

Keeping it Human: A focus group study of public attitudes towards AI in banking

Abstract. While there is substantial interest in ethical practice relating to Artificial Intelligence (AI), to date there has been limited consideration of what this means in the banking sector. This study aimed to address this gap in the literature through a qualitative study of public attitudes and perceptions of current and potential future uses of AI in banking. A series of focus groups were conducted with diverse members of the public. Focus group participants were largely positive about the role of AI in speeding up financial processes and increasing efficiency. Yet, they also expressed a number of concerns around potential negative impacts on society and consistently emphasized the importance of human judgement or oversight. The findings suggest a potential cognitive dissonance where people use new services due to perceived convenience or immediate benefits, while disliking or distrusting those services or holding concerns about their impacts on society. The findings illustrate that participants' concerns did not typically relate to private or individual interests but more often to wider ethical and social concerns. The focus groups demonstrated the value of qualitative, deliberative methods to explore the nuances of public responses and highlighted the importance of taking account of conditions for public acceptability - rather than just customer uptake - in order to develop ethical practice and establish a social licence for uses of AI in banking.

Keywords: FinTech, Artificial Intelligence, Public Acceptability.

1 Introduction

Innovation in artificial intelligence (AI) is underpinning the development of new products and services across a wide range of industries. This is particularly true in banking where digital innovation and new data practices are fuelling interest and investment in FinTech (financial technology) and changing the banking sector as a whole (1).

AI in banking is used for a number of purposes including developing automated chatbots for customer services, efficient processes for detecting fraud and money laundering and improving automated processes that utilise large volumes of data (e.g. client risk profiling or credit scoring [2]).

While, there is substantial interest in ethical practice relating to AI and growing consideration of what it means for AI to be trustworthy (3,4), as well as some evidence of growing interest in what this means in the financial sector¹. The existing literature typically does not engage with public – rather than professional stakeholder – views and experiences. Studies which have examined public attitudes or responses have typically focused on customer uptake of FinTech products (e.g 5-7). In doing so they have tended to focus on customers’ motivations for using FinTech services, and largely neglected non-customers’ reasons for not using FinTech services, or the reasons why some FinTech offerings have been unsuccessful (8). Where studies have considered trust relating to FinTech (e.g. 5, 6) they have typically focussed on customer trust and willingness to use FinTech services and largely overlooked considerations of what it means for a FinTech to be trustworthy.

While there is a rapidly growing literature relating to social and ethical dimensions of AI (e.g. 9,10) across both industry and academia. There is a lack of research examining the extent to which FinTech practices align with public values and interests, in particular there is very limited evidence relating to public acceptability of uses of AI in banking (1). Therefore, this study aims to address this important gap in the literature through a qualitative study of public attitudes and perceptions of digital innovation in banking with a particular focus on current and potential future uses of AI in banking.

The research was undertaken through a series of focus groups which engaged with diverse members of the public. Focus group participants are referred to as members of the public rather than customers or consumers since the study aimed to engage broad perspectives rather than solely those of current or potential FinTech customers. Indeed, the aim was not to identify customer experiences or reasons for using/not using FinTech services, but rather to understand public attitudes, concerns and interests relating to uses, and potential uses of AI in banking. This is important since finance is an area that affects most – if not all – members of society and therefore digital innovation in this field is likely to have broad, and diverse, impacts. Such impacts might be positive (e.g. opening up financial services to unbanked or underbanked populations [11]) or negative (e.g. creating new opaque systems through which access to finance is determined or increasingly necessitate citizens’ participation in the Big Data society [1]). Therefore, ethical practice requires considering broader social acceptability – not just customer uptake - of innovation.

2 Methods

A series of five focus groups were held between September and November 2019. Participants were recruited through pre-established groups including students, meetup

¹ For example, through the work of organisations such as the Finance Innovation Lab: financeinnovationlab.org or industry oriented publications on developing ethical AI.

groups for senior citizens and young professionals, and via community centres. Focus group participants were sampled through purposive sampling focused on maximising diversity across the focus groups to access a broad range of viewpoints and perspectives. The aim was to have a diverse, rather than statistically representative, sample (12). It was important that individuals within each of the groups shared common traits or interests as this meant that they felt comfortable and able to discuss the issues freely (12).

A semi-structured approach was taken. A topic guide was developed to ensure a level of consistency between the focus groups. This was very flexible and allowed participants to raise issues and/or concerns which they considered to be relevant. The semi-structured design also meant that topics of discussion did not always arise in a pre-determined order, and that the focus groups were able to explore unanticipated areas of interest. As is recognised to be an advantage of focus group research, this approach allowed for a responsive, conversational style resulting in open and frank discussions and enabled individuals to engage with aspects of the topic which were previously unfamiliar to them (12).

The focus groups began by asking participants about the ways in which they currently access banking services and their experiences using digital banking and mobile banking. Following this the focus groups then focused on digital innovation in banking. A series of vignettes were introduced to engage the participants in discussions of the ways in which digital innovation is used in banking and how this may change banking services in the future. Through these discussions participants were encouraged to reflect on how they felt about these areas of innovation, whether they would be interested in using them and what potential benefits or risks they anticipated relating to these. Following the focused discussions on each of the vignettes the focus groups then led into broader discussions of attitudes, preferences and concerns around future directions of digital innovation in banking and, in particular, the role of AI in banking.

The focus groups typically lasted around one hour. They were audio-recorded with consent from participants. The recordings were professionally transcribed. The transcripts were reviewed and corrected by MA and any potentially identifying information was removed (e.g. any names or details about workplaces) before the transcripts were analysed.

The transcripts were coded in NVivo. This coding process identified key themes in the focus group discussions and highlighted emerging areas of interest and/or concern. Following an inductive approach the coding focussed on bringing out high level themes which emerged from the data, to structure a narrative account of the discussions. The themes were mapped to illustrate their interconnections and the ways in which different topics discussed related to one another. The mapping exercise then informed the process of writing up the findings by indicating the connections between different areas of interest.

2.1 Sample

Five focus groups were conducted these are referred to in the following discussion by reference to the common characteristics of the participants in each group:

1. Seniors (members of a meetup group for senior citizens)
2. Rural (a focus group held in a rural town 20 miles outside Newcastle)
3. Young Professionals (recruited via a meetup group for young professionals)
4. International Students (recruited via student forums at Newcastle University)
5. Postgraduate Students (recruited via student forums at Newcastle University)

The focus groups all took place in Newcastle, except for one which was held in a rural town 20 miles outside of Newcastle. It was important to hold a focus group in a more rural location as it was anticipated that rural residents may be likely to have different experiences relating to accessibility of banking services and digital connectivity.

A total of 23 participants took part in the focus groups. While the majority of participants were under the age of 40, the focus groups engaged a diverse range of age groups (the youngest participants being in their early 20s and the oldest over 75). Participants came from a range of professional backgrounds, the largest group were students (39 per cent), 13 per cent of participants were employed full-time and equally 13 per cent were employed part-time, 13 per cent were jobseekers, 9 per cent of participants were self-employed and equally 9 per cent were retired. Sixty per cent of the sample identified as white, 21 per cent identified as Asian or Asian British, 4 per cent identified as Black or Black British, 4 per cent as mixed and 8 per cent as other.

While the focus groups were conducted with a range of groups which were non-gender specific, in all instances female members were more likely to volunteer to participate in the study. This resulted in the sample being almost 75 per cent female (17 women, 6 men). This is a limitation of the present study which will be addressed in future research.

In the following discussion individual participants are identified by the name of the group they participated in followed by either F or M (to indicate Female or Male) and a number which remains the same for each participant throughout.

3 Findings

While the focus groups discussed a variety of forms of digital innovation in banking, this paper reports findings relating to responses to AI in banking. Across the focus groups participants discussed current and potential roles of AI in banking. In particular this was explored in relation to: virtual money coaches, chatbots and algorithmic decision-making.

3.1 Virtual Money Coaches

Virtual money coaches was introduced as a vignette to spark discussions on potential uses of AI in banking. Participants were given some brief information about current applications (apps) which can track customers' spending habits and provide advice on how to meet savings goals and/or which banking services would be most suitable (e.g. credit cards, loans, savings accounts). The information given noted that these apps may use AI to process customers' information and give advice and that AI can be used to

create personalised “virtual money coaches” which can track spending habits and offer advice.

Participants had mixed responses to this vignette. Some reported already using apps that had these functions, or potentially being interested to use them in the future. However, a common area of concern was the frequency of alerts that such apps send, and that they can become a nuisance.

Other participants were strongly opposed to using virtual money coaches. A key factor distinguishing responses of participants who were enthusiastic about virtual money coaches from those who were opposed to using them was whether they were perceived to give customers more or less control over their personal finances, for example in one discussion a participant described virtual money coaches as being useful for people who “want to have their lives run for them”:

F1: I think it's quite a good idea, because my stepdaughter- I always withdraw cash a couple of times a week £40 pounds here or there, but she just uses her card for everything and she's got no track of what she spends, so I think that for someone like her it would be good

F2: She doesn't even keep her statements does she?

F1: No she puts her statements straight on the fire, she's terrible! So I think something like that for her, people like her, would be really good

F2: If they want to have their lives run for them it's ideal

F1: But I like to keep track of my own money

F2: I do as well

(SENIORS)

Some participants felt that such tools relinquished personal control or decision-making:

“I think it's good but I always have the opinion you should do something yourself about it and not have others influence you. If it's a virtual money coach or your mum or whatever, you should have you own control and find your own ways of how to use your money because this is not really a way to learn how to deal with money or how to spend your money or how to save money. It's just literally someone in the back always coaching you so you're never really independent in whatever you do.” (INTERNATIONAL STUDENTS: M2)

Conversely, other participants felt that virtual money coaches gave individuals more control over their finances:

F2: I think it's a good idea more so than getting alerts on you're reaching your limit or whatever, but getting coaching on right you've been to Costa coffee twice already this week, did you know that if you swapped that for something else you could save this much

F1: Yeah, if it was kind of helpful like that. Maybe not just an alert like you've spent this much. If you're spending this much eating out this week, maybe just make a sandwich!

(RURAL)

A number of participants were concerned that data collected through virtual money coaches could be used for marketing or commercial purposes. Therefore, they would be more inclined to use the service when it was provided by a trusted organisation or a non-profit third party.

"I'd feel a lot happier if the virtual money coach is set by the government or by Watchdog or somebody that's fully impartial like Citizen's Advice Bureau type thing. If it was just giving sound financial advice, that's brilliant. If it's a start-up and it's minor information that you're giving to it selling that to third parties or if it's sold through a bank and they have a vested interest and say, "You'll save £100 if you buy this financial product," then I think that's a lot murkier but then I also think that if it was done by the government, the artificial intelligence might not develop quite as quickly"

(POSTGRADUATE STUDENTS: M2)

Another area of concern relating to virtual money coaches was the extent to which an automated, AI-powered process could take account of the personal, emotional nature of finances:

"the money coach doesn't really know your own situation. It doesn't know I have 500 Euros cash, someone else is going to pay me something, someone else has to pay me something back or these kind of things they are not taking into consideration your own situation for when I'm like I have, I don't know, 30 Euros in cash so I'm paying 30 Euros from my savings, they're like, "You can't touch that kind of thing." I'm like, "I have cash," and those kinds of things. That's what I think." (INTERNATIONAL STUDENTS: M2)

M1: I think it's good to have the human touch behind it, because sometimes if you're going through some sort of vulnerability or problem and you're trying to discuss with this money coach, it doesn't understand that vulnerability

F1: yeah, I would like to speak to a person if I was having money problems

F2: Yeah I would definitely

M1: Yeah

MA: Is that because a person has empathy?

F2: You're not talking to a machine

M1: It could make things worse talking to a machine when you've got problems

(SENIORS)

This emphasis on human involvement to take account of personal and emotive factors is a theme which was also salient in relation to chatbots.

3.2 Chatbots

In discussing the role of technology in banking, participants frequently raised examples of experiences with chatbots in customer services. Chatbots were not introduced as a vignette in the focus groups but were frequently referred to by participants, suggesting this was a topic that participants had some familiarity with and were comfortable discussing. Typically participants discussed this in relation to experiences with automated telephone services and reflected on their experiences and frustrations with current automation in customer services. In particular, participants frequently reported frustrations relating to trying to get answers to questions that were not routine queries, resulting in slower processes with calls eventually being redirected to human handlers:

F1: but you could ask it something that it doesn't understand

F2: does not compute

F1: when you do telephone banking and you get the press 1, 2, 3 sort of thing sometimes they don't understand what you're asking. If it's anything that deviates from what they're programmed to do

F2: yeah, yeah I've had problems like that. I think I had to do that when I was making enquiries about my card possibly being scammed, I didn't answer with so many words

F1: and then they don't understand

F2: they don't understand

F1: They do put you through to somebody eventually

F2: am I going to get somebody real on this phone

F1: yeah, I want to speak someone real

(SENIORS)

"every time we need to call to the customer service centre to say, "Okay, my card is lost, please freeze." So it takes time because by the time you want to reach them, you go through, I don't know, ask you, "Okay, you have to press one, press two," so it takes a lot of time. Then you're worrying at that time so it's very frustrating." (INTERNATIONAL STUDENTS: F1)

These same frustrations were also reported in relation to participants' experiences with chatbots. The majority of participants reported disliking chatbots:

"One of the reasons I said I left M&S is because of their customer service and one of it is because it's a chat bot. It looks like you're going to be speaking with a human and it's not. You ask it a question and it just gives obviously this default standard answer which is really frustrating when you've got an actual question that needs answering. Then you have to ring up, doubling the time. That's not for me. I can't stand chat bots. I like to speak to a real person" (RURAL: F1)

F1: I always get frustrated by chatbots because every time they just ask you one question and then you click the answer and then they pop up other things. Then you go

through... at the end it doesn't really solve my problem so it's very frustrating I could save my time and talk to someone else.

M2: You sometimes also have very specific questions that can be answered, like they would analyse it as this is a general question and you're like, "Yes, it is a general question but it has a completely different context," which the machine wouldn't really understand. That's the problem I have a lot.

(INTERNATIONAL STUDENTS)

"I'm not a lover of chat bots and stuff because I think if you've got a query that's not, what's the word, not ordinary, it's complex in some way, they don't always have the right answer so you end up ringing up and speaking to somebody anyway"

(POSTGRADUATE STUDENTS: F2)

Many participants did not like to use chatbots as they felt it was important to speak to a human when discussing personal finances. This was in part because the data disclosed was sensitive and also because participants considered human judgement and empathy to be important when providing advice on financial matters.

M2: I totally agree because it's such sensitive data, it's such a personalised thing, I don't want to talk to a robot about my finances. I'd rather have a person who has experience with it and can basically react to my specific case because the chat bots are only there to take all these questions that can easily be answered from the service team, which is a good idea for the bank because it saves time but I would still...

F3: It saves them money for customer care.

(INTERNATIONAL STUDENTS)

M2: I would still prefer to talk to someone in person because you know who is sitting there, you know the person who is telling you this, who is behind it. If you just talk to the machine then it's like you have no point of relation to the machine or to the bank or anything. It's more trustworthy to just go to the person who is representing the bank as well.

M1: I think at the end of the day, even if you physically talk to a person, the person is only looking at the information on the screen. He is just interpreting that in a human friendly way. So at the end of the day it is machines and I think all of those systems are, in one way or the other, linked to artificial intelligence so I'm not sure. Even for example for the loan side or anything, even if you explain to him for example about the credit score, if the machine says no, if the system says no, I don't think he physically can do anything about it so he's just saying it maybe in a nicer way.

F4: I think for me myself, I prefer talking to a person because somehow a person has more emotion and can consider the consequences, he can tell you more things than the artificial intelligence can do. Some people just like to have this relationship where you just want to talk to that person face to face and then they can consider you and they can sympathise with you. They give you the best solution and all that. People can take advice and they will tend to follow what the other person is advising them.

M2: Because otherwise you're just a case, like the case opens, they analyse you, the case closes, okay next one. You're like, "Okay, can I have this?" "No," where else to go or what else to do or who to talk to or whatever. That's what I'm missing with this human connection.

(INTERNATIONAL STUDENTS)

However, there were a small number of participants who preferred speaking with chatbots compared to human customer service operatives. For those who preferred chatbots the primary reason was being able to get answers to questions quickly and easily:

M2: I find the online chat really useful. It's outsourced to people but you can get to meet to have a look over your history and any questions that you've got or if something isn't quite right, I can go into the actual bank itself.

F3: Yes. For me it's anything that I can avoid. If I can avoid going into the bank, I will. I think now, I can only speak for myself but I think I've got quite a lot of trust even if it's not a person behind, even if it's a bot, especially with machine learning and there's such advanced AI capabilities I'm okay speaking with a robot as long as I get my answer quickly. So any services that prevent me from going out to stand in a queue or to the bank is a bonus for me.

(POSTGRADUATE STUDENTS)

While most participants felt that the sensitive nature of personal finances required human consideration, these participants felt more comfortable disclosing sensitive – or potentially embarrassing – information to a machine:

“For me, if we get the security sorted and all of that sorted, I would probably use it just because I think I'd feel a bit more confident saying things which I don't want to say to a human being or person. For example if nowadays you go on to Google and you search anything, sometimes silly questions or anything which you wouldn't be able to go to someone and ask them, you just don't feel confident whereas with Google you know it's a machine. It will give you the right things you are looking for, probably better than a human being but even give you that confidence that no one is going to know about it but it gives you that more privacy I'd say.” (INTERNATIONAL STUDENTS:M1)

However, even participants who were more positive about using chatbots had concerns. In particular, some participants felt uneasy with the idea of “talking to a robot” or with AI imitating human emotions:

“It sounds a bit strange, like as if you're talking to a robot. So I don't know if it's that trust that it's not a human you're talking to, but if they're intelligent. I guess it would just be like getting an alert but they're speaking to you. I mean maybe because people do really like Alexa and stuff” (YOUNG PROFESSIONALS: F1)

The discussions around experiences and perceptions of chatbots revealed the importance of human involvement. Participants who did not like chatbots typically

stressed the importance of human judgement and human relationships in providing appropriate, responsive and empathetic customer service. Conversely, participants who were happier using chatbots typically did not feel that such human features were necessary, indeed in some cases a chatbot was preferred if it did not imitate human interaction (e.g. people may feel more confident to disclose information to a machine which does not have human judgement or emotional responses).

3.3 Algorithmic Decision-Making

Participants were told that algorithms – or computer programmes – are increasingly being used to quickly process large amounts of financial information and that these can identify patterns in the information and identify which loan applications are high risk or low risk. Participants were told that at the moment algorithms are used to assess suitability of loan applicants but the final decision is overseen by a human member of the bank's staff.

Some participants were fairly ambivalent about algorithmic decision-making:

“I think as long as there is some kind of appeals process I'd be fine with it because if you could see all of the information they have access to and you can say, "Yes, that is my credit card. Yes, that is the thing that I did there," as long as that's there, as long as there is a way to say, "I think something is not right, can I speak to a human being," I think it's fine. A lot of cases are quite simple and some cases are complicated. So the complicated situations might just be escalated but as a general thing, I think they're just taking numbers and data anyway. It's not like they're doing a personality assessment. So I think it's generally fine, in my opinion.” (POSTGRADUATE STUDENTS: F5)

Most participants felt that automated processes could be positive in increasing speed and efficiency of decision-making processes and were satisfied provided that final decisions were taken by humans:

M1: for me I think it's good the way it is, to have the algorithm assisted by the humans

F1: yeah

M2: because then in a way you've got two different opinions whereas if you had just the human they could make a biased or prejudiced decision, or even the algorithm itself. So I like it how it is

F2: yes, I suppose if you wanted a loan or a mortgage it would be helpful so long as there was a backup

F1: yeah, it would be quicker

F2: yeah, I suppose it would be a good idea so long as someone is keeping an eye on things

(SENIORS)

Across all the focus groups human oversight was described as a vital condition for acceptability of algorithmic decision-making. Participants typically stressed that they would not be comfortable with final decisions being made by an algorithm. This was

largely due to a sense that algorithms cannot fully take account of people's individual circumstances or the complexities of people's lives:

F1: People freelance as well, they work for themselves, you have to show the last three months. Your income might be one thing for three months, huge amounts for the previous but an algorithm is not going to pick up whereas you can see it in person, however in the last year I earned this. There is no way to do that when it's all just done on algorithms. I think you need a human person making the ultimate decision.

F3: Exactly. I don't like this. It's very black and white. Like you say, life is not black and white. You do need to have that human interaction to make decisions and to look outside of the digital box. People don't just tick boxes everyday because life isn't like that. So I disagree.

(RURAL)

"I would be scared if it's completely automatic in the future because [...] the algorithm can't say, "Is this good? What is the reason for this?" and so on. So I would always prefer to have a human overseeing the whole process because algorithms can also be completely faulty like with You Tube where they just recommend completely old videos or videos you've already watched. It's just complete mistakes, the evaluation of the video or anything like that. So I don't think I would like to have it completely. Maybe as a help but always have a human before you with that kind of thing"

(INTERNATIONAL STUDENTS: M2)

It was frequently acknowledged that all systems – whether based on human judgement or algorithms – contain bias and as such neither automated nor human decision-making were considered to be free of bias. However, it was largely suggested that it was easier to deal with and address human biases, either through discussion or by seeking a different person:

"I think you do find biases in algorithms in certain things, whereas a person can have biases but that may vary person to person. So if you try to get a different person, you might get a different result, whereas I hate technology and I think personally I would prefer it to be a person who made the final decision rather than an algorithm."

(POSTGRADUATE STUDENTS: M2)

Participants also felt that automated processes would reduce the opportunities for customers to discuss their needs and circumstances, appeal decisions or ask for advice:

F1: I think it's kind of good that there's a human still involved

F2: Yeah

F1: because people might have previously been in bad financial situations and a computer – I don't know how smart these computer are – they might not then look at it from a human perspective and say well they were like this, but look at them now this is their

salary and they're doing a lot better, they can kind of prove that they can afford a mortgage or whatnot. But an automated system might just say no, they were bad previously, so they're bad. But I don't know how smart they are about it.

F2: I think when you've still got a human like at the moment, that kind of fall back to know the particular circumstances, so you can go back and say I know I was rejected for these particular terms but this is what I need them for, this is what I want them for and I don't think a computer could do that

F1: Or you can ask what do you recommend. What are the options. You can't really talk about that with a computer.

(YOUNG PROFESSIONALS)

"I think it depends what the algorithm looks at because if it's a physical person, they can weigh up the pros and cons whereas an algorithm, it's either yes, no dependent upon your credit score or whatever. But I know somebody who had entries from somebody else's credit file on there so bringing their credit score down. So with the computer they wouldn't have gotten anywhere" (POSTGRADUATE STUDENTS: F2)

In summary, while some participants were positive about the potential for algorithmic decision-making to speed up processes, there was considerable concern about automated decision-making. In particular, it was widely agreed that human judgement and oversight are essential to ensure that decisions adequately take account of individuals' circumstances and needs.

3.4 Broader Themes

Through discussion of particular examples a number of broader themes emerged. In particular there were some key areas of concern that emerged consistently across the focus groups.

Pace of Change

Across the focus groups there was considerable discussion of the rapid rate at which technological innovation is advancing. For many this was an area of concern:

"urgh. It's too much, it's like sci-fi. All these stories you read. Don't laugh: it could happen! Machines taking over and running the world. Everything's moving so fast" (SENIORS: F1)

"Things are maybe moving a little bit too far [...] So it really is about, I think I'm being a Luddite as well in pushing back against technology a little bit with banking. I have a feeling people are going to do that a bit more, even the younger generation when they figure it out, will be doing that as well." (RURAL: M1)

"I think probably banking technology has gone as far as I need it to go because maybe I'm traditional or old fashioned but a lot of this is just a bit of a step too far, for me

personally. I don't know. I don't want everything to be, all the decisions, I don't want that all to be by computer and not people” (RURAL: F3)

“I think it is quite scary how things are evolving now in the future but I think there is no stopping it. Things are constantly being evolved and companies are constantly evolving to be easier and better which is to the detriment of, this is maybe pessimistic but it takes out that human to human interaction of a bank and stuff which maybe some elderly people, it's the only interaction they have in a week or whatever.” (POSTGRADUATE STUDENTS: F6)

However, it was frequently mentioned that while new technologies can be unnerving, typically when they are convenient or beneficial people adapt to using them relatively easily:

F1: I mean probably in the next twenty years everybody will be computer literate

M1: it's about easing people into change, because we're all afraid it, but 9 times out of 10 it's ok so there needs to be something that makes us all feel-

F2: safe

M1: yeah, safe

(SENIORS)

“But everyone for example in this room said at the beginning that all of us use the banking apps but then when we talked about it everyone didn't like it. But when you just see the final product, everyone likes it. I think they probably didn't like the journey or when it was developed. People probably 15-20 years ago they didn't like it but us, we're just born with it, we think everyone just enjoys it. We wouldn't imagine life without it.” (INTERNATIONAL STUDENTS: M1)

“I think it really depends heavily on the context you grow up in and the time period because some of the things that we use every day like mobile banking, if we pitched that idea to our parents and grandparents they would say, "Well no, I'm going to go into a bank for the rest of my life. That's a crazy, stupid idea." If we have this conversation in ten years, all of these things that are hypothetical now could be part of everyday practice.” (POSTGRADUATE STUDENTS:F5)

Slippery Slopes

Related to discussions of the pace of change, there were significant concerns about potential “slippery slopes”. A number of participants were concerned about the ways in which data used in financial apps such as virtual money coaches might potentially be reused in the future. Beyond particular data uses, there was also concern about potential “slippery slopes” around the role of AI and impacts on society. Participants expressed concern that increasing reliance on AI in automating decision-making processes could erode human capacities and expertise:

F1: I think it's kind of worrying if now we have this one and then we depend on AI to give us advice. Say AI is getting better and then everyone can, let's say AI can really does the job and then everyone just talks to AI and then no more human interactions and then we are all relying on this AI to give us whatever advice. It's going to take over our thinking. We cannot think as a human but it's like we let them think and then they tell us. I'm just worried.

M2: Yes, that's what I meant especially with teenagers and people who grew up then in that time because they don't know anything else.

F1: They will always think, "Okay, there is an app, AI, that can always give me that."

M2: They depend on it.

M1: At the same time, these apps, these systems need humans behind the scenes developing them. They wouldn't just happen from nowhere. You'd need so many brains behind the scenes developing, debugging, expanding those systems.

F3: When they take over, suppose we are asking finance questions, obviously finance people are only in the development phase, they are only telling them, "You have to advise this." [...] I think if I'm not using my brain for ten years and suddenly, even if I'm a finance person and for example you come and ask me and I'm like I don't know, I don't remember

(INTERNATIONAL STUDENTS)

This related to concern that AI could "take over":

"If this did come about, how long would it before artificial intelligence takes over the whole thing?" (SENIORS: F2)

The risks of AI and automation increasingly replacing human judgement and oversight was a major theme which emerged consistently across all the focus groups and which was articulated in relation to each of the examples discussed above.

"I think personally that there's got to be some sort of human interaction no matter how far the technology goes because a human can spot if there's something wrong, whether it's through feeling or patterns or whatever, which technology might not. Say someone has got dementia and they've applied for five credit cards, is the algorithm going to spot that or is someone physically looking at stuff going to be like, "Well something is not right here"?" (POSTGRADUATE STUDENTS: F2)

"I also believe that no matter what we should rely on ourselves. Technology is good but human connection is actually important in our life I think." (INTERNATIONAL STUDENTS: F4)

Related to this a number of participants discussed the potential for increasing use of AI to result in job losses and bank branch closures:

F1: they're going to put people out of work aren't they?

F2: yeah, that's a good point. They do away with all- oh you can laugh but it's true! They're doing away with lots of jobs. I mean there's banks closing all over the place. I think my local one will be next
(SENIORS)

"It's also like because people will just get lazier and lazier and lazier. Also if everything they can do with AI and then humans will have no job at all so in society they will do nothing and they will have no earning, no income and all that." (INTERNATIONAL STUDENTS: F4)

"I think it's just becoming more convenient, isn't it, to everybody. I love my banking app because I can do stuff while I'm on the move and you're not standing in queues all the time but is that necessarily a good thing because you're taking jobs off people."
(POSTGRADUATE STUDENTS: F2)

These discussions indicated tensions between people's concerns about societal impacts of AI in banking, and their own willingness to use technologies which were perceived as increasing convenience. The discussions also highlighted that ethical concerns were most often raised not in relation to impacts on individual service users or bank customers, but rather on wider society. This has important implications for the ways in which ethical considerations may be conceptualized and addressed.

4 Discussion

The focus groups covered a wide range of topics relating to digital innovation in banking. Throughout the focus groups, and in conversations after the groups had formally ended, participants often remarked that these were not topics they typically gave much consideration to, yet many participants said that they enjoyed the discussions and that it had been helpful to consider these areas and develop their own views. Moreover, focus group participants remarked that they did not generally think about banking in their day-to-day lives and indeed they did not want to have to think about banking. As one participant said they "don't want to be too involved with [their] bank" (RURAL: F1). Where people described current frustrations with banking these typically related to inconveniences or friction in accessing services. Therefore, digital innovation was more likely to be viewed positively if it increased convenience and resulted in faster or smoother experiences. As has been observed previously (13), this suggests that new technologies or data practices which lead to frictionless banking services have a high chance of being adopted.

However, this points to a tension in the focus group findings. On the one hand, participants want faster, frictionless services and do not want to have to think about those services in their day-to-day lives: they want banking just to "fit in" with their lives. Yet on the other hand, there was considerable concern across the focus groups that increasing automation of services has negative impacts for both individuals and society. In particular there was concern about the potential for reliance on automated processes

and AI to erode human capacities and skills. There was quite significant concern that this might reduce future generations' abilities to make sound financial decisions, or to make decisions independently, as well as reducing the skills, expertise and authority of employees in the banking sector. As such there was a clear tension between participants' desire for greater convenience and their concerns for potential negative societal impacts of increasing automation. This may suggest a potential cognitive dissonance (14) where people use new services due to perceived convenience or immediate benefits, while disliking or distrusting those services or holding concerns about their impacts on society.

Concern about a lack of human judgement or oversight was a major theme which emerged consistently across all the focus groups. Participants typically stressed the importance of human judgement and interaction in relation to financial decision-making (e.g. determining loan applications), customer services (e.g. in contrast to chatbots) and advice (e.g. in contrast to virtual money coaches). AI, or automated processes, were largely perceived as unable to take account of the complexities of real lives or the particularities of individuals' financial circumstances. There was a preference for human interaction in order to enable decision-making and/or advice which engaged more fully with these dimensions. Additionally, human involvement was seen to enhance accountability and trustworthiness which gave focus group participants more confidence to discuss financial matters or access services.

Many of the concerns raised reflect what Taddeo and Floridi (9) have described as the invisibility of AI. This relates to the scope for AI to exert increasing, but imperceptible influence over our lives and identities and to undermine individual or society's control. These concerns were consistently expressed throughout the focus groups in discussions of the ways in which AI might reduce human capacities or individual's ability to make sound decisions. This highlights that public concerns were not limited to short-term or immediate impacts of new technologies, rather they reflected an awareness of – and concern about – long-term societal impacts.

Nevertheless, most participants were not opposed to automated processes or AI in banking. Automation was widely recognised to have a valuable role to play in increasing efficiency and convenience of banking services. Yet, across the focus groups human oversight was consistently articulated as a vital condition for acceptability of AI and automation. In particular, where algorithms are informing decision-making processes, focus group participants expressed a clear preference for final decisions to be made by humans. Similarly, while automated processes were recognised to be useful at answering frequently asked questions, human interaction was considered important when customers had unusual queries or sensitive matters to discuss.

5 Conclusions

The focus group discussions demonstrated participants' enthusiasm to engage with the subject of AI in banking. The topic guide for the focus groups was intentionally flexible to allow unanticipated areas of interest to emerge, while particular examples presented worked well as catalysts for discussions, across the focus groups this typically led into

broad discussions raising considerations of impacts on society and future generations. This illustrates that participants' concerns did not typically relate to private or individual interests but more often to wider ethical and social concerns, highlighting the importance of taking account of conditions for public acceptability - rather than just customer uptake - in order to develop ethical practice and establish a social licence (1) for uses of AI in banking.

Parallels can be drawn with research into public attitudes towards data-intensive innovation in other sectors. For example, research into public responses to uses – and reuses – of patients' health data has found that members of the public are primarily concerned with the extent to which data will be used in ways which bring public benefits and conversely the possibility for indirect negative, societal impacts of data use (15). There may be some value in looking at the wider literature around public attitudes to data practices in order to reflect on conditions which need to be met in order to establish a social licence for data-intensive innovation in FinTech and banking.

Pursuing a social licence recognizes that there can be meaningful differences between what is legally permissible and what is socially acceptable (1). Therefore, establishing and maintaining a social licence requires going beyond legal compliance in relation to safeguarding customer's data and instead taking steps to align with public values and expectations. This research has indicated that doing so may necessitate a focus on long-term, indirect impacts of innovation and taking approaches which are aimed at societal – rather than just individual – benefits.

The focus groups demonstrated the value of qualitative, deliberative methods to explore the nuances of public responses. In particular, they revealed a potential cognitive dissonance whereby individuals are concerned about new areas of innovation and yet use these services due to perceived convenience or immediate benefits for individuals. This highlights the importance of considering wider public responses - rather than simply reasons for customer uptake of new FinTech services – in order to identify and address ethical concerns.

Further research is needed in this field to inform emerging practices in the fast-paced industry and to ensure public concerns and interests are addressed. To date, the limited literature around public, or customer responses to FinTech and/or digital innovation in banking has tended to focus on individual concerns (e.g. regarding privacy and security) or perceived benefits to individual customers. However the focus groups have demonstrated that such matters form only a small proportion of the concerns held by members of the public. This suggests that establishing a social licence for AI in banking (1) requires much more than consideration of privacy and security or direct customer experiences, but rather necessitates attending to broader social considerations.

As is a common feature of focus group research, this study enabled a deep understanding of the nuances of public attitudes and the conditions for public acceptability. However, the focus on depth of understanding meant that the number of participants in the study was small and it cannot claim to be representative of wider public views. This represents an important first step in building the evidence base in this field. Further, qualitative and deliberative research is needed to build on this study and further explore wider public attitudes towards AI in banking.

6 Acknowledgement

This research was funded by the EPSRC, grant reference: EP/R033595/1.

7 References

1. Aitken, Mhairi, Ehsan Toreini, Peter Carmichael, Kovila Coopamootoo, Karen Elliott, and Aad van Moorsel. "Establishing a social licence for Financial Technology: Reflections on the role of the private sector in pursuing ethical data practices." *Big Data & Society* 7, no. 1 (2020): 2053951720908892.
2. Maskey, Sameer. "How Artificial Intelligence Is Helping Financial Institutions." *Forbes* (2018).
3. Floridi, Luciano. "Establishing the rules for building trustworthy AI." *Nature Machine Intelligence* 1, no. 6 (2019): 261-262.
4. Toreini, Ehsan, Mhairi Aitken, Kovila Coopamootoo, Karen Elliott, Carlos Gonzalez Zelaya, and Aad van Moorsel. "The relationship between trust in AI and trustworthy machine learning technologies." In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, pp. 272-283. 2020.
5. Chuang, Li-Min, Chun-Chu Liu, and Hsiao-Kuang Kao. "The adoption of fintech service: TAM perspective." *International Journal of Management and Administrative Sciences* 3, no. 7 (2016): 1-15.
6. Gulamhuseinwala, Imran, Thomas Bull, and Steven Lewis. "FinTech is gaining traction and young, high-income users are the early adopters." *Journal of Financial Perspectives* 3, no. 3 (2015).
7. Gulamhuseinwala, Imran, Matt Hatch, James Lloyd, T. Bull, and S. Chen. "EY FinTech Adoption Index 2017: The rapid emergence of FinTech." *Ernst & Young Global Limited* (2017).
8. Kavuri, Anil Savio, and Alistair Milne. "FinTech and the future of financial services: What are the research gaps?." (2019).
9. Taddeo, Mariarosaria, and Luciano Floridi. "How AI can be a force for good." *Science* 361, no. 6404 (2018): 751-752.
10. Mittelstadt, Brent. "Principles alone cannot guarantee ethical AI." *Nature Machine Intelligence* (2019): 1-7.
11. Demirguc-Kunt, Asli, Leora Klapper, Dorothe Singer, Saniya Ansar, and Jake Hess. *The Global Findex Database 2017: Measuring financial inclusion and the fintech revolution*. The World Bank, 2018.
12. Barbour, Rosaline. *Doing focus groups*. Sage, 2008.
13. King, Brett. *Bank 4.0: Banking everywhere, never at a bank*. John Wiley & Sons, 2018.
14. Festinger, Leon. *A theory of cognitive dissonance*. Vol. 2. Stanford university press, 1957.
15. Aitken, Mhairi, Jenna de St Jorre, Claudia Pagliari, Ruth Jepson, and Sarah Cunningham-Burley. "Public responses to the sharing and linkage of health data for research purposes: a systematic review and thematic synthesis of qualitative studies." *BMC medical ethics* 17, no. 1 (2016): 73.