University of Tartu
Institute of Philosophy and Semiotics

THE PROBLEM OF CREEPY PERSONALISED RECOMMENDATIONS IN BANKS

Master's Thesis in Philosophy

Palina Yaroshyk

Supervisors: Alexander Stewart Davies, Lennart Kitt

Number of characters: 110953

Year
2025

Table of contents

Abstract

Banks face a growing challenge: while personalised recommendations are extremely useful for both banks and their clients, recommendations that are personalised sometimes come across as “creepy,” undermining trust and rendering the recommendation ineffective. To address this, the thesis asks two questions: (1) What makes some personalised recommendations creepy? and (2) How can banks prevent this creepiness while maintaining personalisation?

The first, and most important, question is answered by showing that creepiness stems from a lack of grounded common knowledge that would make recommendations warranted and expected. What is lacking is a client’s awareness that a bank knows that a client wants to have a solution to their specific issue. While the bank knows, the client does not realise that this need has been identified, making the recommendation feel abrupt and creepy, as it arrives suddenly.

This insight leads directly to answering the second question: creepiness can be mitigated by designing micro-interactions (such as surveys, tests, or chatbot conversations) that invite clients to signal interest before receiving recommendations.

Kokkuvõte

Pangad seisavad silmitsi kasvava väljakutsega: kuigi isikupärastatud soovitused on potensiaalselt kasulikud nii finantsasutustele kui ka nende klientidele, võivad need mõnikord mõjuda häirivalt või “kõhedust tekitavalt”. See võib õõnestada klientide usaldust ja muuta soovitused ebaefektiivseks. Käesolev töö keskendub sellele probleemile, esitades kaks peamist uurimisküsimust: (1) Mis muudab osa isikupärastatud soovitusi kõhedust tekitavaks? ja (2) Kuidas saavad pangad vältida selle kõhedustunde tekkimist, säilitades samas soovituste isikupärasuse?

Esimesele, olulisemale küsimusele vastatakse, näidates, et kõhedustunne tekib eelkõige maandatud ühisteadmiste puudumisel – klient ei mõista, miks või kuidas soovitus tekkis, ning see muudab soovituse ootamatuks ja seeläbi häirivaks. Eelkõige puudub kliendil teadlikkus sellest, et pank mõistab kliendi soovi lahendada konkreetne probleem. Kuigi pangal on vastav teave ning tahe tegutseda sellele tuginedes, jääb kliendile mulje, et soovitus tuleb ootamatult ja põhjendamata, loob selline käitumine kliendile kõhedustunde.

Sellest arusaamast tuleneb vastus ka teisele uurimisküsimusele: pangad saavad sellist kõhedust leevendada, kujundades teadlikult mikrointeraktsioone (nt lühikesed küsitlused,

enesehindamistestid või vestlused juturobotiga), mis võimaldavad klientidel ise väljendada oma huve ja vajadusi enne, kui personaliseeritud soovitused neile esitatakse.

Annotation

This thesis explores why some personalised banking recommendations are perceived as creepy and how this reaction can be avoided without reducing personalisation. It argues that creepiness arises when recommendations are delivered without grounded common knowledge, specifically, without client awareness of their specific need recognition by a bank. The proposed solution is to introduce micro-interactions that allow clients to express interest before receiving recommendations, thereby making them feel expected and appropriate.

Keywords: banks, creepiness, personalised recommendations, common knowledge, grounding

1. Introduction

Personalised recommendations that banks give to their clients are often very helpful, as such recommendations take into consideration the specific circumstances of an individual. In some cases, however, banks seem to cross the line between helpful and creepy when giving personalised recommendations. By pursuing ways to make recommendations more personalised to an individual, banks sometimes overlook something crucial, which causes their efforts to improve client engagement to become *creepy*. Creepy recommendations create a negative experience for clients, which in turn affects the relationship with customers and a bank's profit. In this thesis, I aim to find out what this crucial something is that banks sometimes overlook and to suggest a structured way to both benefit from personalisation and ensure that banks do not risk making personalised recommendations creepy. Having stated the problem banks face and the aim of this investigation, in order to better understand what the issue of creepy personalised recommendations in the banking context is, let us look at it in more detail.

It will be useful at this stage to consider why banks give recommendations to their clients at all. To start from afar, people want to manage their finances effectively. They want to know how to accumulate more and waste less. Banks, as financial institutions, have a lot of experience and specialised knowledge that, if shared with clients, can help them navigate their financial strategies. One way to facilitate this is through recommendations, which allow banks to share their expertise with customers. When banks provide guidance, clients gain access to insights that can help them make better financial decisions. Here's an illustration of what a recommendation in a banking context looks like:

Account Options

In an email from the bank to a client: "To ensure you're not overpaying for an account that may not suit your needs, check out our account options alongside the terms and conditions on our website."[1]

Ideally, recommendations bring benefits to both parties: banks and clients. Banks, with the aid of recommendations, can advertise their services, gain profit, and earn the trust of their customers by supplying the demand for information and guidance. Often, a lot of services are available in banks upon direct request, but clients may miss the exact service they need because they get lost in an overabundance of information. So, clients, then, can

[1] Disclaimer: The cases presented in boxes throughout this thesis were generated using ChatGPT (OpenAI, 2025), refined by the author, and approved after the review for real-world plausibility by Lennart Kitt, Head of Customer Analytics and Data Science at SEB Eesti.

learn about relevant offers, make more informed decisions and, overall, improve management of their own finances by acting on recommendations from those who know better.

Banks give two kinds of recommendations: general and personalised. General recommendations, like the one mentioned in Account Options, provide only broad financial advice that is directed to a wide audience.[2] Even though general recommendations offered by banks may be useful, they fail to consider individual financial situations, which makes some of the recommendations irrelevant to a client's needs, wants, and circumstances. Moreover, even if a certain recommendation is relevant to a client, they may not fully understand its potential relevance to their case. This lack of relevance not only affects the usefulness of general recommendations but also contributes to customer frustration and disengagement, as evidenced by recent research. LXA reports that 70% of millennials feel frustrated when receiving irrelevant offers (Powell, 2022). Similarly, Statista notes that 78% of young consumers dismiss irrelevant offers, suggesting that general recommendations often fail to engage customers effectively (Statista, 2024).

Personalised recommendations, on the other hand, take into account the unique circumstances of the client and offer suggestions based on factors such as income, spending habits, financial goals, risk tolerance, recent account activity, purchases, and many more. Here's an illustration of what a personalised recommendation looks like:

Fifteen-euros

In an email from the bank to a client: "Hi Sarah, we've noticed that your transaction patterns mean you're paying around €15 more in fees each month than necessary. Switching to our FlexSaver Account could help you save up to €180 a year—contact us if you'd like a personalised review to optimise your banking!"

Personalised recommendations, unlike general ones, enhance customer satisfaction and strengthen the relationship between a bank and a client (Waupsh, 2022). When the content of recommendations is personalised, it increases client engagement: HubSpot reports that personalised call-to-actions perform 202% better than basic ones (HubSpot Staff, 2013). In simple words, a call-to-action is a specific phrase like "buy now", "subscribe", "share", "contact us", "check out", etc., that encourages a client or user to perform the directed action. The results of the statistics show that personalised

[2] I acknowledge the distinction between advising and recommending, as described by Habgood-Coote (2024); however, for the purposes of this thesis, I will not differentiate between them and will use terms such as "advising" or "suggesting" to mean roughly the same as "recommending."

call-to-actions convince many more people who view them to take the directed action compared to the persuasion success of non-personalised call-to-actions. Moreover, personalisation prevents clients from being bombarded with unnecessary, irrelevant information, and it makes recommendations more persuasive (Yoo et al., 2013, p. 4). Given the evidence, it becomes clear why banks and clients would prefer personalised recommendations over generalised ones: personalisation enhances recommendations' relevance to clients' specific circumstances, engages them, and increases satisfaction—all of which strengthens banks' relationships with clients and ultimately leads to greater profitability.

Once we take a look at how banks make their recommendations personalised, we will see that some personalised recommendations are more personalised than others. More precisely, depending on the data analysis method a bank uses, recommendations will have a greater or lower chance of being better suited and more specific to a client's circumstances. Importantly, we will see that the more personalised a recommendation is, the higher the risk that it can cross the invisible line and cause a negative experience to clients. Let us then take a look at what banks do with the data that clients consent to give away in exchange for a useful personalisation.

At first, in order to have recommendations personalised to a specific client, banks usually ask for consent to have access to their data and permission to process it. Then, once consent is obtained, banks analyse the data provided by clients by using various data analytic methods, for example, *scoring* (assigning numerical values to assess risk or creditworthiness) (Lyn et al., 2002), *profiling* (creating customer categories based on behaviour and selected characteristics) (Wangler et al., 1999), *decision trees* (using branching logic to make automated decisions) (Plapinger, 2017), or *trigger-base systems* (initiating recommendations based on specific customer events) (Goic et al., 2016).

Consider the following pair of personalised recommendation scenarios. They are going to illustrate the difference between (1) personalisation that uses very little of a client's data and (2) one that uses much more of it. In the first case, the recommendation is based on a minimal amount of data, specifically, an inference drawn from a single piece of information that triggered the recommendation system (RS) to send a personalised suggestion:

Sarah's Investment Account

Five minutes after a client opens an investment account, she receives the following email from a bank: “Hi, Sarah, congratulations on your investment account! Would you like to explore our investment opportunities?”

This recommendation is personalised, as it takes into account the client’s recent account activity. It is triggered by a single event, the opening of an investment account, which prompts a relevant suggestion. Such recommendations are often helpful, as they are timely and directly related to the client’s actions. They also serve as a useful source of additional information that the client may find beneficial.

In the following case, however, the recommendation is more personalised. Here, a personalised recommendation is not triggered by a single piece of data, but rather by a complex analysis of multiple data points, like in the profiling data analysis method, making the suggestion more uniquely personalised to this specific client’s circumstances:

Peter’s Divorce

After noticing changes in Peter’s spending patterns, such as separate apartment rent, reduced joint account usage, and increased legal consultation fees, etc., the bank sends him a personalised offer via email: “Dear Peter! We would like to offer personalised financial counselling for recently divorced individuals, along with a recommendation to open an individual account for greater financial independence. We're always here to help you take control of your financial future.”

Even though the recommendation given in Peter’s Divorce is more personalised than the one in Sarah’s Investment Account, and it has a greater chance to “hit the spot” and be precisely what Peter needs because it uses more data, the recommendation seems quite unsettling. When we realise that Peter did not share his upcoming change in marital status with the bank, but somehow the bank found it out by analysing Peter’s data and giving a personalised recommendation, this recommendation comes to seem creepy.

Immediate reactions may be “How did they know?” and “Why does the bank give me such recommendations?”, and, perhaps, “It is none of their business”. Whatever the precise thought Peter may have had, this recommendation probably caused a negative experience of interacting with a bank for Peter as a client. It seems that banks missed something important that turned well-intended and potentially useful personalised recommendations creepy. And these are the exact kinds of cases of creepy personalised recommendations in banks that will be our interest in this thesis — the kind that provokes such reactions of surprise and discomfort when banks make unexpectedly accurate inferences about clients.

Considering both that more personalisation gives recommendations a better chance of being what clients and banks want, and that greater personalisation risks making recommendations seem creepy, this creates a tension that leaves banks wondering what to do with their RS. The tension is as follows: how to make personalised recommendations as precise as possible while nonetheless avoiding the risk of creepiness it can cause at the same time. Since we can now clearly see the problem banks have, we can consider three options for what banks can do about it. Let us take a look at them.

One approach banks can take to avoid such cases like Peter's Divorce is to refrain from offering personalised recommendations based on complex analysis of customer data altogether. This would eliminate the risk of providing clients with recommendations that they will consider creepy, but, at the same time, this would also mean that they miss out on a lot of potentially highly valuable insights that could improve their financial decisions. For instance, a client might overlook the fact that they're paying €15 a month in unnecessary fees, as in Fifteen Euros, because generic and primitive single-trigger personalised advice cannot identify such specific insights. Banks, in turn, would lose an opportunity to enhance customer engagement and satisfaction through personalised services and, therefore, drastically lower their profit. This approach, then, even if effective, is not the most optimal one if banks and clients wish to continue benefiting from personalisation.

The second solution to the problem is to downplay the size of the problem posed by creepy personalised recommendations. However, after we look at the statistics, we will see clearly that this is not a feasible approach either.

The data shows that creepy personalisation drastically worsens client engagement. We are not talking about 2 or even 10 percent, but more than half of customers will unsubscribe, and 38% will stop doing business with a company if they find personalisation efforts creepy (Gartner, 2019). That is a huge number of clients walking away. Far too many for banks to ignore. If they keep pushing creepy personalised recommendations, they risk losing a massive chunk of their customer base and damaging their reputation.

Moreover, if banks close their eyes, their clients won't. We can see that the problem of creepy personalisation is publicly addressed: recent news reports show that some customers are voicing concerns and even taking legal action against companies over creepy recommendations (Ghaffary, 2024; Dean, 2025).

Our approach in this thesis is to go for a third, and most promising, solution. This solution is to confront the problem directly by investigating what makes personalised recommendations creepy, then identifying methods for eliminating the source of creepiness without reducing personalisation. This approach would allow banks to eliminate the risk of creepiness while maintaining the high value that personalisation provides.

To adjust personalised recommendations, we need to first of all get clear on what exactly we should adjust. Once we know what exactly banks should do to lower the risk of creepiness in personalised recommendations, they will be able to solve the previously stated issue, namely the tension between usefulness and creepiness of some personalised recommendations.

In the next few paragraphs I will show that the creepiness of personalised recommendations in the banking context is a unique issue that requires a new approach, as it cannot be reduced to two well-known problems in data processing: (1) epistemic errors in data analysis and low data quality or (2) legal violations in data handling. This suggests the issue of creepy personalised recommendations requires an alternative perspective that will explain what kind of flaw in the banking context creepiness is. Let us slowly examine these two possibilities, in turn.

Let's consider first (1): that a creepy personalised recommendation is an epistemic error in data analysis and data quality. In simple terms, research data and data analysis concerns how good the information about the client was (was it outdated or not, was it full or partial, was it correct or faulty, etc.) and how well one reasoned through this information (how good the machine learning algorithm used is, or how well human did the analysis in the case of human advisors). Recall Peter's Divorce. Suppose that Peter is not actually divorced. In that case, a personalised recommendation that has flaws in the data or data algorithm (and it has epistemic flaws too, as it is based on a false belief) is practically useless because it will not be relevant to Peter's situation. It seems that if there are no flaws in the data and data analysis, as in Peter's Divorce, and he is actually getting a divorce, the recommendation would be potentially useful. However, the creepiness will not disappear. It will arise partially in virtue of this personalised recommendation being accurate and relevant to Peter's situation. Making recommendations less personalised, as we mentioned, is not the solution we want to apply; however, as in this case, both banks and clients would miss out on the many benefits personalisation brings. Therefore, we must look elsewhere to find a way to make recommendations without the risk of being creepy.

Let's now consider (2): that a creepy personalised recommendation is a flaw in legal compliance. Is creepiness a legal flaw? While violating consent and other regulations provided by privacy laws like GDPR (General Data Protection Regulation) can result in a recommendation that is both illegal and creepy, creepiness isn't inherently tied to illegality.[3]

For example, GDPR Act 6 implies the need for explicit user consent before personal data can be processed for personalised services. However, even when a bank operates within legal boundaries, with a client consenting to their data being analysed for personal recommendation purposes, a recommendation can remain creepy, as we can see from Peter's example. Legal compliance ensures that data processing is formally justified, but it doesn't address negative emotional responses to personalised, creepy recommendations. This means that when users provide consent for the use of their data, the creepiness they experience with personalised recommendations can't be framed as a legal consent violation. Peter cannot legitimately object on the grounds that he never consented to this specific recommendation, as consenting to specific personalised recommendations is implied when one consents to receiving personalised recommendations in general. Therefore, the solution must go beyond legality, too.

Thus far, we have seen that creepiness in personalised recommendations is neither solely a result of flawed data processing nor a legal shortcoming. The question, then, remains: what flaw is it?

To sum up our detailed introduction, the main problem we want to address in this thesis is that while personalised banking recommendations offer significant value, they sometimes cross a line into creepiness, even when no flaws in data analysis, data quality, or flaws in legal compliance are present. We are now in the position to describe the strategy this thesis is going to apply. The aim is to answer two main questions:

(1) What makes some personalised recommendations creepy?
(2) How can banks maintain high personalisation and creepiness-free recommendations?

This thesis takes the approach that knowing what exactly the problem is is already half of the solution. So, once we learn what is the cause of creepiness in personalised recommendations in banks, we will be able to propose possible solutions that banks can

[3] On the official GDPR website GDPR is explained to be "the toughest privacy and security law in the world. … it imposes obligations onto organizations anywhere, so long as they target or collect data related to people in the EU. The GDPR will levy harsh fines against those who violate its privacy and security standards, with penalties reaching into the tens of millions of euros" (Wolford, n.d.).

implement. While this approach requires more effort than simply abandoning personalisation altogether or ignoring the problem, it represents a worthwhile endeavour. It would allow both banks and clients to continue benefiting from personalised services while ensuring these recommendations remain free from creepiness.

In this thesis, I am going to use the following methodology: before answering the first question, which is, "what makes some personalised recommendations creepy?", I will define relevant concepts. Firstly, I will define what a good personalised recommendation is, so that when we look at creepy personalised recommendations, we will be able to clearly see what criteria it did not fulfil to be considered good. Then, I will define what creepiness is, so that we know how something can become creepy, which will later help us to identify what makes some personalised recommendations, in particular, creepy. To answer the second question, which is "how can banks maintain high personalisation and creepiness-free recommendations?", then, we are going to apply insights from the answer to the first question, establish a strategy banks can use, and show how the solution can be applied in practice. In the remainder of this introduction, I will describe the structure of the thesis.

The thesis firstly defines good personalised recommendations (Section 2). It starts with looking at recommendations in general, then moves on to describing what is essential to personalised recommendations specifically, as well as what constitutes a good personalised recommendation. Next, it defines creepiness (Section 3) by first reviewing literature on creepiness and later defending the best way to describe what creepiness is. With these definitions in place, the thesis then looks at the intersection of the two—creepy personalised recommendations (Section 4). Section 4 returns to Peter's Divorce and looks at three variations of this case where the personalised recommendations are not creepy. By analysing what unites these three cases in lowering the risk of creepiness and what differentiates them from Peter's Divorce, we will find a creepiness-inducing factor that Peter's Divorce, in particular, and similar cases of interest share. This insight is going to answer the question of what makes some of the personalised recommendations in banks creepy. Finally, Section 5 translates these insights into practical strategies for banks, demonstrating how to redesign recommendation systems in an innovative way to maintain personalisation while eliminating creepiness. The thesis concludes by acknowledging limitations and reiterating the broader implications of its findings.

Ultimately, I will answer the two questions stated in the following way:

Firstly, what makes some personalised recommendations in the banking context creepy is the lack of grounded common knowledge, i.e. creepy personalised recommendations are those given without explicitly or implicitly informing clients that the bank knows that the client wants a solution to a specific problem, and this makes personalisation unexpected and disturbing.

Secondly, banks can improve their strategy of giving personalised recommendations by introducing micro-interactions that invite clients to actively participate in the recommendation-generating process before a recommendation is presented. These interactions can take the form of tests, surveys, conversational prompts via chatbots, or other small tasks that would nudge clients to express their desire to solve their specific issue before a personalised recommendation is given. By doing so, banks ground common knowledge that establishes a clear context, where recommendations will be socially acceptable, expected and welcomed, and thereby significantly lower the risk of creepiness.

2. Personalised Recommendations

To answer our first question stated in the introduction, which is "what makes some personalised recommendations creepy?", we must begin by clarifying two things: what good personalised recommendations are and what creepiness is, since we cannot fully understand what makes some personalised recommendations creepy unless we first understand the basic practice of recommending and what it takes for something to be creepy.

So, this section starts by looking at recommendations in general, then moves on to describing what personalised recommendations have to have in order to be called *personalised* recommendations, as well as what constitutes a *good* personalised recommendation, both from the perspectives of recipient (what good personalised recommendation should do) and a recommender (what good personalised recommendation should achieve).

A recommendation, in simple terms, can be defined as a communicative act in which one person suggests a course of action to another (Habgood-Coote, 2024, p. 13). A key feature of a recommendation, for example, is that it must, in principle, be possible to accept or reject. Unlike an order or a command, which demands compliance, a recommendation presents an option while acknowledging the recipient's agency (Habgood-Coote, 2024, p. 4). Imagine a recommendation was framed in a way that makes refusal impossible. In this case, the utterance risks no longer functioning as a recommendation. The absence of optionality and other essential features doesn't just make a recommendation bad, but it can make it cease to be a recommendation altogether, which will make it impossible to assess it as either good or bad.

Essential features of recommendations are also present in personalised recommendations, as a special kind of recommendation. However, since our focus is on cases where recommendations are already clearly present and recognised as such (even if they sometimes fail in other ways, like being creepy), we won't need to explore all these foundational features in depth. Instead, we'll take for granted that recommendations, in our context, are properly formed and identifiable, so we can concentrate on what makes some of them creepy.

When it comes to personalised recommendations, this basic definition requires further refinement, as a personalised recommendation is a special kind of recommendation. Personalisation introduces an additional dimension that we should consider. Precisely, personalised recommendations do not just offer general advice that could apply to anyone, as established in the introduction, but they are tailored to the specific wants, needs, or circumstances of a particular individual.

So, let us now turn to defining what a personalised recommendation is. To detect in Section 4.1 how exactly creepy personalised recommendations fail, we need to understand what makes a good personalised recommendation. To achieve this, we need to answer the following three guiding questions:

1) What features are essential to a recommendation being called personalised, or what must be present for it to count as a personalised recommendation at all?
2) What should a good personalised recommendation do for the person receiving it?
3) What outcomes or effects should a good personalised recommendation achieve, or how can a recommender judge whether it has succeeded in the best case?

Primarily, as trivial as it sounds, personalised recommendations should include personalisation. As we mentioned in the introduction, in banking, personalisation comes from using a client's data they consent to share, but to generalise, we can say that personalisation is based on information about the recipient of this recommendation.

In contrast to personalised recommendations, general recommendations do not need to be tailored to the specific characteristics, needs, or circumstances of a particular individual. One can make a general recommendation that is completely irrelevant to the recipient. In such cases, where a recommendation is irrelevant to the recipient, it would simply be a bad recommendation.[4] However, if a supposedly "personalised" recommendation fails to account for the recipient's specific situation, it would be incorrect to call it merely a bad personalised recommendation. Rather, the personalisation would be entirely absent, meaning it wouldn't properly qualify as a personalised recommendation at all, as it would just be a bad recommendation (perhaps mistakenly labelled as personalised).

Having established what makes an utterance count as a personalised recommendation in the first place, we now know what a personalised recommendation is. Our ultimate aim, however, is not simply to help banks make personalised recommendations, but rather to help banks make good personalised recommendations. So, in this part of the section, I am going to present four conditions that need to be present at once to ensure that a personalised recommendation is good.

Firstly, it should align with the person's specific goals.[5] A clear example would be advising an aspiring pilot to enrol in flight school. This advice directly supports their objective. However, a personalised recommendation becomes useless when it completely disregards these goals, like suggesting ballet classes to someone whose main objective is to become a pilot. Such a personalised recommendation is a bad one because it doesn't meaningfully contribute to the person's goals and is, therefore, useless.

Secondly, it must be realistically actionable for the recipient. A suggestion to "take the stairs" to reach a hilltop café fails to be a good personalised recommendation when given to a wheelchair user, as it's impossible to follow, given their circumstances. Even the most well-intentioned personalised recommendation becomes useless if the recipient lacks the means or ability to act on it.

Thirdly, for a recommendation to be counted as a good, personalised recommendation, it should be sincere. That is, the person making the recommendation should genuinely believe that the suggested course of action would be beneficial for the recipient. Without sincerity, a recommendation risks becoming manipulative, as it would no longer reflect a genuine attempt to promote the recipient's best interests, but instead might serve the recommender's hidden goals. This issue is particularly important in commercial contexts, such as banking, where institutions may have strong incentives to propose actions that are profitable for themselves but not necessarily optimal for the client.

[4] If we talk about groups of people, then we can say that a bad general recommendation is one that is irrelevant to every member of this group.

[5] It will be too demanding to say that good general recommendations should also fit specific aims of an individual. Rather, they should fit general aims of people to be considered good.

A recommendation that appears to offer financial advice but is actually motivated by the bank's pursuit of profit, rather than by a sincere concern for the client's well-being, would fail to meet the standard of a good recommendation. In such cases, even if the recommendation is presented in a persuasive way and seems actionable, it falls short because it lacks the intention to benefit the recipient.

These criteria explain why some personalised recommendations succeed and are considered "good" while others fail and are considered "bad". Even when they have all the essential features to count as a personalised recommendation, it can still be bad if it ignores the recipient's goals, circumstances, or is insincere.

This question brings us to the final layer of understanding good personalised recommendations, which is their ultimate purpose, aim, or, in other words, what effect they should produce. The produced effect can be counted as a fourth criterion of what makes a good personalised recommendation. However, I mention it separately for a key reason. The fourth criterion evaluates recommendations from the recommender's perspective. In contrast, the first three, which are the alignment with the client's goals, actionability, and sincerity, are assessed from the recipient's perspective. By highlighting this distinction, it becomes clear that a personalised recommendation only counts as good in our context, where banks provide recommendations to clients, if it meets the criteria from both the recommender's and the recipient's perspective.

Fourthly, consider: what if a personalised recommendation complies with the goal of the recipient, it is actionable upon, and is sincere, but it is delivered in such a way that a client dismisses it despite its utility. Would we consider it a good, personalised recommendation in the banking context too? Well, no, and here is why. Normally, we want to help a recipient and, at the very least, make them consider our suggestion and, in the best case, follow it. We want to persuade and compel the recipient to do the action that we think is going to be beneficial for them. This means that a good personalised recommendation should be compelling enough to prompt reflection or action.

Even the best personalised recommendations might be rejected. Just because someone doesn't follow it doesn't automatically make it a bad personalised recommendation. That is why we need to clarify that there is a difference between a recommendation not being taken into consideration because of the delivery and because of some circumstances that are independent of the recommendation itself. People might say "no" for many reasons that have nothing to do with the quality or delivery of the recommendation itself. So while we aim to make personalised recommendations persuasive, we can't judge them only by whether they are accepted. A good personalised recommendation is one that's helpful, makes sense for the person's situation, and is presented in a convincing way, even if the person ultimately decides not to follow it.

To sum up, the four criteria for a good personalised recommendation are presented in this table:

Criteria for a Good Personalised Recommendation	**Description**
Alignment with Goals	The recommendation matches the recipient's specific objectives.
Actionability	The recipient can realistically act on the suggestion given their circumstances.
Sincerity	The recommender genuinely believes the suggestion benefits the recipient.
Persuasiveness	The recommendation compels meaningful consideration or action.

Based on the definition of recommendation provided and the table we can define a good personalised recommendation as is a communicative act in which one person suggests a specific course of action to another (Habgood-Coote, 2024, p. 13), grounded in a true belief of what would benefit the recipient, aligned with their goals and circumstances, realistically actionable, sincere, and presented in a way that invites meaningful consideration.

To note, the definition mentioned that a good personalised recommendation has to be grounded in a true belief, and not, for example, a false belief. And there is a reason for that: imagine a case where a recommendation is well-formed in its structure but is based on false premises. Could we still call it good? In abstraction, perhaps. We might say it is good but mistaken. However, since banks aim to provide recommendations that are not just theoretically sound but genuinely useful for the client, such a recommendation would not qualify as good in practice. After all, if grounded in false beliefs, it would misunderstand a client's needs, misalign with their goals, and likely prove useless, or even harmful to them.

Since we wouldn't call a creepy personalised recommendation a good personalised recommendation, we can assume that creepy personalised recommendations fail in one of the specific ways we discussed. However, at this point, we don't know in what way exactly. In order to know that, our next step would be to establish what creepiness is and what it does to recommendations. So, to answer this question thoroughly (Section 4.1), we need to begin by defining creepiness, and only later can we find out what specifically makes recommendations creepy, and this is exactly what we are going to do in the following sections.

3. Creepiness

In the previous section, we described what a good personalised recommendation is. However, knowing what a good personalised recommendation is is not enough for us to determine in what way it fails as a creepy recommendation specifically. For that, we need to define what creepiness is and what causes creepiness. This section aims to fill this gap.

This section is divided into two subsections. The first subsection (3.1) will identify the causes of creepiness found in previous research, providing a necessary foundation for understanding how creepiness operates in Section 4. The second subsection (3.2) will then defend the most effective way to define creepiness, building directly on the findings from the first subsection. Once we establish this clear definition of creepiness, we can then examine what specifically makes some personalised recommendations in banking contexts creepy. Finally, after providing necessary definitions, we will be positioned to offer concrete suggestions to banks on how to refine their recommendations in Section 5.

3.1 Literature Review on Causes of Creepiness

If there is no consensus among researchers on what constitutes creepiness, then we will have to try to define, or at least find the best way to define, creepiness ourselves. To try to find a commonly accepted feature that causes creepiness, I've selected several studies that either explore creepiness as a general phenomenon or investigate it in a more specific context, like human-technology interactions. By comparing researchers' perspectives, we'll see that while there are recurring mentionings of similar features, scholars don't fully agree on what exactly constitutes creepiness or which factors matter the most. The absence of a unified definition described in this subsection justifies our endeavour to defend what constitutes the best way to define creepiness in the next subsection.

Now, let us take a look at what kinds of things are creepy, and what researchers identify as the cause of creepiness. Manifestations of creepiness vary widely. We can feel creeped out while watching a horror movie or hear the squeaking of the floor tiles while sitting home alone. We can also be creeped out by different people or their actions. At first glance, there seems to be little in common between, for example, a creepy clown and a creepy late-night call from a stranger. This conceptual fluidity makes creepiness difficult to pin down.

Faced with such a tricky phenomenon, current research offers fragmented insights rather than a unified definition, and scholars characterise factors that cause creepiness in different ways. Here are some features that scholars mention as potential sources of creepiness (the list is not exhaustive):

1) *Uncertainty of potential threat* (McAndrew & Koehnke, 2016, p. 10; Demetriou, 2024, p. 2; Bernstein & Nolan, in print, p. 7; Kjeldgaard-Christiansen, 2024, p. 2; Langer & König, 2018, p. 2; Herder & Zhang, 2019, p. 2). An example of this phenomenon occurs when walking at night on a street with a reputation for being dangerous, where one might potentially get shot. It is precisely this uncertainty of threat that creates the feeling of creepiness. By contrast, when the threat becomes obvious, when someone points a gun directly at you, the situation becomes scary rather than creepy.
2) *Moral insensitivity* (Fischer & Fredericks, 2020, p. 199). In simple words, it is a disposition to act or think in ways that disregard basic moral or ethical standards.
3) *Failed humanisation* (Demetriou, 2024, p. 2, Kjeldgaard-Christiansen, 2024, p. 4, Langer & König, 2018, p. 1, Bernstein & Nolan, in print, p. 8). Cases where humanisation of somehting fails usually refer to cases when AI bots are designed to closely mimic human behaviour, but the imitation is imperfect or unnatural. This creates an unsettling effect, often known as the “uncanny valley” (this concept was first introduced by Masahiro Mori in 1970), where people feel creeped out by AI that seems almost human but not quite right.
4) *Lack of transparency* (Watson, 2020, p. 8; Tene & Polonetsky, 2015, p. 94; Langer & König, 2018, p. 8). A lack of transparency is present when one’s intentions or actions are obscured or hidden. For example, when a client does not know how their private data is used in order to generate personalised content.
5) *Low controllability* (Langer & König, 2018, p. 9). An example of low controllability is when an internet user tries their hardest for their data not to be used for personalised advertising, but advertisers somehow still obtain information about them and try to push personalised content (for example, a friend’s data can be used to recommend similar products to another friend).

6) *Violations of social norms* (Tene & Polonetsky, 2015, p. 61; Langer & König, 2018, p. 2; Bernstein & Nolan, in print, p. 17; Herder & Zhang, 2019, p. 2). It is something that comes out of the ordinary way of things for a particular community. For example, an outstanding style of clothing, excessive staring, inappropriate behaviour, or violation of social interaction patterns, like telling one's health issues to a complete random stranger.

Given the wide variety of potential causes and features identified in the literature, one might be tempted to conclude that creepiness is too vague or subjective to study systematically. If creepiness were entirely subjective, we could offer no meaningful guidance to banks, as any recommendation might be perceived as creepy by someone. However, if people demonstrate consistent judgments about what counts as creepy, if there exists *some degree* of intersubjective agreement, this would justify our search for identifiable causes of creepiness in personalised banking recommendations. Only in this latter case could we develop practical improvements that reduce creepiness systematically.

As research shows, individuals not only can experience this sensation of creepiness themselves but can also reliably recognise when others feel it (Fischer & Fredericks, 2020; Bernstein & Nolan, in press, p. 16). This shared capacity is important for this thesis because it proves that there are observable patterns in what people can perceive as boundary violations, which gives us, as independent observers, a chance to evaluate the creepiness of situations. So, while academic definitions and identifications of what factors cause creepiness may vary, in practice, there exists some degree of consistency in socially recognisable triggers that make certain recommendations feel creepy.[6] However, what we currently have at hand is an open-ended list of sufficient conditions for creepiness that, in one way or another, contribute to making something creepy. The key question, then, is whether and how this list can be meaningfully systematised.

3.2 Why Creepiness Resists Simple Definitions: A Prototype-Based View

In the previous subsection, we found several candidates for the causes of creepiness (such as the uncertainty regarding a potential threat, moral insensitivity, failed humanisation, etc.), but at the same time we found that, although there is some partial agreement, there is no strong consensus among researchers on the core attributes. In this

[6] So, creepiness is not entirely subjective, as there are some shared patterns in what people call creepy. However, we should keep in mind that it is not entirely objective either, because, as we will show in Section 4.2., some cases of creepiness may vary across communities.

subsection, we are going to make use of this discovery by suggesting that it is possible to unify previous research under one theory: the prototype theory.

Strong categories can be extremely useful for determining a great many things. For example, whether a woman is pregnant or not is a binary categorisation, as she is either pregnant or not. However, not all categories are so simple. Sometimes, some classifications require more than one attribute to be met. For example, to be considered a bachelor, one must check two boxes: being male and not married. Satisfying only one condition is not enough to make a person a bachelor. An unmarried woman is not a bachelor (she is a bachelorette, but that's a different category), and a married man is not a bachelor. So in some cases, several criteria must be met before something can be clearly defined.

Creepiness, however, is far more complicated than this. There are no fixed boxes that one must necessarily tick for something to be considered creepy. As we can see from the list presented, based on the literature on creepiness, creepiness can be caused by a wide range of factors, and different combinations of different features can result in creepiness.

We cannot analyse creepiness by using the Aristotelian, or, in other words, "classical", theory of necessary and sufficient conditions either (Van der Auwera & Gast 2012, pp. 170-171), which follows the intuition expressed by Bernstein and Nolan (in press, p. 8). A necessary condition is something that must be true for something else to be true. For example, being a mammal is a necessary condition for being a dog. Every dog must be a mammal. But just being a mammal isn't sufficient to make something a dog; there are other animals that are mammals too. A sufficient condition is something that, if it's true, guarantees that something else is true. For instance, being a golden retriever is a sufficient condition for being a dog. If something is a golden retriever, then it must be a dog. But at the same time, being a golden retriever is not necessary for something to be a dog, because there are many other breeds of dogs besides golden retrievers.

Creepiness doesn't work in the same way. We can come up with a lot of things which are creepy, like spiders, clowns, dark streets, some personalised recommendations, etc., but we cannot find just one common feature that is going to be necessarily present in all of them. Moreover, they do not always share the same set of features either, as each instance can be creepy in its own way. An interesting thing about creepiness is exactly that it cannot be tracked or captured using the classical model of necessary and sufficient conditions.

Bernstein and Nolan argue that creepiness is response-dependent. Creepiness arises as a reaction in the mind of the person experiencing it. In other words, something is only creepy because someone perceives it that way. However, authors note that while creepiness depends on personal responses, there are still common patterns in triggers of creepiness in many people. (Bernstein & Nolan, in press, p. 9). They state:

> Creepiness characterizes a wide range of situations and objects, not just those that suggest danger or a violation of social norms. This wide range can be unified in a response-dependent way, by reference to the characteristic response of being creeped out, though since the connection between the relevant reaction and creepy things is not straightforward, we have suggested that spelling out the exact connection calls for a process of bootstrapped inquiry, mutually illuminating both being creeped out and creepiness itself. (Bernstein & Nolan, in press, p. 17)

The authors recognise that the connection between something creepy and the resulting feeling of being creeped out is complex. Importantly, however, their strategy does not claim to provide a final, conclusive theory of creepiness. Rather, it points us toward a direction that is going to be for future research: studying the connection between emotional responses and patterns in features that induce creepiness. This is precisely the starting point from which we will proceed to defend the best way to define creepiness.

To sum up, we have examined various ways of categorising things, for example, through binary distinctions, multiple necessary conditions, and the classical theory of necessary and sufficient conditions, but we have shown that none of these approaches effectively capture the fluid nature of creepiness. Unlike these rigid classifications, creepiness cannot be reduced to a fixed set of shared features. Creepiness arises from a wide range of context-dependent factors. That is why I propose an alternative way of understanding creepiness, where different combinations of features can elicit the same reaction, which is being creeped out. This approach allows us to account for the diversity of cases that people find creepy and, at the same time, account for the recurring patterns that contribute to the experience.

The way forward for our analysis will be to adopt a prototype theory approach. It offers flexibility for understanding categories not as rigidly defined by strict boundaries, but as networks of overlapping features (Van der Auwera & Gast 2012, pp. 174-175). However, it’s important to note that prototype theory can be represented in two distinct ways, and we must assess which version best suits our investigation of creepiness in recommendations.

The first approach is based on the identification of a single, most typical member of a category, the prototype, against which all other members are compared. In this model, category membership is determined by how similar a particular thing is to a central example. While this works well for well-defined concepts, it does not capture the variability and fluid nature of creepiness. There is no single, ideal example of a “creepy thing," since creepiness arises from diverse combinations of factors (such as uncertainty regarding potential threat, moral insensitivity, or failed humanisation, etc.). Thus, this version of prototype theory will not be the best choice for our purposes.

Instead, we will use a second model: the category membership prototype approach. In this model, categories are defined not by a central prototype, but by a set of frequently occurring features that function as sufficient conditions for creepiness. Each feature can independently cause creepiness, although none of them is strictly necessary. The most typical members of the category emerge from the intersection of these features, making them representative instances of a category. This model aligns perfectly with our understanding of creepiness, since different combinations of factors can independently or jointly induce a feeling of creepiness, with creepiness intensifying as more sufficient conditions converge. From this point onward, whenever we refer to "prototype theory," we will exclusively mean this second, feature-based model.

This idea can be imagined as a map of creepiness: At the centre of this map, we find those recommendations that are most typically creepy, which are the cases where several creepiness-inducing features (each sufficient on its own) intersect. These are the recommendations that people are most likely to agree are creepy without hesitation. As we move away from this centre, cases become progressively less typical, exhibiting fewer sufficient conditions. In between these extremes, there is a spectrum ranging from the most typical cases, where social consensus is strong, to borderline or ambiguous cases where judgements may vary between individuals. The borderline cases have only one weak sufficient condition present, making creepiness more subjective and context-dependent. These cases show that something may feel slightly creepy to one person and not creepy at all to another, depending on how strongly the observed features align with the common intersections within the category.

More abstractly, creepiness is like beauty, in some sense. By making this comparison, we see that although in some cases people may disagree about whether something is creepy or beautiful, there is a clear pattern, even if it is unclear what exactly contributes to this pattern, that shows that some cases are more agreeably creepy and more agreeably beautiful. There may be cases in which two people disagree about whether something is beautiful or not, but such disagreements are not consistently frequent across all cases of beautiful things. Some paintings, such as those depicting picturesque views, are more universally regarded as beautiful, whereas others, like some abstract art, are less so. In a similar way, something that is creepy to one person may not be considered creepy by another, but such controversies arise more frequently in cases that lack multiple sufficient conditions of creepy recommendations. The more typical a case of creepiness is (i.e., the more sufficient conditions like lack of transparency or low controllability are present), the more universally it will be perceived as creepy.

After defending the best way to look at and categorise something as creepy, we can now define creepiness by using both the insights from the previous subsection, where researchers mentioned certain potential causes of creepiness, and the insights from this subsection, where we established the best method for identifying something as creepy. Creepiness can be defined as an emotional response to different combinations of sufficient conditions, such as an uncertainty regarding a potential threat, moral insensitivity, failed humanisation, lack of transparency, low controllability, and violations of social norms, among others. These factors can independently or jointly induce a feeling of creepiness, with creepiness intensifying as more sufficient conditions converge.

Up to now, we first defined good personalised recommendations in the previous section and defined creepiness in general in this section. Even though the definition of creepiness is useful, however, if we want to make recommendations to the bank on how to enhance their personalised recommendations so that they will not be creepy, a general definition of creepiness would be too wide to give a good understanding of what should be changed in the recommendations that banks give in order to make them free from creepiness.

For example, failed humanisation as a sufficient condition does not explain the creepiness in the cases we are interested in, which are legally and technically sound personalised recommendations that elicit reactions like "it's none of the bank's business". After all, personalised recommendations in the banking context usually do not have to involve anything related to human-like bots, as, for example, email recommendations are actually pre-written by humans, yet can still cause creepiness, like the recommendation in Peter's Divorce. There must be another, more relevant sufficient condition that better explains this phenomenon.

So, what we need to do at this point is to figure out which factor makes personalised recommendations in the banking context creepy. More precisely, we need to find out which specific sufficient condition contributes to making some personalised recommendations creepy. And this is what I am going to do in Subsection 4.2 after combining the definition of good personalised recommendation (Section 2) and the definition of creepiness (Subsection 3.2).

4. Creepy Personalised Recommendations

To summarise our progress so far: we have defined what a good personalised recommendation should be like and what creepiness is. In this section, we aim to answer three questions:

First, we will return to the issue raised in Section 2: how creepiness makes recommendations not good. Specifically, we will answer whether creepy recommendations fail to align with a client's goal, whether it is not possible to act upon, whether a personalised recommendation is insincere, or whether a personalised recommendation fails to persuade. Knowing that will help us understand, from a technical perspective, what happens with a recommendation that becomes creepy and why it is not good.

Second, in Subsection 4.2, we will specify which sufficient condition makes personalised recommendations in the banking context creepy, and, therefore, not good. We will direct our focus toward only one sufficient condition from our prototype theory framework that is relevant to the cases of our interest.

Third, in Subsections 4.3 and 4.4, we aim to determine what causes creepiness in personalised recommendations within the banking context specifically. For that, we will turn back to Peter's Divorce, proposing three alternative scenarios where the same recommendation is made under different circumstances that carry a lower risk of being creepy. By analysing what unifies these cases and comparing them to the original Peter case, we will identify a special case of the sufficient condition established in 4.2, giving us a clear understanding of what makes personalised recommendations in banks creepy, in particular.

Knowing answers to these questions will allow us to propose in Section 5 the strategy banks can use to minimise creepiness while keeping recommendations personalised.

4.1 How Creepiness Makes Personalised Recommendations Ineffective

In this subsection, we examine what kind of failure creepiness introduces into personalised recommendations. For that, we should return briefly to the conditions of good personalised recommendations indicated in Section 2. Recall that we defined good personalised recommendation as a communicative act in which one person suggests a specific course of action to another (Habgood-Coote, 2024, p. 13), grounded in a true belief of what would benefit the recipient, aligned with their goals and circumstances, realistically actionable, sincere, and presented in a way that invites meaningful consideration. To have a clear understanding of how creepy personalised recommendations make banks' efforts to increase customer satisfaction and engagement fail, I will assess whether, for each condition of being a good recommendation, this condition is violated in the creepy bank recommendation cases.

To begin, let us assess firstly whether creepiness causes the recommendation to fail in terms of alignment with the client's goals. Returning to Peter's Divorce, the recommendation provided by the bank was well-aligned with a presumed financial goal, which is efficient money management after the divorce. And it was personalised in a way that took into consideration the unique circumstances Peter is in. This evidence demonstrates that the creepiness of some personalised recommendations does not affect how well the recommendation aligns with clients' goals.

Secondly, the recommendation given to Peter appeared actionable, assuming the client had the means to follow it. This suggests that creepiness does not inherently render a recommendation misaligned with the recipient's objectives or impossible to implement. The case proves that even when a suggestion is creepy, it may still logically fit the recipient's needs and remain executable. Therefore, the failure of a creepy recommendation cannot be reduced to a misalignment with goals or an inability to be acted upon.

Thirdly, the fact that a recommendation is creepy does not inherently undermine its sincerity. While it's possible that in some cases, a bank's malicious intentions behind giving certain recommendations may contribute to creepiness, creepiness itself does not necessarily indicate a lack of sincerity in the recommendation. Peter's example confirms this: even completely sincere recommendations, those genuinely intended to benefit the client based on the bank's best understanding of their situation, can still come across as creepy. The emotional discomfort triggered by creepiness exists independently of the recommendation's truthful intention. Thus, we establish that creepiness does not affect the sincerity of a recommendation.

Now we reach the fourth dimension of failure previously mentioned in Section 2: the persuasive power of the recommendation. We cannot always predict exactly how someone will react to what we say. In fact, it is very common for people to react in ways we do not expect at all. Listeners interpret what they hear based on their own experiences, moods, and assumptions, which may not align with what the speaker intended. However, certain factors can influence how likely it is that a message will be understood in the intended way. For example, if there are multiple possible interpretations of what is said, a speaker can include clarifications to avoid misunderstanding. Common everyday phrases such as "don't take this the wrong way," "just to clarify," or "I don't mean to offend you" are used to guide the listener's interpretation and reduce the chances of confusion. These phrases are meant to increase the likelihood that the listener reacts in the way the speaker hopes.

The creepiness of recommendations, on the other hand, works in the exact opposite way. Instead of making the client more likely to follow the suggestion, it tends to push them away, making it less likely that they will accept the recommendation as the recommender intended. While it is true that even a well-thought-out, non-creepy recommendation will not always be received in the way we want, by learning what makes something creepy, we can take steps to prevent the risk of negative reactions. This distinction is crucial for understanding why some recommendations fail. In this way, while we cannot guarantee that every non-creepy personalised recommendation will be received positively, we can significantly lower the chances that it will have the unintended side-effect of making the client feel creeped out.

To conclude, unlike misunderstandings, which can often be corrected, creepiness leaves a lingering sense of unease that can negatively affect future interactions. If personalisation is meant to serve the client's best interests, then a creepy recommendation, despite being technically and legally sound, betrays that purpose by violating a certain condition that makes interactions feel appropriate. This creates not just short-term rejection of the recommendation, but long-term harm to trust. This condition that creates a negative experience for clients is what we must now identify.

Having established how creepiness makes personalised recommendations not good (Section 4.1), we must now pinpoint which sufficient condition from our prototype theory of creepiness explains this lack of persuasiveness in banking contexts specifically. This is the central question that will be addressed in the next subsection.

4.2 Relevant Sufficient Condition for Creepiness in Personalised Recommendations in Banking

As shown in the introduction, creepiness caused by personalised recommendations in banks cannot be adequately explained by shortcomings in data quality or analytical processes, nor can it be reduced to mere non-compliance with legal standards. This section seeks to determine which sufficient condition from our prototype theory framework best explains creepiness in these specific banking cases.

First, we will examine a view that defends the idea that a social norm violation and, a bit more precisely, a social norm of how banks should give personalised recommendations, operates as the primary sufficient condition in our cases of interest. Second, an ethical/moral view will be presented. This view will argue that it is an ethical norm violation that may function as a sufficient condition for creepiness in personalised banking recommendations.

a) Creepiness In Personalised Recommendations In Banking Context Is A Social Norm Violation

At first glance, claiming that banks have morally wronged Peter by offering creepy personalised recommendations seems to be a too strong claim. Instead, the more modest interpretation is that the creepiness stems from a misalignment between the recommendation and Peter's expectations of what is appropriate for a bank to do, just as we established in Subsection 4.1. This misalignment does not necessarily involve harm, but rather a breach of context-specific social norms. The discomfort arises because a bank's actions deviate from what a client assumes to be "normal" behaviour for financial institutions.

As Brennan et al. explain, social norms are "practice-dependent" (Brennan et al., 2013, pp. 71-72). This means that they derive their force not from universal principles, but from the everyday practices and expectations of a particular group. In other words, they reflect how things are typically done within a given context (Brennan et al., 2013, pp. 67-68). Rather than existing independently of human activity, social norms are deeply tied to specific communities and their habitual ways of interacting. They evolve alongside social practices, gaining strength from repetition and collective recognition. A helpful illustration is the social rule that discourages telling a dirty joke at a funeral. It may not be illegal or unethical, but it disrupts behavioural standards that the group implicitly upholds during such ceremonies (Brennan et al., 2013, p. 61). Similarly, and crucially for our argument, certain personalised recommendations by banks may feel creepy because they breach clients' implicit expectations of how banks should behave. Subsection 4.4 will highlight which specific expectations were breached in Peter's case.

Creepiness in personalised recommendations in the banking context, then, can be understood as an alarm that signals that a social norm of how banks should give personalised recommendations has been violated. Social norms structure our sense of what is normal and create frameworks of accountability. They tell us, often without articulating explicitly, what kinds of actions are fitting or unfitting within particular social contexts.

When institutions behave in ways that conflict with these expectations, we experience discomfort (Brennan et al., 2013, pp. 87-88), which can make clients dismiss recommendations altogether.

To further substantiate this claim, we can draw an analogy with cross-cultural differences in how creepiness is perceived and what counts as a norm. As we mentioned in Section 3.2, the perception of creepiness, even if not random, allows a variation from person to person to some extent (Bernstein & Nolan, in press, p. 16). Moreover, research shows that creepiness can also differ to some extent from culture to culture. For example, a study found that Americans are more creeped out by increasingly human-like robots compared to people from Japan (Castelo & Sarvary, 2022, p. 1871). The "German stare", a phenomenon that describes cases where Germans can't help but stare at other people without even noticing it, is normal among the citizens, while for visitors, the amount of unwanted attention crosses the line into a creepy zone (Bradley, 2006). These cases show that what can breach a certain norm and become creepy depends, at least partially, on what a certain community counts as being within a certain commonly shared norm.

Although the cases mentioned do not directly explain why the cases like Peter's Divorce are creepy, as there is no specific study on cultural differences in how bank recommendations are perceived, it nonetheless justifies us in drawing a parallel: it suggests that the way Peter, and we as readers, interpret the recommendation may differ from other social groups, where standards of "normal" communication with financial institutions may vary significantly. Therefore, we can conclude that Peter's discomfort is fundamentally explained by social practices, or "how things are done here" (Brennan et al., 2013, p. 67), specifically the breach of social expectations about how banks should deliver personalised recommendations. So, we should focus on the violation of social norms as a sufficient condition that will explain why some personalised recommendations in the banking context are creepy.

However, violations of social norms about appropriate client interactions and about how banks should give recommendations remain too broad to be used as a guide to what exact changes banks should make. Social norm violation serves only as an umbrella covering more specific instances of social norm violations. These instances are special cases of social norm violation, and they are what we want to discover in Subsection 4.4. It is this specific feature that will give banks the exact target they need to adjust to optimise their recommendations.

This social norms interpretation offers one compelling explanation for what norms creepy personalised recommendations violate. However, there exists another potential interpretation that focuses not on violated social norms, but on the ethical norms. Let us now turn to this alternative perspective.

b) Creepiness In Personalised Recommendations In Banking Context Is An Ethical Norm Violation

One may argue, however, that reducing Peter's Divorce creepiness to only a social issue risks missing something important.[7] While social explanations account for why the recommendation felt out of place, they may not capture the full moral weight of Peter's negative experience. To frame his discomfort as "just" a response to a social norm breach might seem to delegitimise the negative experience. There is a risk, in other words, that interpreting Peter's case too narrowly as a social norm violation may obscure what could also be an ethical violation.

Moral norms, in contrast to social ones, are "practice-independent" (Brennan et al., 2013, p. 72). This means that, for example, the norm against murder is not justified by people disapproving of it, but by constituting unjustifiable harm to another human being. However, the fact that a recommendation violates a social norm does not make it impossible for it to also constitute a violation of a moral norm, especially if the recommendation causes distress, negative emotions, undermines dignity, or respect. To say the experience of receiving a personalised recommendation from a bank was creepy, according to this interpretation, is not to say it was just something unusual, uncomfortable and harmless.

Several ethical and moral theories would defend the view that creepiness caused by personalised recommendations in a bank is an ethical norm violation. One of them is ethical relativism. Very briefly, ethical relativism would exactly say that what is right or wrong depends on what one's culture believes is right or wrong (Gowans, 2021). If Peter's cultural or social context evaluates the bank's recommendation wrong, then relativist ethics would allow us to treat that reaction as morally meaningful. Thus, in this way, Peter's negative emotional response is part of what makes the act ethically problematic.

Hedonistic utilitarianism would also argue that Peter's concern is not just social, but is, in an important way, ethical. Shortly, hedonistic utilitarians would say that actions are morally right or wrong depending on the pleasure or pain they produce (Weijers, n.d.). In

[7] I owe thanks for pointing this out to Professor Francesco Orsi.

this view, the fact that the bank's recommendation triggered a negative emotional experience is ethically significant. This means that the outcome of the bank's action was morally questionable, especially if a similar objective, such as suggesting a financial product, could have been achieved without causing one to feel creeped out.

Nissenbaum's theory of contextual integrity would also explain why Peter's experience of creepiness signals an ethical violation. Contextual integrity tells us that keeping something private does not equal keeping something a secret, but rather ensuring that information flows appropriately within specific social contexts (Nissenbaum, 2004, p. 119). According to Nissenbaum, contextual integrity is harmed when either norms of appropriateness or norms of distribution are violated (2004, p. 138). Here, the bank's actions disrupted the conventional norms of the banking context, and Peter's reaction of creepiness reflects a broader social sense that the information flow was inappropriate, making the action ethically problematic.

According to this view, to treat Peter's reaction as ethically irrelevant would be to adopt an unduly narrow view of harm. Reducing the creepiness he experienced to a mere matter of social customs risks downplaying the genuine impact of his negative emotional response. Even if the bank's actions did not breach any legal boundaries, they may still represent a failure to uphold ethical norms.

c) Discussion

Now, as we have two perspectives at hand, we can evaluate whether the sufficient condition that creepy personalised recommendations in the banking context is either a social or an ethical norm violation. The first interpretation frames creepiness as a violation of some social norm where, broadly speaking, creepiness stems from a mismatch between how clients expect banks to give personalised recommendations and what they actually encounter. The second interpretation treats creepiness as an ethical norm violation where creepiness stems from some moral standards that some personalised recommendations in banks fail to adhere to.

Firstly, let us discuss the ethical view. As various ethical approaches demonstrate, there are compelling reasons to treat creepiness as an ethical problem. But not any ethical theory will be relevant to our cases of interest. What specific kind of ethics plays a role in personalised recommendations in the banking context is specifically data ethics.

Main questions of data ethics revolve around transparency, agency, consent, etc., yet not all ethical theories mentioned are equally equipped to address these concerns. For example, hedonistic utilitarianism and ethical relativism might acknowledge creepiness as morally relevant, however, these theories operate at a high level of abstraction, evaluating general outcomes or cultural norms rather than providing tools to evaluate institutional data practices. So, hedonistic utilitarianism and ethical relativism fall outside of scope of theories are able to answer standard questions that are asked in data ethics.

In contrast, Helen Nissenbaum's theory of contextual integrity does engage directly with data ethics. So, contextual integrity and similar kinds of data ethics theories could potentially explain why certain personalised recommendations are creepy. However, there is an important gap in the literature. Although researchers focused on issues like contextual integrity, which helps to explain why certain data flows feel inappropriate, they have not focused on the issue of creepiness. So, perhaps, it could be an ethical issue, but at this point, it has not been addressed in current data ethics discussions on creepiness.

Not being able to fully address creepiness as a violation of some ethical norms is not a problem for the thesis, since the alternative we have is social norm violation as an explanation. Given this gap in data ethics research, this thesis will adopt a view that creepiness in banking recommendations is best understood as a violation of social norms, slightly more precisely, the violation of a social norm of how banks should give their recommendations to clients. This stance is well-supported by the evidence, as the analysis demonstrates that Peter's discomfort arises from a breach of context-specific expectations. The cross-cultural variability in creepiness perceptions further underscores that the phenomenon is rooted in socially contingent expectations. Importantly, by addressing social norm violations, the argument indirectly resolves the ethical concerns raised in part b), as we will lower the risk of creepiness or, in the best case, eliminate it altogether. So, in this thesis, we are going to focus on social norm violation as a sufficient condition.

With this conclusion in place, we can clarify the relationship between this view and the broader framework proposed by Bernstein & Nolan (in press), as discussed in Section 3.2. They argue that creepiness is a wide-ranging affective response that characterises diverse situations and objects, not all of which involve danger or social norm violations. At first glance, our account may seem to contradict theirs, since I claim that some personalised recommendations in the banking context are creepy because they breach some contextually specific social norm governing how banks should interact with clients. However, this is not a contradiction but a difference in scope. Bernstein & Nolan aim to describe creepiness as a general phenomenon, which includes cases as varied as baby spiders and sexual misconduct. By contrast, our concern is much narrower, as we are only attempting to characterise what makes personalised recommendations from banks creepy.

Now that we have established that creepiness in banking recommendations stems from violating some social norm, we have a clearer direction in what kind of things banks should change in order to prevent creepiness. However, this framing remains too broad to be useful. Banks cannot eliminate creepiness simply by being told to "avoid some social norm violations", as such advice is too vague to be actionable.

To offer a practical solution, we must identify exactly which social norm is being violated. What specific feature makes certain recommendations feel inappropriate? In the next section, we will analyse variations of Peter's case to isolate the precise cause of creepiness.

4.3 Examples of Personalised Recommendations Done Right

In this subsection, we will examine three alternative cases that serve as counterparts to the original Peter example. It is important to note that this subsection will not yet provide definitive answers or conclusions. Rather, its purpose is to establish a foundation for analysis that we will refer back to when addressing the first main research question in the subsequent subsection.

To better understand what causes creepiness in personalised recommendations in the banking context, we will adopt a step-by-step approach. First, we will examine cases where the risk of creepiness is significantly reduced or entirely absent. This will allow us to identify features that protect recommendations from being perceived as creepy. In Subsection 4.4, we will then analyse what these cases have in common, aiming to uncover the underlying principle that helps to prevent creepiness from arising.

To begin, let us consider three counterfactual scenarios in which the same recommendation, offering personalised financial counselling to recently divorced individuals, along with the suggestion to open an individual account, is presented, but the risk of creepiness is significantly lower compared to the original Peter's Divorce:

Case 1: Peter's Divorce - Test

Peter visits a well-known personal finance website that offers a diagnostic tool to help users understand their current financial situation. The tool asks about recent life events, types of income, current accounts, future goals, etc.. Peter fills it out honestly, mentioning things like a recent change in living arrangements and uncertainty about future financial plans. At the end of the test, the system generates a personalised recommendation tailored to his profile: it suggests financial counselling for individuals who have recently divorced, and recommends opening a separate account to manage finances independently.

Here, the recommendation is clearly framed as a response to Peter's voluntary and structured self-disclosure. He is aware that the system is drawing conclusions from what he has chosen to share.

Case 2: Peter's Divorce - Request

One day, Peter decides to book an appointment with a financial advisor at a local firm. During the meeting, he explains that he's recently gone through a divorce, is now living separately, and wants to make sure he's handling his finances responsibly moving forward. The advisor listens, takes notes, and assures Peter that they'll follow up with resources. A few days later, Peter receives a personalised message suggesting services aimed at recently divorced individuals and a recommendation to open a separate bank account.

In this situation, Peter initiates the interaction, openly shares the relevant personal information, and expects personalised advice in return. The recommendation feels appropriate and is, more importantly, expected as it is a natural response to requesting it.

Case 3: Peter's Divorce - Close Relationship

Peter is having dinner with his older sister, someone he has always trusted and was very close to. Over the course of the evening, he talks about his recent separation, the challenges of living alone for the first time in years, and how he's trying to figure out his financial future. A few days later, his sister sends him a message saying she's been thinking about him and suggests he consider speaking to a financial counsellor who specialises in post-divorce planning, and perhaps open a separate bank account to manage things more easily.

Despite the fact that Peter didn't ask for help or advice, the suggestion comes from a deeply trusted source. There was no need for explicit requests, as the sister's status as a close person made such behaviour expected. Once his sister understood that her brother was going through some hardship, she made a recommendation.

Although these three cases differ significantly from each other in context and the nature of interaction, they all share a notable and important characteristic for the aim of this thesis: the risk of creepiness is significantly reduced or even entirely absent. Despite their diversity, something about the way each recommendation is delivered makes it feel acceptable, appropriate, and even welcome. This brings us to the central question: what is it that something that unites these cases and protects them from the creepiness we observed in the original "Peter" scenario? This is the question we are going to explore in the next subsection.

4.4 Why Some Personalised Recommendations in Banks Feel Creepy

Now we are reaching one of the most important parts of the thesis. This section will answer one of the key questions: What is the cause of creepiness in personalised banking recommendations? To do that, we will first employ a contrastive analysis of cases proposed in Section 4.3, identify what unites them, and compare them to the paradigmatic example of a creepy personalised recommendation given by a bank, which is Peter's Divorce. Then we would propose what could be the factor that is present in these three cases but is absent in the original Peter's Divorce case, making the difference in the risk of creepiness. We are going to consider 3 explanations of what factor makes a difference: emotional closeness, explicit request, and whether a personalised recommendation was based on sufficiently grounded common knowledge (what is meant by that is going to be explained later in this subsection). Finding out what the problem is, as was mentioned, is half of the solution. So, once we know what causes creepiness, in Section 5, we will be able to answer the second key question, which is: How can banks maintain high personalisation and creepiness-free recommendations?

To start, one possible explanation is that, in all three cases, Peter appears to have a certain degree of trust or emotional closeness to the recommender, making the suggestion feel less creepy. This is most obviously true in Case 3, where the recommendation comes from a close family member, someone with whom Peter has a strong, emotionally supportive relationship. However, this explanation falls short when we consider Cases 1 and 2. In Case 1, the recommendation is generated by an impersonal online tool; in Case 2, it comes from a professional advisor whom Peter has likely never met before. Neither interaction involves a close relationship, yet both still avoid the creepiness associated with the original Peter scenario. Thus, emotional closeness might contribute to a feeling of comfort, but it does not answer why all three cases have a lower risk of creepiness.

A second possibility is that in these cases, Peter has explicitly requested a recommendation, which makes it appropriate to respond with a personalised recommendation. In Case 2, this is clear: Peter initiates a meeting with a financial advisor and openly shares his situation, expecting a personalised recommendation in return.

However, this explanation also has limits. Case 3 presents a situation where Peter does not explicitly request help. His sister listens, reflects, and later makes an unsolicited recommendation, and yet, the recommendation still does not feel creepy. So, while voluntary disclosure and request for help can reduce creepiness, they too are not a universal requirement.

We now turn to the third and most compelling possibility: what unites all three non-creepy cases is that the recommendations Peter receives are not experienced as sudden. They do not appear to "pop out of nowhere" but rather emerge as the natural result of a preceding interaction in the given context. It seems that both the recommender and the recipient share a belief that the recommendation given is warranted and is appropriate. Specifically, the recommender knows Peter wants a solution to his specific problem, and Peter knows the recommender recognises this need. Peter can make sense of how the recommendation came about because it naturally flows from a recognisable pattern of preceding interactions. In all the mentioned contexts, there is a coherent sequence that Peter recognises, where his disclosure or behaviour explicitly or implicitly signals his needs, creating *common knowledge* (Vanderschraaf & Sillari, 2023) between him and a bank that motivates the personalised recommendation he later receives.

This sharply contrasts with the original case, where Peter receives a personalised recommendation from his bank without any easily recognisable interaction that would account for it. In the original Peter's Divorce, the recommendation was based on a large amount of data that is just too hard to keep track of for Peter, as he may not recall every single interaction with his account that could make it obvious that he is getting divorced. Crucially, while the bank inferred his needs from the data, Peter had no opportunity to confirm or contextualise this inference through his own actions. While he may have previously consented to receive personalised offers, this general consent does not sufficiently *ground* (Clark & Brennan, 1991) the specific recommendation he receives. The recommendation arrives unexpectedly. It is not necessarily unwelcome in content, nor is it necessarily based on illegitimate data, yet it triggers creepiness.

In the last two paragraphs, I used some new terms that have not been sufficiently explained yet. Let us clarify now what we should understand by grounding, common knowledge, and what should be common knowledge in order to make personalised recommendations in the banking context creepiness-free.

First of all, what we can see from these three cases they showed that the recommender and the recipient commonly shared some knowledge that made the subsequent recommendation-giving process feel natural, expected, and, most importantly, not creepy. More explicitly, these three cases shared a 3-level structure that defines common knowledge:

(1) A wants a solution to a specific problem x,

(2) B knows that A wants a solution to x, and

(3) A knows that B knows that A wants a solution to x.

Importantly, if only conditions 1 and 2 are satisfied, we can say that person A and person B only have *mutual knowledge* (Vanderschraaf & Sillari, 2023). They only individually know that A wants a solution to a specific problem x. However, there is no shared common knowledge.

Let us see how this approach works by returning to our three alternative cases and the original Peter's Divorce. What will be shown is that Peter (the client), in these three cases, has a common knowledge of the fact that the recommender knows that he wants a solution to his financial issues connected to divorce (specific problem). In contrast, the original Peter's Divorce will be shown to have only mutual knowledge, which makes the recommendation socially unacceptable, unexpected and creepy.

In Case 1, the interaction is structured by Peter's voluntary engagement with the diagnostic tool. He actively participates by answering questions about his financial situation, explicitly signalling his desire for a specific issue to be resolved and creating a clear, logical sequence that leads to the recommendation. Because Peter knowingly provides this information, he is aware that the system recognises his desire for the solution. In this way, the interaction satisfies all three levels required for common knowledge:

(1) Peter wants a solution to his financial concerns regarding divorce (true)

(2) The system knows Peter wants a solution to his financial concerns regarding divorce (true)

(3) Peter knows the system knows that Peter wants a solution to his financial concerns regarding divorce (true)

This co-construction of knowledge through his active collaboration makes the recommendation meet social expectations of how personalised recommendations should be delivered: based on information both a recommender and a recipient share, expected and warranted, avoiding creepiness.

In Case 2, the common knowledge of Peter's financial goals and worries comes from the interpersonal exchange between Peter and the financial advisor. Peter initiates the meeting, shares his concerns directly, and thus sets the conversational frame within which personalised recommendation is expected. The recommendation, arriving shortly after the meeting, fits into the context that was previously grounded by arriving at the meeting, initiating a dialogue, and mentioning struggles he has. This case also satisfies all three levels required for common knowledge:

(1) Peter wants a solution to his financial concerns regarding divorce (true)

(2) The advisor knows Peter wants a solution to his financial concerns regarding divorce (true)

(3) Peter knows the advisor knows that Peter wants a solution to his financial concerns regarding divorce (true)

The predictability of the recommendation that is caused by Peter's explicit request ensures that all parties share the same common knowledge, which makes the recommendation given feel contextually appropriate.

Finally, in Case 3, the common knowledge is achieved by Peter's casual disclosure of his worries, his sister's attentive listening, and their mutual recognition of each other's knowledge and understanding that Peter wants a solution to his problem. The three-level knowledge is easily seen, again, through this structure:

(1) Peter wants a solution to his financial concerns regarding divorce (true)

(2) The sister knows Peter wants a solution to his financial concerns regarding divorce (true)

(3) Peter knows his sister knows that Peter wants a solution to his financial concerns regarding divorce (true)

In the original Peter's Divorce case, the interaction lacks the crucial third component of common knowledge: while both Peter and the bank independently knew about Peter's divorce-related financial needs, Peter had no way of knowing that the bank was aware of this specific issue. This created a situation of mere mutual knowledge where both parties possessed relevant information, but they didn't share this information. To illustrate it:

(1) Peter wants a solution to his financial concerns regarding divorce (true)

(2) The bank knows Peter wants a solution to his financial concerns regarding divorce (true)

(3) Peter knows the bank knows that Peter wants a solution to his financial concerns regarding divorce (false)

Even if Peter had consented to personalised recommendations in general, this broad consent did not translate into shared common knowledge for this specific suggestion. The bank's inference about Peter's divorce was based on a lot of data that Peter could not keep track of. As a result, the recommendation appeared to materialise out of nowhere from Peter's perspective, violating expectations of how personalised recommendations in a banking context should be given. To resolve such situations, banks must ensure that common knowledge is established before delivering personalised recommendations. This means creating contexts where a client can recognise that banks recognise their desire for a solution.

However, even though common knowledge is important to making sure that the recommendation does not violate the social norm of appropriate recommendation-giving, which is not to give personalised recommendations based on information that only one party possesses, what we also must focus on is how this common knowledge is achieved. For example, what if a bank decides to compensate for the lack of common knowledge by making a recommendation fully transparent? Such a strategy would require banks to disclose precisely what data was used, how it was analysed, and which data analysis mechanisms were applied. Imagine a situation where a bank, in a single email, lists everything it knows or could potentially infer about a client. Such transparency, motivated by good intentions, however, may, on the contrary, creep clients out. In some cases, this kind of blunt explicitness can even increase the sense of creepiness (Torkamaan et al., 2019, p. 1). Even if the client has consented to personalised recommendations and is aware that their data is being used, such exhaustive and blunt transparency might evoke a feeling of being exposed or closely surveilled. This shows that in order to benefit from the low risk of creepiness that the establishment of common knowledge provides, it is important to ground this knowledge in a proper way.

Clark and Brennan (1991, p. 128) propose to define grounding as a "collective process by which the participants reach this [common knowledge]".[8] What should be highlighted in this definition is that common knowledge or common belief is achieved collectively, and not by one-sided interactions, like in the case where a bank would bluntly list everything it knows and could potentially infer from the data they have about its client. It is exactly this collaboration, the active participation of both a bank and a client, that can ground the recommendation and guarantee that the recommendation given will be welcomed.

[8] Clark and Brennan's original phrasing as "reach this mutual belief" was replaced with "reach this common knowledge", as they did not make a distinction between "mutual" and "common" in their language. It is, however, clear from their text that what they meant is that this belief is collective (common) rather than individual (mutual). For our thesis, however, this distinction is crucial.

Grounding, unlike transparency, rather than overloading the client with data, creates a context through interaction that implicitly secures, so to speak, micro-consent and reaffirms the general consent that was given on the topic of receiving personalised offers. This establishes that the client is both prepared and willing to receive a particular kind of personalised information.

The way in which a certain interaction is grounded differs depending on the medium of communication (Clark & Brennan, 1991, p. 128). Normally, in face-to-face interactions, the conversation can be grounded by verbal cues like “mhm”, repeating what a person heard to confirm, noticing that a person stutters, asking clarificatory questions, etc. Or also visual cues like nodding, facial expressions, etc. In communication channels like emails, phone notifications, or tests, etc., that banks often use while giving recommendations, however, the mentioned cues from face-to-face interaction are impossible to use, as this kind of interaction is asynchronous. So, if banks want to continue giving personalised recommendations via these channels, they need to adjust their strategy to first ground the interaction. Examples of how to do it will be shown in Section 5.

The three alternative cases allow Peter to either actively contribute to establishing all three levels of mutual understanding or easily recognise how they were achieved. These interactions share a common structure: the recommendation feels like a response, rather than a surprise. It is this grounding that gives the interaction a socially acceptable character. In Case 1, Peter voluntarily took a diagnostic test, explicitly signalling his needs and enabling the system to generate a recommendation he could trace back to his input. In Case 2, he directly requested help from a financial advisor, creating a clear expectation for a tailored response. In Case 3, his close relationship with his sister, built on trust and past interactions, provided an implicit warrant for her unsolicited yet appropriate advice. In contrast, the original Peter case lacked grounded common knowledge, making the personalised recommendation feel sudden, unanchored, and ultimately creepy.

At this point, we are now in the position to answer the first question we aimed to answer in this thesis. Creepiness in personalised recommendations, as an emotional response, comes from personalised recommendations in the banking context, lacking the three-level common knowledge conditions that are progressively established through grounded interaction.

To unify our discoveries with the prototype theory that was described in Section 3.2, there are a few remarks needed to make sure that the stance of the cause of creepiness in some personalised recommendations in banks becomes clear. Prototype theory of creepiness says that the most typical instances of creepiness stem from the combination of overlapping creepiness-inducing sufficient conditions, and in the banking context, it is the lack of grounded common knowledge that becomes a particularly important feature contributing to the experience of feeling creeped out. This means that creepiness caused by banks doesn't represent the most typical case of creepiness. Being creeped out by a bank giving some personalised recommendations is creepy, but not as creepy as cases that involve a combination of a large number of creepiness-inducing factors. Rather, it represents a milder but still legitimate form of creepiness that fits the broader, response-dependent understanding of creepiness defended earlier. Thus, a personalised recommendation that was given without proper grounded common knowledge can rightly be called creepy, though of a socially situated and less intense kind compared to, for example, suspecting that a serial killer walks behind a person in one's neighbourhood.

Moreover, to make it clear, when we say that lack of grounded common knowledge makes some personalised recommendations creepy, we do not claim that other creepiness-inducing factors do not play any role at all. Rather, we suggest that in the cases discussed, it is the absence of grounded common knowledge that emerges as the most significant feature, while other factors still contribute to the overall feeling of unease, though more supplementary. For example, clients may also realise that banks act in their own financial interests, which do not always align with those of the client. The recommendation then becomes manipulative, where the client must remain alert to defend their own interests. Banks have a significant influence over the client's financial behaviour, which can make clients concerned that they are not being tricked into some services that go against their best interests. In this way, the uncertainty of potential threats may contribute alongside the lack of grounded common knowledge, which will intensify creepiness.

To lower the risk of creepiness in personalised recommendations is to change the context in which these recommendations are delivered. Rather than reducing the degree of personalisation itself, which would undermine its potential benefits, banks can alter the circumstances surrounding giving personalised recommendations in a way that clients take an active part in generating personalised recommendations. Crucially, this will allow us to reach the main goal of our thesis: provide a strategy to banks that, by reshaping the context rather than the content of recommendations, will allow them to preserve high levels of personalisation while mitigating the social and ethical violations resulting from creepiness. The last question remains: How can banks implement this discovery in practice?

5. Solutions

As we have discussed, creepy recommendations have faults from both the perspective of the recommender and the perspective of the recipient. From the recommender's perspective, the fault of creepy recommendations is that they don't persuade the recipient to consider or act upon the recommendation. From the perspective of the recipient, or, in our context, a client, a recommendation is creepy because it violates some social norm. This justifies the need for recommendations to be adjusted. Building on the prototype theory framework introduced in Section 3.2, we established that creepiness arises due to the presence of a certain sufficient condition, and in the context of personalised recommendations in banking, the primary sufficient condition is the lack of grounded common knowledge between the bank and the client about the client's desire for a solution to a specific issue (as outlined in Section 4.4). With this background in mind, we can now turn to suggesting various strategies banks can use to lower the risk of creepiness in their recommendations.

To illustrate the practical value of this framework, I will demonstrate how it can be used by banks to improve their personalised recommendations. Instead of relying on trial and error while trying to make recommendations less creepy, this framework allows banks to adjust personalised recommendations systematically by having a concrete plan, which is to ensure that the three-level common knowledge structure, which is established through grounded interactions:

(1) The client wants a solution to a specific issue x
(2) The bank knows that the client wants a solution to a specific issue x
(3) The client knows the bank knows that the client wants a solution to a specific issue x

This means creating contexts where the client actively participates in signalling their needs before receiving a recommendation, thereby making the recommendation feel expected and justified rather than sudden and, therefore, creepy. When, in the introduction, it was suggested that Peter might have an immediate reaction of "none of the bank's business," what we need to do is create a situation where the context makes discussing a certain problem a "bank's business". The solution lies in designing micro-interactions that nudge clients to express their willingness to address a specific issue, as proposed in Section 4.4. This approach will transform personalised recommendations into a collaborative process and lower the risk of creepiness.

1) Solution in Practice: A Test

Emails or push notifications are one of the popular ways to deliver recommendations. One big downside of them, however, is that they are essentially a one-sided way of communicating with a client. This means that while banks send out information via this medium, clients do not really have an opportunity to reply. And since grounding necessitates some kind of collaboration and active client involvement in the recommendation-making process, delivering highly personalised recommendations in this way may not be the best option unless the interaction was grounded beforehand in one way or another.

However, we should not completely dismiss the usefulness of these communication mediums. They can still be used for personalised, non-creepy recommendation-making, but a way to do that could be by using them to invite clients to visit another platform via hyperlinks, for example. Once the client visits another platform, like a website with tests designed for personalisation of recommendation-making, the grounding of common knowledge can begin.

One solution that banks can use on this new webpage is to structure the recommendation process through a short consultation-style test, offering clients a controlled way to engage with personalisation. The way it can be executed is that when a bank's system identifies potential opportunities for improvement in a client's financial situation, this would trigger an invitation for a client to take the test. The invitation could be framed as a collaborative opportunity. So, if a client is interested in it, one will be able to express this interest.

The test serves as the crucial grounding mechanism that establishes all three levels of common knowledge required for non-creepy recommendations. As the client answers questions about their financial situation and goals, they become aware of their own needs, satisfying the first level. The bank simultaneously registers these needs through the test responses, achieving the second level. Most importantly, because the client voluntarily provides this information through the structured test, they know that the bank knows their needs, which fulfils the essential third level of common knowledge.

After completing the test, the bank generates recommendations that reference only the information the client explicitly provided. This contextualisation ensures clients can always trace the logic behind recommendations.

By implementing this test-based approach, banks can deliver highly personalised recommendations while avoiding the creepiness that renders them ineffective. Moreover, it preserves the benefits of personalisation (Section 2) by maintaining alignment with client goals, actionability, and sincerity, while ensuring the recommendation is compelling (Subsection 2.4).

2) Solution in Practice: A Chatbot

One more idea for a new recommendation system design proposes the creation of an optional, specialised personalised recommendation service, for example, a conversational chatbot advisor. Clients would have the opportunity to explicitly sign up for using such a tool.

The use of a chatbot that builds a specific recommendation in collaboration with a client tackles one of the biggest problems of the sense of suddenness and unexpected inference that causes creepiness. Instead of a recommendation "popping up" seemingly out of nowhere, the client will be invited in a neutral way, for example: "We've noticed a possible opportunity to strengthen your financial strategy. Would you like to discuss it?" Only if the client says yes would the conversation begin. In the conversation itself, the chatbot would not immediately reveal the detailed inferences the bank has made. Rather, it would nudge the client through a series of context-building questions to share their thoughts about the recommendation they expect to receive. This serves to anchor the interaction in the client's own perspective, allowing the bank's prior data inferences to be either confirmed or adjusted through open dialogue.

This interaction with a chatbot systematically builds a common context. This design ensures that the 3-level knowledge structure, where (1) the client wants a solution to a specific problem, (2) the bank knows this, and (3) the client knows the bank knows, is achieved through active collaboration with the conversational AI assistant. The chatbot guides the client to articulate their needs while simultaneously confirming the bank's inferences. Such a strategy co-constructs the common knowledge and understanding necessary to avoid creepiness.

This conversational approach ensures that grounded common knowledge is firmly established before any recommendation is made. Because the shared context emerges through interaction, and not through, so to speak, the one-sided declaration of a recommendation, when personalised suggestion eventually comes, it will feel naturally entailed.

Another key benefit is that clients stay in full control. They decide whether to start the conversation, how far they want to go with it, and they can stop the interaction at any time without any pressure. This optionality significantly increases the chances that the resulting recommendations will not only be accepted but welcomed. This idea is supported by empirical evidence: “Researchers demonstrated that, paradoxically, if individuals are given more control over the publication of their private information, their privacy concerns decrease and their willingness to publish sensitive information increases” (Tene & Polonetsky, 2015, p. 80). In the context of a chatbot, control translates to the client’s ability to shape the interaction, and a bank’s role becomes one of a collaborative partner.

As demonstrated through these examples, there are many ways to apply the theory of grounding common knowledge in a creative and practical manner. The solutions presented here represent just a few potential approaches and are certainly not an exhaustive list, but they clearly show the kinds of strategies banks can use to address the issue. What is most important, however, is the consistent use of the theory itself: Our analysis proves that if banks are willing to develop various methods based on the principle of creating a meaningful and recognisable context before delivering personalised recommendations, and maintain this grounding throughout the interaction, then properly executed, these efforts should similarly lower the risk of creepiness, or, in the best-case scenario, eliminate it altogether.

Limitations

Returning to the prototype theory of creepiness discussed earlier, it is important to acknowledge that the strategy developed in this thesis addresses only a specific subset of creepy personalised recommendations, the ones that are legally compliant, technically sound, but nevertheless induce reactions such as “How did they know?”, “Why does the bank give me such recommendations?”, and “It is none of their business”. Prototype theory holds that creepiness arises from various features, and therefore, not all instances of creepiness share the same cause. While restricting the interest in particular kinds of cases allows for a targeted and practically useful intervention, it does not claim to exhaust the full range of possibilities by which personalisation may become creepy.

There are, of course, other ways in which personalised recommendations might evoke a feeling of creepiness, even when grounded common knowledge is present. Prototype theory also leaves room for less typical cases. For instance, a recommendation could be creepy not because it emerges without grounding, but because it carries inappropriate or sexually suggestive language. In such scenarios, the grounding strategy proposed here would not be sufficient to eliminate the creepiness, since the source of creepiness lies elsewhere, which is the nature of the recommendation's content rather than its manner of delivery.

Thus, it must be stressed that the purpose of this thesis was not to propose a universal solution to every instance of creepy personalisation. Rather, it aimed to diagnose restricted kinds of cases.

Conclusion

This thesis set out to investigate two primary questions: (1) what makes some personalised recommendations from banks feel creepy, and (2) how banks can maintain highly personalised recommendations while eliminating creepiness. To address these questions, the thesis employed a strategy of first establishing a theoretical framework, examining cases where the risk of creepiness is lower or absent compared to the same recommendation given in different context, and, finally, proposing some practical solutions for banks based on the insights gained, aiming to preserve the benefits of personalisation and eliminate the harms of creepiness.

The analysis revealed that the main cause of creepiness in personalised recommendations is the lack of grounded common knowledge about a client's desire to solve a specific issue. Recommendations that seem to arise suddenly, without a clear or meaningful sequence of interaction, trigger feelings of being creeped out, because without sufficient prior interaction, such straightforward exposure to potentially highly accurate information violates social expectations of how such information should be delivered. It is crucial for smooth communication with clients to make sure that they perceive personalised recommendations within a coherent narrative that allows them to anticipate why a personalised recommendation is delivered to them.

Recognising the cause of creepiness allowed us to propose some solutions that banks can implement aimed at restructuring the delivery and, particularly, the context in which recommendations are given, instead of abandoning personalisation altogether. Introducing some kind of micro-interactions where a client is actively involved in the recommendation-generating process, such as short tests or conversational chatbot consultations, provides the necessary contextual framework that allows clients to feel more comfortable with personalised recommendations they receive. These methods ensure that clients experience recommendations as responses to their own actions, thus reducing the risk of creepiness.

The usefulness of this research extends not only to the particular case that we have discussed but to many others, too. First, because a thorough investigation was completed into what causes creepiness in personalised recommendations in a banking context, other institutions could use the strategy for identifying the cause of creepiness we proposed in their specific circumstances, too. While we focused on recommendations, the same strategy could help improve marketing and general advertising by making them feel less creepy. Second, at the very least, this research shows that studying creepiness is worthwhile, as it proves the effort pays off and can inspire deeper investigations into how to make personalised recommendations more comfortable for clients.

The potential positive side-effect thesis is that it could also help reframe the way companies perceive the issue of social norms in recommendations. By demonstrating that understanding and addressing these norms benefits both banks and clients, rather than merely adding another regulation for banks to comply with, it may shift the perception of this research from being an unnecessary burden to an opportunity for improvement. So, hopefully, this research could spark greater interest in the topic and encourage further research.

Bibliography

Bernstein, S., & Nolan, D. (in press). Creeped out. In U. Kriegel (Ed.), *Oxford studies in philosophy of mind,* 5. Oxford University Press.

Bradley, K. (2006, May 25). The Germanic stare-down: Watcha lookin' at, granny? *Spiegel International*. https://www.spiegel.de/international/the-germanic-stare-down-watcha-lookin-at-granny-a-418068.html

Brennan, G., Eriksson, L., Goodin, R. E., & Southwood, N. (2013). Moral and social norms. *Explaining norms*. Oxford Academic. https://doi.org/10.1093/acprof:oso/9780199654680.003.0004

Castelo, N., & Sarvary, M. (2022). Cross-cultural differences in comfort with humanlike robots. *International Journal of Social Robotics*, 14: 1865-1873. https://doi.org/10.1007/s12369-022-00920-y

Clark, H. H., & Brennan, S. E. (1991). Grounding in communication. In L. B. Resnick, J. M. Levine, & S. D. Teasley (Eds.), *Perspectives on socially shared cognition*, 127-149. American Psychological Association. https://doi.org/10.1037/10096-006

Dean, G. (2025, March). Facebook user sues Meta over targeted ads. *BBC News*. https://www.bbc.com/news/articles/c1en1yjv4dpo

Demetriou, D. (2024). Sexual creepiness [Unpublished manuscript]. Department of Philosophy, University of Minnesota.

Fischer, J., & Fredericks, R. (2020). The creeps as a moral emotion. *Ergo: An Open Access Journal of Philosophy*, 7: 191-217.

Gartner. (2019, March 11). Gartner survey shows brands risk losing 38 percent of customers from poor personalization efforts. *Gartner Newsroom*. Retrieved from: https://www.gartner.com/en/newsroom/press-releases/2019-03-11-gartner-survey-shows-brands-risk-losing-38-percent-of

Ghaffary, S. (2024, August). Amazon's creepy health care integrations are just getting started. *Vox*. https://www.vox.com/technology/369302/amazon-one-medical-pharmacy-prescription-drugs

Goic, M., Rojas, A., & Saavedra, I. (2016). *Investigating the effectiveness of triggered email marketing*. SSRN Electronic Journal. https://doi.org/10.2139/ssrn.3217130

Gowans, C. (2021). Moral relativism. E. N. Zalta (Ed.), The Stanford *Encyclopedia of Philosophy* (Spring 2021 ed.). Stanford University. https://plato.stanford.edu/archives/spr2021/entries/moral-relativism/

Habgood-Coote, J. (2024). Thinking together: Advising as collaborative deliberation. *Philosophers' Imprint*, 24 (1). https://philosophersimprint.org/

Herder, E., & Zhang, B. (2019, September 17). *Unexpected and unpredictable: Factors that make personalized advertisements creepy*. In 23rd International Workshop on Personalization and Recommendation on the Web and Beyond (ABIS '19), Hof, Germany. Association for Computing Machinery. https://doi.org/10.1145/3345002.3349285

HubSpot Staff. (2013, April 25). 15 call-to-action statistics you need to know about to increase your conversion rate. *HubSpot Blog*. https://blog.hubspot.com/marketing/personalized-calls-to-action-convert-better-data

Kjeldgaard-Christiansen, J. (2024). What is creepiness, and what makes ChatGPT creepy? *Leviathan: Interdisciplinary Journal in English*, 10. https://doi.org/10.7146/lev102024144284

Langer, M., & König, C. J. (2018). Introducing and testing the creepiness of situation scale (CrOSS). *Frontiers in Psychology*, 9. https://doi.org/10.3389/fpsyg.2018.02220

Lyn, T. C., Edelman, D. B., & Crook, J. N. (2002). *Credit scoring and its applications*. SIAM.

McAndrew, F. T., & Koehnke, S. S. (2016). On the nature of creepiness. *New Ideas in Psychology*, 43: 10–15. https://doi.org/10.1016/j.newideapsych.2016.03.003

Nissenbaum, H. (2004). Privacy as contextual integrity. *Washington Law Review*, 79(1), 119-158. https://digitalcommons.law.uw.edu/wlr/vol79/iss1/10

OpenAI. (2025). ChatGPT (January 2025 version) [Large language model]. https://chat.openai.com

Plapinger, T. (2017, July 29). What is a decision tree? *Medium*. https://medium.com/data-science/what-is-a-decision-tree-22975f00f3e1

Powell, A. (2022, January 12). Personalisation: Stats and trends for 2022. *LXA Hub*. Retrieved from: https://www.lxahub.com/stories/personalisation-stats-trends

Statista. (2024). Worldwide: How brands fail to engage with customers 2023. *Statista*. Retrieved from: https://www.statista.com/statistics/1425554/brands-fail-to-engage-with-customers/

Tene, O., & Polonetsky, J. (2015). Theory of creepy: Technology, privacy and shifting social norms. *Yale Journal of Law and Technology*, 16: 59-102.

Torkamaan, H., Barbu, C.-M., & Ziegler, J. (2019). How can they know that? A study of factors affecting the creepiness of recommendations. *Proceedings of the 13th ACM Conference on Recommender Systems,* 423-427. ACM. https://doi.org/10.1145/3298689.3346982

Vanderschraaf, P., & Sillari, G. (2023). Common knowledge. In E. N. Zalta & U. Nodelman (Eds.), *The Stanford Encyclopedia of Philosophy* (Winter 2023 Edition). Stanford University. https://plato.stanford.edu/archives/win2023/entries/common-knowledge/

Van der Auwera, J., & Gast, V. (2012). Categories and prototypes. In J. J. Song (Ed.), *The Oxford handbook of linguistic typology* (Online ed.), 166-189. Oxford Academic. https://doi.org/10.1093/oxfordhb/9780199281251.013.0010

Wangler, B., Johnson, C., & Smith, D. (1999). A customer profiling framework for the banking sector. H. Jaakkola, H. Kangassalo, E. Kawaguchi, & I. A. Hamid (Eds.), *Frontiers in artificial intelligence and applications*, 61: 57-73. https://doi.org/10.13140/2.1.3888.1603

Watson, H. J. (2020). Avoid Being Creepy in the Use of Personal Data and Algorithms. *Business Intelligence Journal,* 24(1): 5-11.

Waupsh, J. (2022). How hyper-personalisation is fuelling digital banking's second wave. *Journal of Digital Banking. Henry Stewart Publications*, 7(2): 122-128, September. https://doi.org/10.69554/DIRI2621

Weijers, D. (n.d.). Hedonism. *Internet Encyclopedia of Philosophy.* https://iep.utm.edu/hedonism/#SH1f

Wolford, B. (n.d.). What is GDPR, the EU's new data protection law? *GDPR.eu*. https://gdpr.eu/what-is-gdpr/

Yoo, K.-H., Gretzel, U., & Zanker, M. (2013). Persuasive recommender systems: Conceptual background and implications. *Springer*. https://doi.org/10.1007/978-1-4614-4702-3

Non-exclusive licence to reproduce the thesis and make the thesis public

I, Palina Yaroshyk ,
(author´s name)

1. grant the University of Tartu a free permit (non-exclusive licence) to

   reproduce, for the purpose of preservation, including for adding to the digital archives of the University of Tartu until the expiry of the term of copyright, my thesis

   The Problem of Creepy Personalised Recommendations in Banks ,
   (title of thesis)

   supervised by Alexander Stewart Davies and Lennart Kitt ;
   (supervisor´s name)

2. grant the University of Tartu a permit to make the thesis specified in point 1 available to the public via the web environment of the University of Tartu, including via the digital archives, under the Creative Commons licence CC BY NC ND 4.0, which allows, by giving appropriate credit to the author, to reproduce, distribute the work and communicate it to the public, and prohibits the creation of derivative works and any commercial use of the work until the expiry of the term of copyright;

3. am aware of the fact that the author retains the rights specified in points 1 and 2;

4. confirm that granting the non-exclusive licence does not infringe other persons' intellectual property rights or rights arising from the personal data protection legislation.

Palina Yaroshyk
15/05/2025