When Do Customers Perceive Artificial Intelligence as Fair?
An Assessment of AI-based B2C E-Commerce

Helena Weith
University of Bern
helena.weith@iwi.unibe.ch

Christian Matt
University of Bern
christian.matt@iwi.unibe.ch

Abstract

Artificial intelligence (AI) enables new opportunities for business-to-consumer (B2C) e-commerce services, but it can also lead to customer dissatisfaction if customers perceive the implemented service not to be fair. While we have a broad understanding of the concept of fair AI, a concrete assessment of fair AI from a customer-centric perspective is lacking. Based on systemic service fairness, we conducted 20 in-depth semi-structured customer interviews in the context of B2C e-commerce services. We identified 19 AI fairness rules along four interrelated fairness dimensions: procedural, distributive, interpersonal, and informational. By providing a comprehensive set of AI fairness rules, our research contributes to the information systems (IS) literature on fair AI, service design, and human-computer interaction. Practitioners can leverage these rules for the development and configuration of AI-based B2C e-commerce services.

1. Introduction

Artificial intelligence (AI) provides businesses with new possibilities for AI-based business-to-consumer (B2C) e-commerce services. Recent examples range from smart mirrors providing style recommendations, to smart fridges autonomously placing grocery orders [1, 2]. AI offers advantages for firms, such as better exploitation of customer characteristics and profiles or higher personalization, and, thus, can increase customer satisfaction [3, 4]. AI-based services may, however, come at the cost of providing lower comparability and transparency of alternatives for customers, thus potentially affecting their fairness perceptions [5, 6].

Previous research on fairness not only demonstrated the positive impact of fair services on service quality perceptions, overall service satisfaction, and customers' re-patronage intention [7] but also provided a well-established fundament of four service fairness dimensions, namely procedural, distributive, interpersonal, and informational fairness, spanning the overall concept of systemic fairness [7, 8].

Research on fair AI has experienced a notable increase over the last years [9, 10], fueled not only by the diffusion of AI across industries and corporate tasks but also by popular examples of AI-induced unfairness, such as the decommissioned recruiting tool by Amazon [11]. Research on fair AI primarily focused on a technical perspective, whereas research on the customer-centric perspective including the customers' perception of fairness remains scarce [9, 10]. The latter is often limited to only two (procedural, distributive) among the four fairness dimensions [9, 12, 13, 14].

While recent research efforts on fair AI are laudable, we lack a concrete operationalization of fair AI from a customer perspective [9]. This is of utmost importance, as AI diffuses across industries and services, such as B2C e-commerce. Once customers perceive these services not to be fair, it can prompt undesired consequences [15] and harm businesses' long-term reputation [7].

Our research, therefore, aims to answer the question: *When do customers perceive AI as fair?* To obtain an understanding of fair AI, we conducted 20 semi-structured interviews with German B2C e-commerce customers. As underlying scenario, we used e-commerce use cases including products of different complexity. The interview guideline was structured along the four fairness dimensions (procedural, distributive, interpersonal, informational), spanning the overarching concept of systemic fairness [7]. Our qualitative research adds a contribution to the information systems (IS) literature on fair AI, service design, and human-computer interaction by understanding the customer perceptions of fair AI. We further provide practitioners with a comprehensive set of actionable fairness rules.

2. Theoretical Background

2.1. Service Fairness Dimensions

The role of fairness for services has been researched across several disciplines, highlighting the relevancy of service fairness and the resultant influence on customer

HICSS

behavior [16, 17]. "Firms failing to project an image of service fairness cannot develop the level of customer confidence needed to establish loyalty" [18]. Fairness research originates predominantly from an organizational and legal practice [19, 20], but can also be applied to the service context [7, 18]. Early research differentiated between only three fairness dimensions, i.e. distributive, procedural, and interactional. Later research [21] divided interactional into two further dimensions, namely interpersonal and informational. While the fairness dimensions were mainly considered individually, Beugré et al. [8] and Carr [7] strengthened and operationalized a holistic fairness perception named "systemic fairness." Based on B2C services, i.e. the interaction of employees with a service provider, Carr [7], demonstrated a positive impact of the four fairness dimensions on systemic service fairness, service quality, service satisfaction, and re-patronage intention.

Procedural fairness refers to the fairness of the process itself to achieve the outcome of a task or service. It is based on eight fairness rules, i.e. process control, decision control, correctability, consistency, bias suppression, accuracy, representatives, and ethicality [7, 19, 22]. Distributive fairness refers to the outcomes of processes, and is shaped by three rules, namely equity, equality, and need [7, 22, 23]. Equity refers to the outcome someone receives in return for the provided contribution, for example, goods received in exchange for money [22]. The equality rule describes the allocation of the identical outcome over all individuals [22]. The need rule implies that "persons with greater need should receive higher outcomes" [19]. Interpersonal fairness, as the first component of interactional fairness, refers to the social aspects of distributive fairness by "showing concern for individuals regarding the distributive outcomes they receive […] [and] focusing on the consequences of those outcomes" [21]. Regarding interpersonal fairness, apologies are illustrated as a tactic to harm the negative effects of an outcome [21]. Interpersonal fairness is shaped by the respect rule and the propriety rule [21, 22]. Informational fairness, as the second component of interactional fairness, refers to information provided about the process and is shaped by the rules of truthfulness and justification, which demand honest and thorough explanations [21, 22]. We base our research on the four fairness dimensions, the corresponding established fairness rules, and the overall concept of systemic fairness [7, 8].

2.2. Fair AI

AI defined as "the ability of a machine to perform cognitive functions" [24] achieved strong research presence over the last years and became strategically relevant in practice, as it enables services across various industries and functions [25]. However, there are several examples of AI that demonstrate unfair outcomes by causing bias and discrimination against individuals or groups. Examples include Amazon's recruiting tool [11] and algorithms favoring men over women for credit card limits [26]. Discrimination by AI can also occur in the e-commerce context, by providing products at different prices based on willingness to pay information derived from customer data [27]. Driven by the increasing presence of AI-based services and the indicated unfairness potential of AI, an increase in research on fair AI in information systems, computer science, and management emerged [9, 10]. While a predominant focus of research is on the field of human resources (HR) [9], there is a lack of focus not only on e-commerce services [4, 28] but also on the perception of AI fairness [9, 10].

According to Feuerriegel et al. [14], "the objective of fair AI is to provide systems that both quantify bias and mitigate discrimination against subgroups." This refers mainly to the procedural and distributive fairness dimensions. There is a notable focus on these two dimensions, and a lack of considering the interactional perspective (interpersonal, informational) [9, 10]. Furthermore, Robert et al. [9] emphasize the neglect of interactional fairness as a potential risk.

Research on fair AI can be split into two dimensions, i.e. technical and social, with technical further divided into pre-processing (prior to modeling), in-processing (point of modeling), and post-processing (after modeling) [29]. The technical dimension approaches the mitigation of biases and disparities of algorithms and data, for example, through a mathematical lens [30, 31]. A recurring challenge is the demand for fair data as base for fair decisions made by the AI and, thus, providing fair AI-based services [32]. Besides the technical dimension, also the social dimension, as well as the interchange of both dimensions, is of importance to achieve a comprehensive understanding of fair AI [33]. To strengthen the social dimension, it is relevant to understand "the human needs on fairness and whether the algorithmic definition of fairness actually map[s] to various stakeholders' perception of fairness" [34]. Furthermore, Feuerriegel et al. [14] mention a demand for understanding the perception of fair AI from a people perspective, for example, by understanding which attributes are considered as sensitive, and, thus, impact the fairness perception. Zhang et al. [35] address this demand by developing eight fairness rules for the design of talent intelligence management systems (TIMS). These rules aim to support an understanding of the employees' fairness perception. Zhang et al.

conducted their research based on interviews and a Delphi study that involved talent management experts.

A number of scholars, such as Marcinkowski et al. [13], have also focused on the perception of fairness, identifying that decisions made by algorithms are rather perceived as fair compared to decisions made by humans. In contrary, other academics argued that AI can also cause higher concerns and, thus, impact the fairness perceptions [6, 12]. The black box characteristics of the AI can serve as an explanation for the above-mentioned higher concerns [36]. Since the mechanisms of the AI might not be accessible and understandable for humans, the result is limited transparency, which further leads to a low perception of fairness [6, 37]. Previous research demonstrates that fairness perceptions also differ across people, based on their previous algorithmic experiences and cognitive styles [38]. Even though research partly explored the fairness perception of AI, Robert et al. [9] highlight the need to operationalize fair AI, particularly from a customer perspective, to understand AI fairness.

The variation of tasks performed by AI can impact the fairness perceptions too. Vimalkumar et al. [12] explored the impact of task complexity, including the automation potential and a resulting lack of transparency, and derived users' fairness concerns. These concerns differ along five different task types that vary in complexity. For instance, a decision task comprises a medium automation potential and high opacity by providing a product suggestion based on multiple possible outcomes. Corresponding fairness concerns included transparency, explainability, visualization, and comparability allocated to distributive or procedural fairness. Vimalkumar et al. envisaged a higher likelihood and variety of fairness concerns, the more complex the performed task is [12].

Although it is crucial to understand the perceptions of AI fairness from a customer perspective in the context of B2C e-commerce to build and maintain favorable long-term relationships with customers [7], the current understanding is thus far scarce, and therefore we address this gap with our research.

3. Methodology

3.1. Data Collection

We applied an interpretivist qualitative approach [39] to understand fairness perceptions for AI-based B2C e-commerce services. Since fair AI is an emerging research topic and particularly scarce regarding the customer-centric perspective, we chose semi-structured in-depth interviews to obtain fundamental insights into the research domain. We recruited our interviewees via social media (e.g. LinkedIn) and private networks.

We defined regular exploitation of e-commerce services as prerequisite for our interviewee selection. Most interviewees stated themselves as regular e-commerce users (minimum one purchase per month). Since we chose e-commerce use cases in the context of voice agents (VA), we asked the interviewees about their related experience. Seventeen of 20 interviewees indicated a general usage of at least one VA (including Alexa via Amazon Echo, Siri via Apple iPhone, and Google Assistant). However, the interviewees make use of the VA mainly to compile shopping lists, select music, or retrieve weather information. One interviewee used Alexa to purchase products.

The interviewees included 12 women and eight men, ranging from 24–65 years (average 35 years), all based in Germany. We conducted the interviews, which lasted between 35 and 60 minutes, via online conference tools. We pretested our literature-grounded questionnaire with three interviewees upfront, and slightly refined the interview questions and structure in response. During the interview phase, we reached a first indication for saturation after 12 interviews. However, to ensure a comprehensive understanding and full saturation, we performed in total 20 interviews.

The interview was structured in three phases, starting with a focus on the interviewee's general e-commerce behavior, his/her current usage of VAs, and potential concerns. Second, we focused on questions along a potential e-commerce service via VA and, thus, indirectly covered the four well-established fairness dimensions. However, at the same time, we provided enough flexibility to account for potential deviations from or extensions to those four dimensions. To ensure a common understanding of potential e-commerce services via VA, we included two videos that demonstrated examples [40, 41]. The demonstrated service covered the following steps: recognition of a product need, start of an interaction between the customer and the VA, presentation of a product suggestion (process outcome) by the VA, and completion of the process. Throughout the entire process, the customers in the videos and the VA interacted with each other. We emphasized that Alexa as represented in the videos, serves only as an example for a VA. We randomly presented interviewees with one of two consumer products of varying complexity as a basis for the e-commerce service. For low complexity, we chose a preference product (dog food), and for high complexity a shopping product (rain jacket) [42].

Third, finalizing the interviews, we invited the interviewees to summarize what they regarded the most relevant aspects pertaining to fairness. We allowed them enough time to reflect on the videos, their experiences, the overall interview, and to evolve further thoughts.

3.2. Data Analysis

Prior to the detailed analysis of our interview data, we prepared a complete and clean transcription of all 20 interviews. Subsequently, we performed our analysis structured in four steps (Table 1) of which the first three were supported by MAXQDA [43].

First, we approached the data in an open-minded manner by coding all the interviewees' statements with in vivo coding [44].

Table 1: Example of quote analysis

Quote			*"It is important that my information is used. The rain jacket suggestion must meet my expectations"* [Int. 8]
Step	1	Code	Expects suggestion to meet personal expectation
	2	Category	Individual and personalized outcome
	3	Group	Distributive
	4	Rule	Individuality

Second, we identified common categories combining the codes. We applied step one and two iteratively to ensure that all codes are represented by one of the categories until saturation was achieved. Third, we identified relationships between the categories and structured them in nine groups. Based on the dimensions established by Carr [7], we named the first five groups: procedural, distributive, interpersonal, informational, and systemic. We named the remaining four groups: usage behavior VA, product and service specifications, technology, and others. The latter we deprioritized in the following, as the content (e.g. sustainability concerns based on e-commerce deliveries) did not add to our research focus. Fourth, we approached the definition of our fairness rules based on the interview results and previously established fairness rules. Besides the rules (Table 2, Footnote 3) Carr [7] dealt with, we expanded our base by further rules (Table 2, Footnote 2) from an organizational perspective Colquitt et al. [22] summarized.

4. Findings

Based on the interviews and the previous literature, we identified 19 AI fairness rules (Table 2) in the context of B2C e-commerce services to provide an understanding about when customers perceive AI-based services as fair. Of the 19 rules, six are entirely new for the AI specific focus, and 13 rules are named in a similar manner or the same as in the previous literature. However, we adapted the descriptions for all fairness rules to the AI context.

Table 2: AI fairness rules

Established Dimension	Established Fairness Rule	No.[1]	New AI Fairness Rule	Description of AI Fairness Rule
Procedural	Process Control[2]	1	Process Control	Process allows for influence along the service steps
	Decision Control[2]	2	Decision Control	Process provides option to accept or deny outcome
	Correctability[2]	3	Correctability	Process enables cancellation and changes along the service steps
	Consistency[3]	4	Adapting Consistency	Process ensures consistency across customers, but AI ensures adaption to one customer over time
	-	5	Anticipation	Process enables anticipation of future demand driven by AI
	Bias Suppression[3]	6	Bias Suppression	Process suppresses bias in the interest of the AI and other profiting parties
	Accuracy[2]	7	Accuracy	Process ensures accuracy of information involved
	Representatives[4]		-	-
	Ethicality[3]	8	Ethicality	Process ensures adherence to individual's morals and values
	-	9	Competency	Process demonstrates competency of AI
	-	10	Security	Process ensures security and protection of data
Distributive	Equity[2]	11	Data-Based Equity	Outcome is allocated based on amount of contributed data and insights derived from data
	-	12	Individuality	Outcome is personalized to each customer
	Equality[3]	13	Equality	Outcome details are the same across customers

	Need[4]		-	-
	-	14	Alternatives	Outcome includes customer relevant alternatives
	-	15	Independency	Outcome is independent of third party interests
Inter-personal	Respect[3]	16	Respect	Outcome is considerate and polite
	Propriety[3]	17	Propriety	Outcome is adequate and proper
Infor-mational	Truthfulness[2]	18	Truthfulness	Process and outcome information are honest and transparent
	Justification[3]	19	Justification	Process and outcome explanations are detailed and plausible

[1] Number of new AI fairness rule
[2] Fairness rule added based on Colquitt et al. [22]
[3] Fairness rule dealt with by Carr [7]
[4] Fairness rule included by Colquitt et al. [22], but not dealt with by Carr [7]. Included in the table for comprehensive overview.

4.1. Procedural AI Fairness Rules

All interviews provided indications of *process control.* The interviewees demanded the opportunity to take control along various steps of the e-commerce process, for example, by deciding about data involved in the process. Another example of process control is the option to adjust the AI proactiveness or the level of information (informational fairness) provided throughout the e-commerce process. Besides control over the process, the interviewees also requested *decision control* over the outcome, i.e. the product suggestion provided by the AI. An interviewee (Int.) outlined the demand for an opportunity to accept or deny the product suggestion: *"I want the opportunity to approve the product at the end before purchasing it"* [Int. 9]. Taking control away from the customer, for example, by not actively asking for a purchase confirmation, would harm the fairness perception.

Whereas decision control refers to the outcome of the e-commerce process, *correctability* focuses on decisions along the process. The interviewees requested the opportunity to change or modify product preferences they had previously provided throughout the process. Further *consistency* of the process, especially across different customers, was demanded. However, the interviewees expected the AI to *adapt* consistency of the process over time based on customer specifications. An interviewee mentioned: "*If it is a product I order regularly, I want the process to be shorter. Only the information about what I ordered last, followed by the question if I want to order the same again. That is enough—I don't need all details again"* [Int. 1].

Since the interviewees associated capabilities like processing big data and identifying patterns with the VA, most interviewees generally perceived it as fair if demand is *anticipated.* Three different levels of anticipation (low, medium, high) resulted from the interviews: alternative suggestions close to the actual product intended for the customer (low), reminder for recurring purchases (medium), proactive suggestions based on customer data (high). The latter could be realized by suggesting a rain jacket purchase due to upcoming calendar entries and the interviewee's vacation plans, combined with weather forecasts. The preference for the level of anticipation varied across the interviewees. However, they shared the preference for anticipation in combination with process control, explained as follows: *"If I buy that product once every four weeks [medium anticipation], it would be fair to receive a reminder maybe after three weeks. But I want to adjust that in the settings"* [Int. 9].

Most interviewees highlighted the expectation of AI *suppressing bias* from any party that potentially profits from the service. Bias suppression should ensure that no other interests, but the ones of the interviewee impact the process of the e-commerce service. Bias suppression is closely related to the later introduced independency rule (distributive), which, in contrast to the process, focuses on the outcome.

Furthermore, process relevant data must always be *accurate.* Our interviewees would hardly accept potential data errors, as they expect the AI to be capable of handling data with proficiency, and error free. The interviewees required the e-commerce process to be accurate not only about the details of the data involved but also regarding the type of data involved. They would not perceive it as fair once data is exploited without their approval. The general incorporation of data by the AI should follow *ethically* acceptable procedures. Ethicality could be harmed by invading one's privacy to gather information, for instance, *"when the VA recommends me products I had a conversation about with my partner, even though I did not mention to the VA at all that I want to buy them"* [Int. 9]. Furthermore, the e-commerce process must demonstrate *competence* of the AI to achieve the interviewees' fairness perception as explained by an interviewee: "*It is a computer; I would expect far more than from a person in store. It knows so much more"* [Int. 10]. This could be demonstrated by providing transparency about the

process and by thorough explanations about the reasoning for suggestions (informational fairness). Lastly, as highlighted by almost all the interviewees, *security* of the IT processes and a corresponding data protection are core aspects to achieve fairness perception.

4.2. Distributive AI Fairness Rules

Focusing on process outcomes, we identified five rules to be part of the distributive dimension. *Data based equity* refers to the outcome of an e-commerce service in return for the data contributed, i.e. the extent and amount to which the customer provides data to the AI and e-commerce provider. If the interviewees decided to withhold their data (process control) and, thus, receive less personalization or a lower match with their actual requirements, they would still perceive the outcome as fair. An interviewee explained it thus: "*Well, it is fair, if you argue that data are the new currency [...]. However, it must be communicated clearly—maybe like a disclaimer, which says that you receive a better matching service the more data you provide*" [Int. 6].

Personalization is at the core of the second distributive AI fairness rule, namely *individualization.* Outcomes which are dedicated to and personalized for the individual are perceived as fair and are even more expected due to the technical capabilities expected from the AI: "*I think it is absolutely fair to use my data. Therefore, I receive advertising truly tailored to me. Since these opportunities are provided, it should also be leveraged. Furthermore, no one is forced to use it; you can also go into the store instead*" [Int. 4].

Even though the interviewees requested an individualized outcome, they expected it to be provided *equally.* This means that a congruent product suggestion should be presented with the same conditions per customer, without preference or discrimination: "*The price has to be the same for both customers. But if the other person receives a different product suggestion, which is maybe more expensive but also fits the demand of that person better, then it is still fair*" [Int. 3]. For a process outcome to be perceived as generally fair, the interviewees demanded *alternative* outcomes (e.g. product suggestions) to be provided, followed by the customer making a final decision. If customers decide to forego (process control) the alternative suggestion (e.g. due to limited time available on their side to complete the process) and, thus, miss a potentially attractive price-value alternative, the equality rule would not be violated. Almost all the interviewees stated that an *independent* outcome of the process, i.e. an outcome not driven by any third party interests, would be perceived as fair and would be favored by them.

4.3. Interpersonal AI Fairness Rules

We defined two AI fairness rules for interpersonal fairness: *respect* and *propriety*. The AI can possibly harm the fairness perception of these two rules, if the AI fails to consider the impact of the e-commerce service outcome on the recipient. "*Regarding clothes, I would hope that the AI does not know how heavy I am. I think it would be okay for me that I mention my size [...]. But I have a smart scale connected to my phone, and it would now definitely be unfair if the VA tells me that I gained weight and that I should buy new clothes in bigger sizes*" [Int. 8]. Such suggestions, which lack respect and consideration of how the customer experiences the outcome, could result in the failure to achieve AI interpersonal fairness. Inappropriate behavior by exerting pressure would further harm the interviewees' fairness perceptions: "*I do not want to feel pressured. I would even appreciate it if the VA inquires whether I really need the product or not. You know, like making me aware that I might already have enough outdoor jackets or so*" [Int. 14].

4.4. Informational AI Fairness Rules

An interviewee stated at the end of the interview: "*One of the most important aspects is transparency—not only transparency, but also the freedom to decide and adjust settings. [...] And besides transparency, I always want to be able to understand the decision process*" [Int. 6]. Based on such statements, we defined two rules for informational fairness, namely *truthfulness* and *justification,* both focusing not only on procedural aspects but also on the outcome of the e-commerce service, i.e. product suggestion. Truthfulness refers to correct information and to transparency about potential alternatives. An interviewee explains it like this: "*I would perceive it as fair, if Alexa points out that I could also buy an alternative product, which is maybe on sale, and does not withhold that information*" [Int. 5]. There is a fluent transition from truthfulness to justification. The interviewees requested information that explains the reasoning behind how and why a certain outcome results, i.e. a product is suggested. They did not demand technological specifications about AI, but rather plausible reasonings, for example, that a suggestion is based on the individual's order history or on reviews of other customers: "*I would want to know why a certain product is suggested to me. And if it is a sponsored recommendation, I would definitely want to know that*" [Int. 16].

5. Discussion

Our research contributes to the understanding of the customer fairness perceptions of AI. We achieved this by providing 19 AI fairness rules in the context of AI-based B2C e-commerce services. We, thereby, built on previously established fairness rules [7, 22], which provided a thorough base, but extended and adapted those to an AI-based service context. The majority of our AI fairness rules are allocated to procedural fairness, which provides an indication for the importance of the dimension [10, 45]. Active demand for the rules of interpersonal AI fairness was the least proactively mentioned by our interviewees. We noticed, however, that this does not imply low relevance, but rather demonstrates that respect and propriety of distributive outcomes are expected to be provided. In line with Dodge et al. [38], we recommend to further consider interactional fairness (interpersonal and informational) as relevant, especially since rude behavior by the AI can occur occasionally [46].

We believe that all rules are relevant to a holistic AI fairness perception of e-commerce services, but the rules might not be limited to e-commerce services. We recognize that certain rules, such as security, are also relevant in the context of other technology-based services [10]. Furthermore, even though the independency and bias suppression rules are of importance and even though the interviewees would appreciate these rules, they are also aware that a provider of an AI-based e-commerce service most likely pursues economic interests. Such an interest could occur in favoring products with high margins. Since customers are skeptical about the service provider's interests, they require transparency and justification even more. This provides us with an indication that rules, i.e. independency and bias suppression, should still be considered partially, if it is not possible to consider them completely.

Compared to the established fairness rules, two are not represented in our AI-based e-commerce context, as the interviewees provided no evidence for them. The representatives rule defines that procedures "should take into account concerns of subgroups" [22], i.e. involve sub-groups as representatives in the process. Since we base our research on B2C customer interviews, we already represent the concerns of customers as "subgroup" throughout our research. The need rule implies that someone who is more needy should receive a higher outcome [19], but this contradicts with our AI fairness rules bias suppression and equality.

The rules and dimensions show a high interrelation within and across the four fairness dimensions. For instance, a highly automated service (e.g. high anticipation of future demand) causes a higher customer demand for transparency (truthfulness). If, however, it is not possible to provide transparency sufficiently, the fairness perception could be violated. This finding is in line with previous research that demonstrates the skepticism toward the unknown aspects of the AI [6, 36, 37]. Moreover, the process control rule indicates a strong relation with a variety of further rules across other dimensions too, such as anticipation, alternatives, and justification. We also expect this strong relation to be driven by the AI black box characteristics [6, 36, 37], as customers may not be able to comprehend the algorithmic mechanisms and, thus, aim to compensate for this inability by controlling the process.

Given the strong interrelation of the rules and the four required dimensions, we advocate a holistic perspective of AI fairness, previously summarized as systemic fairness [7, 8]. We further suspect a variation in importance of individual rules. Certain rules even appear to incorporate the role of deal breakers, which, in our definition, are AI fairness rules with the power to interrupt the systemic fairness independently of any other rule. Ethicality would be such a rule, as an unethical process could disrupt the overall fairness perception. Such an unethical process would occur once the AI makes use of data from customers, they did not agree to (e.g. data retrieved by a VA listening to conversations without customer consent). Even though the outcome of the e-commerce service itself might be perceived as fair, the unethical process could overrule the fairness perceptions that result from all other rules. This can further lead to customers refraining completely from the AI-based service [47].

We also noticed a certain degree of discrepancy regarding the demands that the interviewees communicated. Discrepancies did not only occur across interviewees; single interviewees expressed contradictory expectations too. On the one hand, process control is demanded. On the other hand, however, interviewees value the AI's capability for automation, adaption, and individuality. Exploitation of the AI capabilities and providing AI-based e-commerce services fully meeting the customer demand, could result in less demand for process control. We assume, however, that process control will not become fully redundant, as customers tend to evaluate AI-based suggestions that they can control as much more positive than having no control [48]. Going forward, it would be of interest to study the demand and progression of the process control rule. Moreover, it should further be investigated how AI fairness rules generally change regarding the importance and their relation to each other, the more presence AI receives.

We lastly noted two determinants that impact the AI fairness perceptions in the context of an e-commerce service: task complexity and customer profiles. Task

complexity within our use cases is driven by the nature of the service (first-time vs. recurring), and also by product complexity. We noticed that recurring services are mostly perceived as fair, if they were already perceived as fair previously [38]. The recurring service could even be provided with less transparency (truthfulness), but still be perceived as fair. Furthermore, a proactive behavior by the AI for a recurring service, results in the process being perceived as fair. Besides the process, the product complexity (low complexity: dog food vs. high complexity: rain jacket) impacts the fairness perception too. In line with process complexity, high product complexity can hamper the fairness perception as well. The task complexity's identified impact on fairness perception is in line with the fairness concerns that Vimalkumar et al. [12] identified for decision tasks, especially the transparency concern and the explainability concern. We see this further anchored in previous research that aims to understand the impact of transparency on AI fairness [6, 12, 37]. We generally expect all rules to remain relevant even for tasks of low complexity. However, certain rules and aspects, such as justification and transparency (truthfulness), might be of less importance with regard to a recurring purchase or a low complex product. In line with Dodge et al. [38], we recognized differences regarding the AI fairness perceptions that are influenced by different customer profiles, by their previous experiences with AI, and also by their general attitude toward AI. Based on indications from our interviews, we assume that customers who make high use of e-commerce services have a higher expectation of rules like adapting consistency, or anticipation.

6. Implications and Limitations

6.1. Implications for Theory

Our research contributes to the IS literature on fair AI, service design, and human-computer interactions. We achieve this not only by improving the understanding of customer fairness perceptions of AI-based services but also by identifying relevant AI fairness rules. Our customer-centric approach contributes to the social dimension of fair AI research [29] and further supports bridging the gap between the social and the technical research dimensions [33]. We address the demand highlighted by Robert et al. [9] and Dodge et al. [38] for a more customer-centric perspective.

By selecting e-commerce use cases, we provide a contribution to expand the AI fairness research in the field of e-commerce [4, 28], compared to a current emphasis on the field of HR [9]. However, we also demonstrate overlaps (e.g. accuracy, bias suppression)

with fairness rules for the design of TIMS by Zhang et al. [35]. Compared to the latter, and to obtain a thorough customer perspective, we based our data collection on potential customers as the recipients of the service.

As suggested by previous research [9, 46], we did not limit our scope to a selection of fairness dimensions upfront, which is contrary to most other studies. Furthermore, our interviews strengthened the demand for AI fairness rules along all four fairness dimensions and indicated their interrelation. We can, thus, confirm the relevance of a holistic fairness consideration based on all four dimensions summarized by systemic AI fairness as introduced by Carr [7] and Beugré et al. [8].

Since we indicated that transparency of AI-based services impacts the customers' perceptions of fairness, we can confirm findings of previous research by Zarsky et al. [6] and Martin [37]. Furthermore, our rules demonstrate overlaps with the user concerns regarding AI fairness that Vimalkumar et al. defined [12]. Their concerns are more likely to emerge, the more complex the task performed by the AI is. Our interviews also provide an indication for task complexity impeding the customer AI fairness perception.

Lastly, our 19 AI fairness rules in the context of AI-based B2C e-commerce services provide a base for future research to develop AI fairness rules dedicated to further AI-based services, whereby we respond to previous calls for research [9, 10, 24].

6.2. Implications for Practice

Recognizing that ignoring customers' AI fairness perceptions can cause long-term B2C reputational risks [7, 15], our research also contributes to practice in several manners. Our operationalization of fair AI supports practitioners to achieve a thorough understanding of the customer perception of fair AI. The 19 actionable AI fairness rules can be applied for the development and in-depth design of AI-based B2C e-commerce services. Besides our use cases, firms can take our AI fairness rules and adapt them for other types of AI-based services.

As our results indicate, it is important to understand that neglecting only one of these rules can harm the B2C relationships. Another concrete implication results from the independency and bias suppression rules. The interviewees were aware that the AI implemented in e-commerce services is most likely not independent, as businesses that follow their own economic interests provide it. The interviewees appeared to tolerate this to a certain degree, also due to lacking alternatives. However, if competitors can demonstrate that their AI acts independently of their economic interests (e.g. by using external certifications), it might put pressure on those companies cannot certify this.

We further provide an indication that high complexity tasks (product and process) can result in less transparency and, thus, harm the fairness perceptions. Firms need to be aware of the impact of complexity and consider it during service design. AI fairness perceptions can further vary across customers, thus emphasizing the need to develop and design customer-specific service design. Businesses should also internalize AI fairness rules not only to develop new services but also to reassess existing ones.

Besides implications for businesses, our results provide a basis for regulatory authorities and customer-interest groups to understand which factors matter for individual's AI fairness perception. They can make use of these factors to judge to which degree services violate customer demands, and develop labels or prepare charges in case of substantial violations. As such, through the development of a common ground for customer fairness perceptions, AI-based services might also help achieve overall trust of AI and, therefore, serve as a critical companion to the future development of AI-based B2C e-commerce services.

6.3. Limitations and Future Research

Our research is not free of limitations. First, the concrete choice of the e-commerce service for our interviews might have impacted the results. Furthermore, the low practical familiarity was challenging, but also not surprising due to the thus far low establishment of e-commerce services via VA. However, we explained a potential e-commerce service in detail in the interviews supported by two videos. Moreover, the interviewees were, thus, not biased by preferences or resentments toward established services.

Second, we could only sample the interviewees' perceptions as a single snapshot of time. However, it is plausible that AI fairness perceptions develop along with AI-based service availability and usage over time. We, therefore, recommend analyzing the development of fairness perceptions, the interactions between the rules, and their impact on businesses over time. Future research could also focus, subject to variations of task complexity or customer profiles, on the confirmation of single rules that by themselves prevent a systemic fairness or identify the weight of individual rules.

7. References

[1] Marr, B., The Magic of Smart Mirrors: Artificial Intelligence, Augmented Reality and the Internet of Things, https://www.forbes.com/sites/bernardmarr/2019/10/04/the-magic-of-smart-mirrors-artificial-intelligence-augmented-reality-and-the-internet-of-things/?sh=5a5c98c4615b, Accessed: 20th March 2021.

[2] Chitkara, H., Samsung and LG are Set to Unveil AI-Equipped Smart Fridges with Food-Identification Capabilities, https://www.businessinsider.com/samsung-lg-unveil-artificial-intelligence-equipped-smart-fridges-2020-1?r=US&IR=T, Accessed: 20th March 2021.

[3] Saffarizadeh, K., Boodraj, M., and Alashoor, T., Conversational Assistants: Investigating Privacy Concerns, Trust, and Self-Disclosure, International Conference on Information Systems, South Korea, 2017, pp. 1-12.

[4] Ameen, N., Tarhini, A., Reppel, A., and Anand, A., Customer Experiences in the Age of Artificial Intelligence, Computers in Human Behavior, Vol. 114, 2021, pp. 1-14.

[5] Robert, L.P., Bansal, G., Melville, N., and Stafford, T., Introduction to the Special Issue on AI Fairness, Trust and Ethics, AIS Transactions on Human-Computer Interaction, Vol. 12(4), 2020, pp. 172-178.

[6] Zarsky, T., The Trouble with Algorithmic Decisions: An Analytic Road Map to Examine Efficiency and Fairness in Automated and Opaque Decision Making, Science, Technology & Human Values, Vol. 41(1), 2016, pp. 118-132.

[7] Carr, C.L., The FAIRSERV Model: Consumer Reactions to Services Based on a Multidimensional Evaluation of Service Fairness, Decision Sciences, Vol. 38(1), 2007, pp. 107–130.

[8] Beugré, C.D., and Baron, R.A., Perceptions of Systemic Justice: The Effects of Distributive, Procedural, and Interactional Justice, Journal of Applied Social Psychology, Vol. 31(2), 2001, pp. 324-339.

[9] Robert, L.P., Pierce, C., Marquis, L., Kim, S., and Alahmad, R., Designing Fair AI for Managing Employees in Organizations: A Review, Critique, and Design Agenda, Human-Computer Interaction, Vol. 35(5-6), 2020, pp. 545-575.

[10] Kordzadeh, N., and Ghasemaghaei, M., Algorithmic Bias: Review, Synthesis, and Future Research Directions, European Journal of Information Systems, published online 6th June 2021, https://www.tandfonline.com/doi/full/10.1080/0960085X.2021.1927212, Accessed: 11th June 2021.

[11] Dastin, J., Amazon Scraps Secret AI Recruiting Tool That Showed Bias Against Women, https://www.reuters.com/article/us-amazon-com-jobs-automation-insightidUSKCN1MK08G, Accessed: 2nd April 2021.

[12] Vimalkumar, M., Gupta, A., Sharma, D., and Dwivedi, Y., Understanding the Effects that Task Complexity has on Automation Potential and Opacity: Implications for Algorithmic Fairness, AIS Transactions on Human-Computer Interaction, Vol. 13(1), 2021, pp. 104–129.

[13] Marcinkowski, F., Kieslich, K., Starke, C., and Lünich, M., Implications of AI (Un-)Fairness in Higher Education Admissions: The Effects of Perceived AI (Un-)Fairness on Exit, Voice and Organizational Reputation, Conference on Fairness, Accountability, and Transparency, Barcelona, 2020, pp. 122-130.

[14] Feuerriegel, S., Dolata, M., and Schwabe, G., Fair AI: Challenges and Opportunities, Business & Information Systems Engineering, Vol. 62 (4), 2020, pp. 379-384.

[15] Fernandes, T., and Calamote, A., Unfairness in Consumer Services: Outcomes of Differential Treatment of New and Existing Clients, Journal of Retailing and Consumer Services, Vol. 28, 2016, pp. 36-44.

[16] Kahneman, D., Knetsch, J.L., and Thaler, R.H., Fairness and the Assumptions of Economics, The Journal of Business, Vol. 59 (4:2), 1986, pp. 285-300.

[17] Parasuraman, A., Berry, L.L., and Zeithaml, V. A., Understanding Customer Expectations of Service, Sloan Management Review, Vol. 32(3), 1991, pp. 39-48.

[18] Seiders, K., and Berry, L.L., Service Fairness: What it is and Why it Matters, Academy of Management Executive, Vol. 12(2), 1998. pp. 9-20.

[19] Leventhal, G.S., What Should Be Done with Equity Theory? New Approaches to the Study of Fairness in Social Relationships, Gergen, K.J., Greenberg, M.S., and Willis, R.H. (Eds.) Social Exchange: Advances in Theory and Research, Plenum Press, New York, 1980, pp. 27-55.

[20] Seidman Diamond, S.S., and Zeisel, H., Book Review: Procedural Justice: A Psychological Analysis. By John Thibaut and Laurens Walker, Duke Law Journal, Vol. 1977(6), 1978, pp. 1289-1296.

[21] Greenberg, J., The Social Side of Fairness: Interpersonal and Informational Classes of Organizational Justice, Cropanzano, R., (Ed.), Series in Applied Psychology. Justice in the Workplace: Approaching Fairness in Human Resource Management, Lawrence Erlbaum Associates, Hillsdale, New Jersey, 1993, pp. 79-103.

[22] Colquitt, J.A., and Rodell, J.B. Measuring Justice and Fairness, Cropanzano, R.S., Ambrose, M.L. (Eds.), The Oxford Handbook of Justice in the Workplace, Oxford University Press, New York, 2015, pp. 187-202.

[23] Leventhal, G.S., The Distribution of Rewards and Resources in Groups and Organizations, Advances in Experimental Social Psychology, Vol. 9, 1976, pp. 91-131.

[24] Rai, A., Constantinides, P., Sarker, S., Editor's Comments, Next-Generation Digital Platforms: Towards Human-AI Hybrids, MIS Quarterly, Vol. 43(1), 2019, pp. iii-ix.

[25] Borges, A.F.S., Laurindo, F.J.B., Spínola, M.M., Gonçalves, R.F., and Mattos C.A., The Strategic Use of Artificial Intelligence in the Digital Era: Systematic Literature Review and Future Research Directions, International Journal of Information Management, Vol. 57, 2021, pp. 1-16.

[26] Thorbecke, C., New York Probing Apple Card for Alleged Gender Discrimination After Viral Tweet, https://abcnews.go.com/US/york-probing-apple-card-alleged-gender-discrimination-viral/story?id=66910300, Accessed: 25th May 2021.

[27] Gautier, A., Ittoo, A., and v. Cleynenbreugel, P., AI Algorithms, Price Discrimination and Collusion: A Technological, Economic and Legal Perspective, European Journal of Law and Economics, Vol. 50, 2020, pp. 405-435.

[28] Guha, A., Grewal, D., Kopalle, P.K., Haenlein, M., Schneider, M.J., Jung, H., Moustafa. R., Hegde, D., and Hawkins, G., How Artificial Intelligence Will Affect the Future of Retailing, Journal of Retailing, Vol. 97(1), 2021, pp. 28-41.

[29] Caton, S., and Haas, C., Fairness in Machine Learning: A Survey, arXiv:2010.04053 [cs.LG], 2020, pp. 1-33.

[30] Haas, C., The Price of Fairness – A Framework to Explore Trade-Offs in Algorithmic Fairness, Fortieth International Conference on Information Systems, Munich, 2019, pp. 1-17.

[31] Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., and Galstyan, A., A Survey on Bias and Fairness in Machine Learning, arXiv:1908.09635 [cs.LG], 2019, pp. 1-31.

[32] Barocas, S., Hardt, M., and Narayanan, A., Fairness in Machine Learning – Limitations and Opportunities, https://fairmlbook.org/, Accessed: 20th March 2021.

[33] Goodman, B.W., Economic Models of (Algorithmic) Discriminations, 29th Conference on Neural Information Processing Systems, Barcelona, 2016, pp. 1-10.

[34] Zhang, Y., Bellamy, R.K.E., Liao, Q.V., and Singh, M., Introduction to AI Fairness, CHI Conference on Human Factors in Computing Systems Extended Abstracts, Yokohama, 2021.

[35] Zhang, X., Zhao, Y., Tang, X., Zhu, H., and Xiong, H., Developing Fairness Rules for Talent Intelligence Management System, Proceedings of the 53rd Hawaii International Conference on System Sciences, Maui, 2020, pp. 5882-5891.

[36] Pasquale, F., The Black Box Society, The Secret Algorithms That Control Money and Information, Harvard University Press, Cambridge Massachusetts, 2015.

[37] Martin, K.E., Designing Ethical Algorithms, MIS Quarterly Executive, Vol.18(2), 2019, pp. 129-142.

[38] Dodge, J., Liao, Q.V., Bellamy, R.K.E., and Dugan, C., Explaining Models: An Empirical Study of How Explanations Impact Fairness Judgement, Proceedings of the 24th International Conference on Intelligent User Interfaces, Los Angeles, 2019, pp. 275-285.

[39] Saunders, M.N.K., and Bristow, A., Research Methods for Business Students, 7th Edition, Pearson Education, Harlow, 2015.

[40] amazon, Voice Shopping with Alexa, YouTube, 2016, https://www.youtube.com/watch?v=mCjvV3iFsuw, Accessed: 5th April 2021.

[41] BusinessWire, Alexa, Order from Pime Now!, YouTube, 2017, https://www.youtube.com/watch?v=7zsd1VN1R9c, Accessed: 5th April 2021.

[42] Enis, B.M., and Roering, K.J., Product Classification Taxonomies: Synthesis and Consumer Implications, Lamb, C.W., and Dunne, P.M. (Eds.), Theoretical Development in Marketing, Marketing Classics Press, Decatur, 1980, pp.186-189.

[43] Kuckartz, U., Qualitative Inhaltsanalyse. Methoden, Praxis, Computerunterstützung, Beltz Juventa, Weinheim Basel, 2018.

[44] Miles, M.B., Huberman, A.M., and Saldaña, J., Qualitative Data Analysis: A Methods Sourcebook, Sage Publications, Thousands Oaks, 2019.

[45] Ötting, S.K., and Maier, G.W., The Importance of Procedural Justice in Human-Machine Interactions: Intelligent Systems as New Decision Agents in Organizations, Computers in Human Behaviour, Vol. 89, 2018, pp.27-39.

[46] Binns, R., v. Kleek, M., Veale, M., Lyngs, U., Zhao, J., and Shadbolt, J., "It's Reducing a Human Being to a Percentage": Perceptions of Justice in Algorithmic Decisions, Conference on Human Factors in Computing Systems, Montreal, 2018, pp. 1-18.

[47] Ebrahimi, S., and Hassanein, K., Empowering Users to Detect Data Analytics Discriminatory Recommendations, International Conference on Information Systems, Munich, 2019, pp. 1-9.

[48] Harper, M.F., Xu, F., and Harmanpreet, K., Putting Users in Control of Their Recommendations, Conference on Recommender Systems, Vienna, 2015, pp. 3-10.