, M. L., & Liao, Y. K. (2017). Understanding factors influencing information communication technology adoption behavior: The moderators of information literacy and digital skills. *Computers in Human Behavior, 71*, 196–208.

Laszlo Horvath is a Lecturer in Political Communication and Director of the MSc in Political Communication programme at Birkbeck College, University of London. His work focuses on digital technologies for governance and politics, and how they affect citizens' political behaviour.

Oliver James is a Professor of Political Science, focusing on topics in public policy including: citizen-state relations (especially participation in policy-making and implementation), evidences and evaluation in the policy process, public regulation, and digital technology and public services.

Susan Banducci is a Professor of Politics at the University of Exeter, a co-director of the Exeter Q-Step Centre and a Turing Fellow. Her research addresses questions of inequalities in political behaviour, bringing together individual and institutional explanations.

Ana Beduschi is an Associate Professor of Law at the University of Exeter, Law School. Her research and teaching focuses on international human rights law, technology, digital law, data protection and privacy, and international migration and refugee law.