



Explaining automated decision-making: a multinational study of the GDPR right to meaningful information

Jacob Dexe^{1,2}  · Ulrik Franke^{1,2} · Kasia Söderlund³ · Niels van Berkel⁴ · Rikke Hagensby Jensen⁴ · Nea Lepinkäinen⁵ · Juho Vaiste⁵

Received: 15 June 2021 / Accepted: 29 March 2022 / Published online: 3 May 2022
© The Author(s) 2022

Abstract

The General Data Protection Regulation (GDPR) establishes a right for individuals to get access to information about automated decision-making based on their personal data. However, the application of this right comes with caveats. This paper investigates how European insurance companies have navigated these obstacles. By recruiting volunteering insurance customers, requests for information about how insurance premiums are set were sent to 26 insurance companies in Denmark, Finland, The Netherlands, Poland and Sweden. Findings illustrate the practice of responding to GDPR information requests and the paper identifies possible explanations for shortcomings and omissions in the responses. The paper also adds to existing research by showing how the wordings in the different language versions of the GDPR could lead to different interpretations. Finally, the paper discusses what can reasonably be expected from explanations in consumer oriented information.

Keywords GDPR · Right of access · Meaningful information · Transparency · Insurance · Automated decision-making

Introduction

In May 2018, the European General Data Protection Regulation (GDPR) came into effect in all EU Member States, and in July that same year also in the EEA-states Norway, Iceland, and Lichtenstein. Broadly speaking, it aims to increase transparency and strengthen the rights of data subjects, thus building trust in the digital single market. Furthermore, EU regulations such as the GDPR can have a substantial impact outside the union as well, for example if companies active in many jurisdictions find it convenient to adhere to EU rules globally to avoid the hassle of

✉ Jacob Dexe
jacob.dexe@ri.se

Extended author information available on the last page of the article



managing a plethora of local rules (for a discussion of the GDPR from this perspective, see Bradford 2020, Chap. 5).

One of the rights established in the GDPR is an individual's right to request "meaningful information about the logic involved" in "automated decision-making" based on the personal data of the individual (GDPR 2016, Article 15). However, this right is not unconditional. The company that responds to such a request has to consider what meaningful information is, what logic actually is involved and at which levels, and how to consider other limitations—including business confidentiality, budget, time, and technical ability. At the same time, the individual who uses the right to get information from a company will already have some preconceptions about the service in question, as well as an understanding as to what kind of processing is reasonable and what meaningful information about this might look like. Furthermore, as argued below, it is currently very difficult to determine the legal meaning of the GDPR in this respect in an objective and exact way, though this may become easier in the future when relevant court decisions appear.

In a poll conducted on behalf of the European Commission in 2017, 61% of EU citizens say that privacy and security feature in their choices of IT products (European Commission 2017). 61% is also the number of EU citizens who have a fairly positive or very positive opinion about AI, the set of technologies which are increasingly often behind the 'automated decision-making' mentioned in the GDPR.

With this in mind, it is worth considering what the right to meaningful information actually looks like in practice: How do companies go about interpreting this right and balancing the interests of both requesters and their own organisation in the absence of clear-cut guidance?

This paper extends our previous work (Dexe et al. 2020). In the previous study, we conducted an empirical investigation of how Swedish insurance companies responded to information requests made by consumers regarding automated decisions in the pricing of home insurance. The paper showed that while it was hard to determine if any company failed to follow the legislation (except in one case, where the response failed to meet the 'undue delay' requirement), there was also no company that ticked all the boxes and provided an answer that could be deemed 'complete'. Thus, while previous research has found large differences in how companies handle article 15 requests (Sørum and Presthus 2020; Dexe et al. 2020) these studies are conducted in single countries. In order to get a broader and more reliable view of how article 15 works throughout the EU, additional studies such as the one presented here are called for.

To do so, we have gathered researchers representing five EU member states, and performed the same data collection as in the previous study. However, in addition to repeating the same study in more countries, the present study also goes beyond the previous study in two respects:

First, while Dexe et al. (2020) offered a description of the history of, and recommendations for, applying the GDPR, the article did not give any further insight into why companies in practice seemed to have interpreted the requirements in such different manners. Second, it provided a limited background and little discussion on how to assess the quality of the responses outside of the lens of legal requirements. In this paper we explore how insights from other fields, such as Human-Computer



Interaction (HCI), might provide different perspectives on how data can be presented and communicated to customers. We address these aspects by using data gathered from the insurance markets of five EU countries. More precisely, the following research questions are investigated:

- RQ1** What do insurance companies in the EU disclose when consumers ask about the logic behind automated decisions?
- RQ2** What differences and similarities can be observed in the responses?
- RQ3** How can the differences and similarities in the responses be explained?

These research questions were investigated by use of actual consumer requests to 26 separate insurers in five EU countries.

While the GDPR applies to all industries, insurance offers an interesting case to study since it is a business that often cites trust as a key value. When a customer enters into a contract with an insurer, no physical goods change hands. Instead, a relation of trust is established: if certain conditions are met in the future, the insured will receive a compensation under the terms agreed to. Clearly, trust plays an important role here, and insurers thus have every reason to be transparent and responsive to customer requests about how their data is used and how decisions are made (see, e.g., Scott 2004). It should be noted that profiling and use of automated decision making in insurance is hardly a new phenomena. Insurance companies often stand on solid legal ground to perform such processing (DAC Beachcroft 2016). This paper builds on this understanding by investigating the companies' explanations of such processing.

The remainder of the paper is structured as follows: The next section introduces related work. The subsequent section explains the data collection methods used, followed by a presentation of the results. These are then discussed, before the final section concludes the paper.

Related work

The study reported in this paper resides at the intersection of several fields; law, information systems, insurance studies, etc. In the following, we review related work from a few different perspectives.

GDPR effects related to transparency and insurance

Since the GDPR came into force in May 2018, scholars and scientists have had the opportunity to study its consequences empirically. Though the time elapsed has been relatively short, the empirical GDPR literature is steadily growing. A study similar to both our present and previous study is Sørum and Presthus (2020), where consumer rights stipulated in article 15 and 20 of the GDPR are used to approach 15 companies on the consumer market. They analyse the time to receive responses, the quality of feedback and compliance with the different paragraphs in articles 15



and 20. The authors note that with respect to Article 15(1)(h), no company provided explanations about algorithms and automated decision-making. Finally, Sørum and Presthus offer advice for consumers wanting to use their rights under the articles—advice which has been followed in the present article.

Other studies requiring manual data collection are rarer. For example, Alizadeh et al. (2020) interview 13 volunteering households who requested their personal data from loyalty program providers, investigating expectations about GDPR and the right to access data. Turning to insurance in particular, Bahşi et al. (2019) conducted interviews with Norwegian insurers and their customers to investigate the effect of GDPR and the NIS Directive on cyber insurance and Dexe et al. (2021) offer a broader interview-based discussion of how the Swedish insurance industry looks at transparency issues from a strategic perspective.

Additionally, there are many studies of GDPR aspects where data can be automatically collected on a large scale (from hundreds or thousands of websites), such as website cookie-consent (Machuletz and Böhme 2020; Sanchez-Rola et al. 2019; Nouwens et al. 2020), app privacy (Momen et al. 2019; Fan et al. 2020), and data portability (Syrmoudis et al. 2021). These studies are similar to ours in the sense that they investigate GDPR effects empirically, but differ considerably in the precise objects of investigation.

To conclude, while the number of empirical studies on the effects of the GDPR in general are growing, the right to meaningful information about the logic involved in insurance has not, to our knowledge, been investigated before, with the exception of our previous article (Dexe et al. 2020) and Sørum and Presthus (2020).

Legal literature on the GDPR right to explanation

With the GDPR being relatively recently adopted, the discussion concerning the interpretation and scope of the right to explanation is still ongoing. An early analysis comes from Temme (2017), who criticises the regulation for insufficiently addressing the challenges to render algorithms more transparent, questions the legal basis of the right to explanation, and points out that even if legally binding, its scope is very limited.

Wachter et al. (2017) are also sceptical about the right to explanation and argue that it does not exist in the GDPR. Instead, they claim, it should be talked about as a right to be informed. Selbst and Powles (2017), more cautiously, argue that the term used should simply be “meaningful information about the logic involved”, exactly as it is worded in the GDPR text. This is the principle followed by us in the title of this paper. Bottis et al. (2019) argue that there is a tension between the right to protection of personal data, on the one hand, and the right of access to information, on the other hand.

Human–computer interaction literature on explanations

In order to be able to determine the reasons behind differences and similarities in the responses, we must also consider the increasingly debated topic of explanation



of automated decision making, within the HCI community. While initially focused on supporting the developers of algorithmic systems to assess the workings of their own systems, recent work has more actively focused on communicating algorithmic trade-offs to the wider public of consumers and citizens (van Berkel et al. 2021).

Of particular relevance to the topic of insurance is prior work by Binns et al., who aim to “find explanation styles which could plausibly meet or exceed the regulatory requirements regarding transparency of automated decisions, in particular the requirement that organisations provide ‘meaningful information about the logic involved’ in an automated decision (GDPR art 15(h))” (Binns et al. 2018, p. 4). The empirical results also suggest that the explanations labeled *sensitivity* and *input influence-based* explanations are more just compared to *case-based explanations*. Especially sensitivity-based explanations resonate with the respondents as actionable explanations—giving suggestions for how to change behaviour to get different results.

Another way to highlight what ‘meaningful information’ can be to look at how explanations work, and to what extent and which explanations are effective. One in-depth look at explanations is by Wilson and Keil (1998), who define an explanation as “an apparently successful attempt to increase the understanding of that phenomenon” (p. 139). However, they note, most explanations that we encounter day to day are surprisingly limited—called the *shallows of explanation* in their terminology. At the end of the paper they ask; if explanations are generally shallow, why do people still seem to accept them? They give a list of four aspects¹ that “give an explanatory sense without yielding precise mechanisms”:

- (1) *Explanatory centrality* is the idea that certain properties are especially important in specific domains. That is, they have explanatory value for many different things in a domain, and would be useless as explanations if they did not have generality. E.g. size matters for many artefacts, location matters for human interaction.
- (2) *Causal power* is the idea that we can understand something not just by listing its properties, but rather by a tendency to behave in certain ways in different situations. For example, “[G]old has a wide array of causal powers that are distinctive to it and [I] expect that any explanation involving it must be in accord with those causal powers.”.
- (3) *Agency and cause* describes the notion that different agency and cause have contextual dependency. For example, “Intentional agency is critical to understanding humans but not earthquakes”.
- (4) *Causal patternings* are related but different from agency and cause and appear as “patterns such as whether causal interactions proceed in serial chains or are massively parallel”.

¹ These aspects have no specific names in Wilson and Keil (1998), and are given names here for the sake of improved intelligibility and ability to refer back to them.



These four aspects give people a way to grasp a problem without knowing the full details of specific problems. The perspectives on explanations offered by Binns et al. (2018) and Wilson and Keil (1998) are employed in the analysis of our empirical results in the “Results” and “Meaningful explanations” sections.

Computer science literature on explanations

The literature review would not be complete without mentioning the growing field of *explainable AI* (XAI) within computer science. Rai (2020) says that the field “explains the rationale for the decision-making process, surfaces the strengths and weaknesses of the process, and provides a sense of how the system will behave in the future.” For the most part, XAI research explores how very large and complex models such as deep neural networks can be approximated by smaller and simpler models such as decision trees. Whereas the former cannot be grasped by humans, the latter are supposedly more understandable, though more empirical research with human subjects is needed (Abdul et al. 2018). Useful reviews of the computer science literature on explanations include Guidotti et al. (2018), Du et al. (2019), Meske et al. (2020) and, as mentioned, Rai (2020). Dellerman et al. (2019) offers a complement to these reviews by highlighting how humans and machines can collaborate to create *hybrid intelligence*.

However, it should also be mentioned here that recent empirical research suggests that even though insurers are convinced about the future value of AI, they are not currently using AI extensively (Dexe et al. 2021). Thus, the relevance of the XAI field to article 15 requests is most probably small at the moment, but can be expected to grow as AI is more widely adopted by insurers.

Method

Since this paper features replications of a method from a previous paper (Dexe et al. 2020), the method will be presented in three steps. First, a discussion about the selection of countries. Second, the standardised approach from the previous paper, which we intend to replicate. Third, the necessary adaptations that have been made in each country.

Selection of countries

While it should be acknowledged that the selection of countries is a convenience sample, driven by the composition of the research team, the sample of the five countries included does cover some interesting differences. First there are differences in digital maturity, with several studies showing Denmark, Finland, The Netherlands and Sweden being among the digital front runners of the EU, with Poland providing an interesting contrast as a less digitised country (European Commission 2020; McKinsey & Company 2020). The difference is particularly striking with respect to AI, where only 28% of Poles had heard or read about AI in the last year compared



to 81% of Dutch and Swedes and 73% for Finns and Danes (European Commission 2017). Second, the insurance markets in Denmark, Finland, and Sweden are a bit more oligopolistic compared to the Netherlands and Poland where the market is less concentrated. Third, the sample includes countries which are relatively small population-wise and one country (Poland) which is bigger in this respect. Fourth, the sample includes a broad range of how long the countries have been part of the EU (and predecessors): The Netherlands are a founding member, Denmark joined in 1973, Sweden and Finland in 1995, and Poland joined in 2004. Fifth, Poland is post-communist, adding yet another kind of diversity to the sample.

Thus, even though the sample of countries is not a random sample which can lay claim to statistical generalisation, it is nevertheless *not* a strikingly homogeneous sample—it does contain a number of potentially meaningful differences, relevant when assessing how much can be generalised from the study.

The standardised approach

The approach can be summarised in the following manner:

- (1) Identify the relevant type of insurance—extant insurance offerings will differ between countries and companies.
- (2) Get an overview of the market—identify main actors and rough market shares (to know which ones are most relevant to include).
- (3) Translate the request.
- (4) Recruit volunteers—making sure that there is only one per company.
- (5) Send requests—volunteers are asked to note date and means of contacting their insurer.
- (6) Gather responses from the volunteers and analyse these.

Insurance schemes will vary between countries depending on a number of variables. The concentration of the market, age (and legacy) of the market actors, the type of homes people have, the ownership structure of those homes, etc. Our specific case study concerns ‘home insurance’, a package of insurance policies related to the home, covering things such as theft and fire (Insurance Sweden [n.d.](#)). This insurance covers slightly different things in each country, which will affect comparability. *Home insurance* was initially chosen for several reasons: (i) multiple market actors offering very similar products, (ii) a straightforward product which is comparable across companies, (iii) it is nearly ubiquitous on the Swedish market (nearly 97% of Swedes have home insurance (Insurance Sweden [n.d.](#)), and (iv) it is based on non-sensitive data which makes it easier for volunteers to share responses (if the study been done on, e.g., health insurance it would probably have been more difficult to recruit volunteers).

We then looked at what the market composition for the specific product. The main reason for this was to reach not only a large number of market actors, but more importantly also a sufficiently large part of the market so as to better represent the



Table 1 Approximate market shares covered in each country

Country	Market share covered (%)
Sweden	90–95
Denmark	68
Finland	72
The Netherlands	45
Poland	40–45

Hi!

In accordance with article 15, section 1h, of the General Data Protection Regulation 2016/679 I would like information on how the premium of my home insurance is determined. This article in the regulation should be applicable if pricing (i) is automated and (ii) is based on personal data (both collected from me and collected by other means).

I would be pleased to receive this information in suitable form (e.g., mathematical formulæ or descriptive text) that meets the requirements of the regulation on meaningful information about the logic involved in automated decision-making. Thanks a lot for your help!

Best regards etc.

Fig. 1 The request for pricing information made to the insurers

hypothetical experience of the population. The approximate market shares thus covered are summarised in Table 1.²

The request, translated into English, is shown in Fig. 1. The actual queries were sent in Danish, Dutch, Finnish, Polish, and Swedish, respectively. Translation was a

² All figures are subject to some uncertainty, as large insurers are separately accounted for in the statistics, while smaller ones are lumped together into ‘other’ categories. In some sets of statistics, numbers for apartments and houses are given separately, forcing interpolation. Swedish figures are based on the official market statistics of Q4 2018 from Insurance Sweden (<https://www.svenskforsakring.se/globalassets/statistik/importerad-statistik/statbranch/branschstatistik/2018/branschstatistik-q4-2018.pdf>). Danish figures are based on Q1 2020 figures from Forsikring & Pension, the Danish industry organization for insurance and pension companies (<https://www.forsikringogpension.dk/statistik/markedetsandele-for-brand-og-loesereforsikring-for-private/>). Finnish figures are based on numbers from Finance Finland in their overview of Finnish insurance companies in 2020 (<https://www.finanssiala.fi/wp-content/uploads/2021/06/Finnish-insurance-in-2020.pdf>). Dutch figures are based on the 2016 market statistics from the Insurers Association (<https://www.verzekeraars.nl/media/3544/verzekerd-van-cijfers-2016-nl.pdf>). Polish figures are based on the official market statistics of Q3 2020 from The Polish Financial Supervision Authority (UKNF) (https://www.knf.gov.pl/knf/pl/komponenty/img/Raport_sektor_ubezpieczen_III_kw_2020_72458.pdf, p. 10).



relatively straightforward step. Since we aim for an interaction with the companies on the level of an above average informed customer, we initially chose to frame it in language that was somewhat ‘formal’ (e.g., pointing to the particular GDPR article) but not specific legalese. That being said, the translations have been made by or together with lawyers in several cases, which also ensures that the relevant legal references are kept in the request.

Recruitment of volunteers was made among colleagues or personal acquaintances of the researchers involved in the study. When a sufficient amount of volunteers had been recruited we proceeded to send one request per insurance company. While sending more requests per company would give a larger sample of how that specific company answers the query, it would also notify the company about the organisation of data requests, and would most probably affect responses. Similarly, approaching the company as researchers might produce a more qualitative answer, but would likely not mimic how they respond to actual customers.

As noted above, the study design is very similar to Sørum and Presthus (2020). Important differences in the approach include (i) the specificity of using home insurance across all requests, (ii) approaching companies in comparable markets but in different countries and (iii) framing the request to only concern automated decisions. This framing, by pointing to a particular section of the GDPR, is also in line with their recommendation for effective use of GDPR rights.

In the following, some country specific adaptations of the methodology are detailed.

Denmark

In Denmark, requests were made in March and April 2021. A total of five Danish companies were contacted by five recruited volunteers, each being a home insurance customer at one of these companies. The five volunteers were recruited through our personal network. One volunteer had home insurance for an apartment, while four lived in houses. All five requests were made digitally; one by e-mail, four through online web forms.

Finland

In Finland, the requests were made between April and May 2021. Four volunteers sent the predetermined request via an online message platform offered by an insurance company, filled a contact form with the query in the customer portal of a company or sent an e-mail request. All of the recruited volunteers had an apartment insurance.

The Netherlands

All requests were sent out in March 2021. The four recruited volunteers were customers at the respective company, and were recruited through our personal network.



Three of the requests were made by mail, whereas one request was sent through an online form. The volunteers live in a variety of apartments and houses.

Poland

The recruited volunteers in Poland sent out their requests between December 2020 and March 2021. The volunteers were customers of six different insurance companies. Two additional companies were approached by our volunteers, but the data collection was not possible in these cases. For one company, the data collection failed due to technical limitations in the contact form. For another company, a problem with identification of the insurance-holder occurred. In the latter case, our volunteer was later informed, upon the telephone contact with the customer service, that the insurance premium was set as a certain percentage of the bank loan, without any further specifications.

The recruitment was made partly through our personal network, and partly with the help of the non-governmental organisation Panoptykon.

The volunteers were instructed to send their requests in the most convenient way, which meant sending an email or filling the request in a contact form. The volunteers had an apartment or house insurance, sometimes in conjunction with other forms of insurance.

Sweden

The seven requests from Sweden were made from December 2018 to March 2019. Recruitment of volunteers was mainly done among colleagues at a research institute, and one non-researcher. For the Swedish study we did not record how requests were made, but did record how the responses came in. Also, three of the seven volunteers had home insurance for apartments, while four had for houses. The house insurance has broader coverage, but there was no obvious difference in the responses for the different types of insurance.

Results

In the following, we present the results from the requests made. In Table 2 we show data related to how the pricing algorithm works, or at least what the companies claim affect the price. The categories are derived from the original analysis of the Swedish replies. In Table 3 data related to how the companies reply to the queries are presented. As with the previous table, the categories are derived from the Swedish replies. In both cases we have had a discussion about whether other groupings were possible, but have chosen not make significant changes.³ Finally, Fig. 2 shows

³ The only alteration is the addition of “No response” in Table 2 to indicate contacted companies that failed to respond, and the removal of “Visualisations” from Table 3 since none of the replies contained visualisations. The reason for including it in (Dexe et al. 2020) was due to a discussion about ways to improve intelligibility of the responses, but that discussion is not included in the present paper.



Table 2 Overview of pricing information obtained, by company

	Living area (m ²)	Family status (e.g., number of people)	Address	Real estate data	Age	Deductible	Indemnity limit	Income and financial data	Security measures (e.g., locks)	Age of policy (loyalty)	Claims history	No response
DK1	L		L	X	X	X	L			X	X	
DK2	L		L	L	X	X	L	X		X	X	
DK3		X	X	X	X	X	X			X		
DK4			X	X	X							
DK5			X				X					
FI1	X	X	X		X	X						
FI2	X		X		X							
FI3	X		X		X	X	X		X			
FI4												X
NL1				X						X		
NL2									X			
NL3	X	X	X	X	X					X		
NL4												
PL1			X	X	X			X			X	
PL2												
PL3	X		X	X							X	
PL4			X	X			X		X	X	X	
PL5	X		X	X			X		X	X	X	
PL6												X
SE1	X	X	X		X							
SE2	X	X	L	X	L	X	X	X	X	L	X	
SE3			L		L							
SE4		X	X	X				X	X		X	
SE5			X		X							



Table 2 (continued)

	Living area (m ²)	Family status (e.g., number of people)	Address	Real estate data	Age	Deductible	Indemnity limit	Income and financial data	Security measures (e.g., locks)	Age of policy (loyalty)	Claims his- tory	No response
SE6	X			X								
SE7		X	X		X					X	X	

X denotes that the type of data was mentioned; L that the type of data and the logic of its use was mentioned. Real estate data is only applicable to house insurance. 'No response' (FI4 and PL6) means that companies did not respond at all, whereas an empty row (PL2, NL2, and NL4) means that a response was received but that it did not reveal any pricing variables



Table 3 Overview of procedural information obtained, by company

	Process description	Fairness	Legal basis	General logic	Contact details	Information on other customers	Business confidentiality	How a profile is created
DK1	X		X	X	X			
DK2	X		X	X	X			
DK3	X				X			
DK4	X		X		X		X	X
DK5	X				X			
FI1	X			X	X			
FI2							X	
FI3	X			X	X			X
FI4								
NL1	X		X	X	X		X	
NL2	X				X			
NL3	X			X	X		X	
NL4	X		X	X			X	
PL1							X	
PL2			X	X			X	
PL3			X				X	
PL4			X	X	X		X	
PL5								
PL6								
SE1								
SE2	X	X				X		X
SE3		X						
SE4	X				X			
SE5	X				X			
SE6	X		X	X	X		X	X
SE7		X		X	X	X	X	

a representative insurer response (translated into English). The rationale for the figure is twofold: (i) it offers a concrete example of what a response looks like and (ii) it illustrates the coding regime employed.

Denmark

The time duration between sending the request to getting the final answer varied from 6 to 22 days, with a median of 17 days. The length of the response varied from 45 to 661 words. Two of the replies had additional attachments describing in depth the insurance company’s personal data policy (a 10-page and a 20-page attachment).

Three companies (DK2, DK4, and DK5) provided generic examples of the data categories used to calculate the home insurance without referring to the specific



Hi!

You have requested that SE7 inform you about how we calculate the premium for your home insurance.

SE7 uses all the data we have access to. It is about, e.g., information about you and your household. Where do you live,¹ how many are there in your household,² how old are you,³ how long have you had insurance in SE7,⁴ how many claims do you have.⁵ But also about other information about other customers, e.g., how many claims come from a certain residential area.

Information on other customers

The purpose is to calculate a premium that is as fair to each customer as possible in relation to the risk. It is not possible to set completely individual premiums because the idea of insurance is to spread the risks over a collective.

Fairness

Not all people suffer injuries, but if an injury occurs, it can be costly if you have to pay for everything yourself. When a collective bears the overall risk, the cost for each individual in the collective is lower. It will be a win-win situation for policyholders and the insurance company.

General logic

We do not provide information on exactly how e.g. what the actuarial formulæ look like. This is a business strategic and critical information that each insurance company keeps to itself. We have no obligation to disclose that information.

Business confidentiality

If you have any questions/comments on the above, youre welcome to contact us. You have all my contact details below.

Contact details

Fig. 2 A representative insurer response (translated into English) with the company named replaced by a pseudonym. The footnotes correspond to the coding of pricing information as summarised in Table 2: 1 Address, 2 Family status, 3 Age, 4 Age of policy, 5 Claims history. The annotations to the right of the text correspond to the coding of procedural information as summarised in Table 3

categories used to calculate the pricing of the enquirer. One company gave three category examples (address, area of living, and price of insurance). In contrast, the two others provided more elaborate examples, e.g., insurer's age, postcode, the value of household effects, type of home (apartment, townhouse or house), and real estate



characteristics (roof, basement, apartment floor). One company (DK4) responded that the price of home insurance was automatically generated but did not disclose the calculation formulae, arguing that such insights are “classified as trade secrets” referring directly to Chapter 6, Article 22, Section 1 of the GDPR. DK2 argued that they do not use automated decision-making when settling on insurance pricing. Nonetheless, they still provided a detailed explanation of the general logic behind their pricing calculations and examples of how different variables could potentially influence these. Their explanation resembles the “input influence-based” explanations: “For example, if you living in a neighbourhood where the risk of burglary is small, we offer you a lower price on your insurance than if you lived in a neighbourhood with many burglaries”. DK5 stated they could not share their “actuary’s mathematical computation models”, giving no reason as to why.

DK1 and DK3 disclosed more detailed information on the different variables used in their calculation of home insurance. They specified which personal data was used when calculating the customer’s insurance price. For instance, DK3 gave examples such as “Number of people: 2, Roofing material: hard roof”, etc.. Both companies disclosed the exact price reductions that were part of the pricing calculations. One of these companies (DK3) provided no further explanation on their process or logic of calculating prices (automated or not). The other company (DK1) gave a general description (resembling “input influence-based explanation”) of the company’s algorithmic calculation stating that calculating insurance price is “based on data and associated statistical models. The higher/lower we assess the risk of getting a damage claim during the insurance period, the higher/lower the price can be for the insurance”. They further exemplified this with how different variables could influence price (living factors such as probabilities of flooding and burglaries in the area compared to others).

Finland

For Finland companies responded in 9 days (FI1), 4 days (FI2), and 23 days (FI3). One company (FI4) failed to respond to the request.

FI1 responded with information about the pricing methods used in calculating home insurance prices. The answer explained that any home insurance pricing is based on variables like living area of the apartment, location of the apartment, size of the city where the apartment is located, the age of insurance holder, and the selected amount of deductible. The response also referred to the terms and conditions document of the company’s home insurances: “[T]he household insurance product description also describes the determination of the insurance premium in home insurance. You can find the product description in the appendix to this message.”

FI2 stated that they are not using automated decision-making for calculating home insurance pricing. The company claimed that “[p]ricing is not a decision, but an offer of the price of a product sold to a customer. The exact determination of the price is a matter of business secrecy.” In addition, they stated that automated decision-making is used only when deciding whether they can grant the insurance or not



and when deciding insurance compensations. However, they revealed that the living area of the customer and the location of the apartment are major variables influencing the premium calculation of the exact type of home insurance the informant has (a discounted home insurance offered for people under 27 years of age).

FI3 said that automatic processes are used for price calculation. Still, decision-making is not considered automatic because an insurance officer takes part in the calculation and makes the final decision. In addition, the FI3's answer included a separate remark about the use of profiling: “[p]rofilng is used, for example, to determine the risk equivalence of the insurance price”. The company did not clarify further how this profiling method affects the price calculation of home insurance.

The Netherlands

The time between sending the request for information and the final answer was 4 to 44 days, with a median of 7.5 days. Response length varied between 81 and 271 words (excluding lengthy attachments).

One company provided examples of the data categories that were relevant in the calculation of the price, including both data collected through public information (the living area, year of construction, and type of home) as well as information provided by the customer (year of birth, family composition). The other companies did not provide such specifics and referred to business confidentiality. One company assured that the use of this information would be exactly the same ‘if done manually by a person’.

All companies further stressed that their operation was in line with existing regulations, oftentimes pointing to additional information provided in attached documents. One company detailed the most relevant points in relation to GDPR being the lawful basis for the data processing and data minimisation. While all companies stated that they met these requirements, details provided were minimal. For example, one company highlighted that a retention policy has been established within the company to ensure that data was not kept for longer than necessary, but details on this policy were missing.

Poland

The replies given by five of the companies were between 78 and 301 words in length. The time interval between the request and the answer varied, in three cases, between a week and a month. Although the time-frame provided for the response by the legislator should normally not exceed 1 month, two of the companies were late with the response (1 week for PL1 and 2 weeks for PL5 respectively), and one company did not respond (PL6) to the request, as of 3 months after the time provided by the legislator.

The content of the responses varied significantly. Four companies provided examples of data categories taken into account in calculating the insurance premiums. One company stated that the decision made was not fully automated, since automation only had a supportive role. As for how the premiums are set, some of the



companies explained that the premiums are determined on the basis of the currently applicable tariffs, but almost all companies emphasised that the details of the tariffs are subject to business confidentiality rules. Additionally, in two cases, the responses contained a reference to national laws governing the insurance sector along with the clarification that insurance companies are obliged to provide the supervisory authority with information on the tariffs and the grounds for their determination. Aside from the six companies presented in Tables 2 and 3, two additional companies were contacted, but are not included in the results due to various problems with the data collection.

Sweden

The time from request to response varied from 2 h to 2.5 months. Response length varied from 50 to 600 words. More details about the Swedish results can be found in (Dexe et al. 2020).

The most frequently provided information was address, age, and family status as contents of the insurance. Two companies provided descriptions of the logic behind the decisions, with examples of how certain variables were used. For instance, SE2 listed a number of variables that might increase premiums (“if you’re young and belong to a category of persons who statistically experience more damages or if you live in an area where the risk of damages for things and persons are higher than in other areas”) as well as variables that might lower premiums (“if you have been a customer with us for a long time without any claims”).

As for the procedural information, the most frequent type of information was process descriptions and contact details. One company offered no procedural information at all. Two offered descriptions of the logic of insurance in general. Some companies offered legal explanations for why they did not disclose all the information, referring either to business confidentiality or other national regulations as well as their specific interpretations of the GDPR. However, these companies were not the least forthcoming in the responses.

Discussion

Similarities and differences

In this section we discuss similarities and differences between the replies, before looking deeper into interpretations of the legislation and then to a section where we discuss how the responses work as explanations.

For a study such as this one, a point of comparison within our sample is differences in practice between the different countries. As can be seen in “[Limitations and future work](#)”, there is no shortage of plausible explanations for such differences. Level of digitisation, how long the countries have been apart of the EU project, when the requests were made, and others. However, no stark differences appear when looking at the results. The Danish companies seem to be more forthcoming



and exhaustive when listing variables, but the sample is too limited to say conclusively that such a country profile exists. One could argue that this is as it should be when all the companies approached are subject to the same legislation.

Another point of comparison would be within-country variety. One would assume that size, age and ownership structure could have an effect on the replies. Specifically, that a larger company would have more resources and routine to give a better answer, an older company might have experience and loyal customers to maintain that might improve the replies, and mutually owned companies would have less incentives to keep in information from their customers that in effect also own the company. However, no such differences show up. There is one start-up in the sample, which is neither less or more forthcoming than significantly larger and older competitors. In the Netherlands, the two largest companies sampled have the fewest marks for pricing information in Table 2, but this trend does not appear in other countries. Finally, there does not seem to be any difference between mutually owned insurance companies and other companies.

There are still interesting variations though. Whether automated decision-making has occurred, and the significance of business confidentiality are two that are especially noteworthy.

In the replies there are different interpretations of what constitutes automated decision-making. Some companies simply respond to the queries, and thereby either confirm the validity of the query by admission or non-negation, or show that they do not understand the full extent of the request. Others deny the existence of automated decision-making in their system. These companies either flat out deny it, or try to explain their process as one where the actual decision is not automated.

The latter argument is the most interesting. While we have not confirmed the extent of automation in the companies' processes, the arguments deserve further consideration. As mentioned in the results from Denmark, one company claimed to not have automated decision-making, but still listed several variables and the logic behind the decision. In other words, they seem to hint at an algorithmic approach to the decisions. Another company, mentioned in the results from Poland, says that the automated system only had a supportive role—it had an “auxiliary, not a decisive, function”. A Finnish company said that the pricing is an offer delivered for the customer, not a decision per se. One interpretation of these examples is that the companies use different communicative strategies to avoid the specific requirements in the GDPR. Another is that they are simply true statements.

Another variation is that the contacted companies often pointed to business confidentiality (Table 3) as an obstacle to providing further details on the way in which customer data is used in price-setting. For example, FI2 states that “The exact determination of the price is a matter of business secrecy”, and PL2, similarly, claims that “In response to the application, we would like to inform you that the premium is determined based on the currently applicable tariff. The details of the tariff are an insurance secrecy.” Yet, somewhat paradoxically, it is not necessarily the case that the companies referring to business secrecy are less forthcoming than those that do not. There are companies that refer to business confidentiality in an obstructive manner, but again, there are also those that only refer to business confidentiality for specific parts of the decision-making process and are otherwise forthcoming. Thus,



one possible explanation for this somewhat counter-intuitive result is that they are transparent with respect to a significant portion of the query because they understand the question, and refer to business confidentiality only for a limited amount of information.

It is also interesting to observe that there is the notable absence of marks for indemnity limits, which would seem to be an important aspect for the price of home insurance. While location and age are well represented, only Denmark has more than two companies that talk about indemnity limits. DK2 says that “if you own many valuable items, your price will be higher”. while DK1 instead says that you can get a lower price if “you have a low sum (content) insured”.

Finally, there is a point to be made about the relationship between trust and the business of insurance. There is a possibility that insurance companies, being in the ‘business of trust’, will act differently than other types of businesses in order to maintain the trust of the consumers. One answer for why the responses are vague could be that the automated-decision making the request references may not reach the threshold for legally mandated right of access, as we will see in the next section. In such a case, insurance companies have responded to requests despite no legal requirement forcing them to do so. That would mean that even if the responses are not as detailed as the consumer would like, they are more transparent and forthcoming than what the law compels the companies to be.

Different interpretations and language versions of Article 15(1)(h) GDPR

One observation made in the course of the study was that the wordings of Article 15(1)(h) GDPR are subtly different in different languages, as illustrated in Fig. 3.

More precisely, the languages differ in how they refer to Article 22 GDPR. In the following, these differences are explained using the English version of Article 15(1)(h). The main question of interpretation is what the phrase ‘at least in those cases’ refers to. Below, we suggest two types of interpretations identified in our discussions:

- The narrow interpretation of Article 15(1)(h) implies that the phrase ‘at least in those cases’ relates to the ‘automated decision-making, including profiling, referred to in Article 22(1) and (4)’ as a whole phrase. This interpretation creates a legal obligation on the part of entities processing personal data to provide ‘meaningful information about the logic involved, as well as the significance and the envisaged consequences of such processing for the data subject’, whenever the automated decision-making or profiling fall within the scope of Article 22 GDPR. Article 22 GDPR governs fully-automated decision-making and profiling, which produce legal or similarly significant effects on individuals subject to the exceptions in Article 22(2) GDPR. Following this interpretation, the question of when the controller needs to provide such ‘meaningful information’ will depend on whether the automated decision is ‘fully-automated’, whether it ‘produces legal effects’ or ‘similarly significantly affects’ the individual, and whether any of the exception in Article 22(2) would apply.



*Article 15***Right of access by the data subject**

1. The data subject shall have the right to obtain from the controller confirmation as to whether or not personal data concerning him or her are being processed, and, where that is the case, access to the personal data and the following information:

[(a)–(g) omitted]

English:

(h) the existence of automated decision-making, including profiling, referred to in Article 22(1) and (4) and, at least in those cases, meaningful information about the logic involved, as well as the significance and the envisaged consequences of such processing for the data subject.

Danish:

(h) forekomsten af automatiske afgørelser, herunder profilering, som omhandlet i artikel 22, stk. 1 og 4, og som minimum meningsfulde oplysninger om logikken heri samt betydningen og de forventede konsekvenser af en sådan behandling for den registrerede.

Finnish:

(h) automaattisen päätöksenteon, muun muassa 22 artiklan 1 ja 4 kohdassa tarkoitetun profiloinnin olemassaolo, sekä ainakin näissä tapauksissa merkitykselliset tiedot käsittelyyn liittyvästä logiikasta samoin kuin kyseisen käsittelyn merkittävyys ja mahdolliset seuraukset rekisteröidylle.

Dutch:

(h) het bestaan van geautomatiseerde besluitvorming, met inbegrip van de in artikel 22, leden 1 en 4, bedoelde profilering, en, ten minste in die gevallen, nuttige informatie over de onderliggende logica, alsmede het belang en de verwachte gevolgen van die verwerking voor de betrokkene.

Polish:

(h) informacje o zautomatyzowanym podejmowaniu decyzji, w tym o profilowaniu, o którym mowa w art. 22 ust. 1 i 4, oraz – przynajmniej w tych przypadkach – istotne informacje o zasadach ich podejmowania, a także o znaczeniu i przewidywanych konsekwencjach takiego przetwarzania dla osoby, której dane dotyczą.

Swedish:

(h) Förekomsten av automatiserat beslutsfattande, inbegripet profilering enligt artikel 22.1 och 22.4, varvid det åtminstone i dessa fall ska lämnas meningsfull information om logiken bakom samt betydelsen och de förutsedda följderna av sådan behandling för den registrerade.

Fig. 3 GDPR Article 15, Section (1)(h), in different languages

For all other cases of automated decision-making and profiling, the provision of the ‘meaningful information...’ would be recommended in light of the principle of transparency in Article 5(1)(a), but would not be mandatory.

- The broad interpretation suggest that ‘at least in those cases’ could refer to ‘automated decision-making’, meaning that profiling in scope of Article 22 GDPR would be regarded as a non-exhaustive example of processing covered by this provision, and that meaningful information is equally applicable to automated decision-making in general. This would require data processors to provide individuals with ‘meaningful information about the logic involved...’ in all cases of automated decision-making and profiling, irrespective of whether the decisions would be caught by the scope of Article 22 GDPR. Recital 63 GDPR would seem to support this reading, even broaden-



ing the requirement further by referring to the word ‘processing’: “[e]very data subject should therefore have the right to know and obtain communication in particular with regard to (...) the logic involved in *any* automatic personal data processing and, at least when based on profiling, the consequences of such processing” (emphasis added). Moreover, this line of interpretation seems to be supported by the European Data Protection Board (EDPB), a body established by the GDPR to ensure the consistent application of the Regulation. The Board indicates in the Guidelines on automated individual decision-making and profiling that Article 15(1)(h) entitles data subjects to receive the information about existence, automated decision-making, including profiling, meaningful information about the logic involved, and the significance and envisaged consequences of such processing for the data subject (Article 29 Data Protection Working Party 2018).

The comparison between the different language versions of Article 15(1)(h) of the countries involved in the study revealed that some of the versions could suggest different interpretations identified above. The consensus among the researchers is that the English language version suggests a more narrow interpretation in comparison with some of the other versions.

In the Swedish version, the lack of the comma after the word ‘profilering’ (Eng. ‘profiling’), suggest that a broad interpretation is more plausible. This was assumed—not considering the possibility of different interpretations—by the researchers when the request in Fig. 1 was first worded and sent to Swedish insurers in the original study (Dexe et al. 2020). Only in the work with the international study did the possibility of the narrower interpretation become clear. The Dutch version seems to closely follow the English version of the text, making the interpretational nuances similar to the English version. The same goes for the Danish version, although the use of ‘som minimum’ (Eng. ‘as a minimum’) instead of ‘at least in those cases’ may open up for a broader interpretation.

The Polish text, however, hints at the reading in accordance with the broad interpretation, since the expression ‘o którym mowa’ (Eng. ‘referred to’) would point to the word ‘profiling’, while the phrase ‘istotne informacje o zasadach ich podejmowania’ would pertain to the phrase ‘decision-making’. In the Finnish version, the Article 15(1)(h) includes the wording ‘muun muassa’ (Eng. ‘among other things’) before referring to the article 22, and afterwards the phrasing ‘sekä ainakin näissä tapauksissa’ (Eng. ‘and at least in these cases’; note the word these instead of those, highlighting that the mentioned cases are referred to as possible cases among others) indicate that the Article 22 is given as a non-exhaustive example, which suggests a broad interpretation.

It should be stressed that as a regulation, the GDPR is directly applicable in all the Member States. No particular language version takes precedence over any other, as all the official languages of the Union have an equally authentic character. Thus, in the absence of case law, it is difficult to say what the effects, if any, of the difference noted above are. However, it is an important observandum in this context, where the wordings in the different national languages may have affected how insurers respond to the requests made.



Meaningful explanations

Using the literature in “Related work”, we can start discussing how to view the explanations given in the responses.

The idealised explanation styles that are common in the academic literature could be an example of what Wilson and Keil (1998) call the *shadows of explanation*, as in the narrow focus of academics that explain a phenomenon in depth but leaves out other aspects that might also affect the phenomena. In reality, they argue, explanations are often found in the *shallows of explanation*, the ‘surprisingly limited everyday’ explanations. Wilson and Keil try to motivate why humans tend to accept and have explanatory sense of a thing despite the explanations being poor, identifying the four aspects *explanatory centrality*, *causal power*, *agency and cause*, and *causal patternings*. For instance, *explanatory centrality* suggests that most people will consider aspects such as location, age, property value and propensity for damage to be central in determining risk for damages. Therefore those data points also have a larger explanatory value than others, and may not need more in depth reasoning about why and how they are used. Furthermore, *causal power* lets us assume that in the context of insurance, customers will be able to make inferences about how the disclosure of a specific data point works within the context of insurance—i.e. moving the price up and down depending on the associated risk.

These aspects could explain both similarities and differences between the responses. It suggests an explanation for the fact that most responses are fairly vague, and seem to omit what a critical reading would hold to be important data points. If, as Wilson and Keil suggest, people and companies generally apply the aspects above, then it may be that the companies assume a level of knowledge about what holds explanatory centrality or causal power in the field of insurance, i.e., that consumers know something about the insurance(s) they have. In that case, the explanations can hold more explanatory value than one would be able to get from a reading of the responses where we assume no prior knowledge.

Second, it suggests an explanation for some of the differences in that if companies can assume that customers have prior knowledge and are aware of the context in which the requests for information were sent, then the companies have less reason to explain every single data point that is included in the pricing algorithm. They can therefore choose a few examples that could serve as a suggestion for the full explanation, and the differences in responses can arise from companies simply using different examples.

Now, we can probably also assume that companies cannot use their assumption of consumers’ prior knowledge as a legal defence for framing their information in a certain way. The aspects above are not meant to provide ‘meaningful information’, and are not an appropriate way to explain the ‘logic behind’ the decisions made. But they force us to consider that the in-depth account of all the mechanisms of the pricing algorithm is not the sole way to increase understanding, and that even incomplete accounts can be acceptable to consumers. Still, being able to infer certain explanatory value from the aspects listed above requires that the consumer has prior knowledge of insurance or what a pricing algorithm does. It may be that such prior knowledge is also a prerequisite for even wanting to ask the question—as in



someone who does not have any knowledge about algorithms would not even know to ask about them, but that is a weak justification for the responses given, at best.

Wilson and Keil conclude their paper by proposing the following about the value of shallow explanations.

To carry the shallows metaphor further, we know that we cannot dive infinitely deep or even stay in any depth for very long; but by using appropriate public charts and supporting frameworks, we can make occasional brief and directed dives, sometimes of surprising depth. So also, in focused and limited ways, we can go to extraordinary explanatory depths with the help of public charts and frameworks of knowledge. But we could never possibly stay at such depths at all times across all domains. (Wilson and Keil 1998)

That is to say, giving full, in depth, explanations to the requests asked in this study might not always be a reasonable approach by the companies. We can still claim that for the most part, the responses by the companies are lacking—but people are unable to always stay in the deep end of the explanatory pool and may be perfectly content with somewhat shallow responses.

As for the explanatory styles presented in Binns et al., the presented results (see Table 2) suggest an overwhelming use of input-based explanation, detailing (in part) the variables that are considered in setting the premiums. However, the participating companies rarely explained the weight with which these variables are considered in the calculations, thereby failing to meet the influence part of *input influence-based* explanations. We did not find any example of explicitly demographic-based, case-based, or sensitivity-based explanation. This indicates a high level of agreement within the insurance market as to what is deemed an appropriate explanation. This could serve to further explain the similarities in style, and perhaps the similarly lacklustre quality of the responses.

Limitations and future work

We next discuss the limitations of our study and opportunities for future work. First, there is a *company-related reliability* aspect: the companies selected do not represent the entire market in each country. As explained in “[Method](#)”, the selection was done in order to cover as large a market share as possible. If additional requests were made in the selected countries, drawn at random from the populations of home insurance policies in each country, chances are high that the corresponding companies are already included in the study, as seen in Table 1. In the case of Sweden, for example, this probability of prior inclusion is some 90–95 %.

Further, we highlight a limitation in the *customer-related reliability* aspect: for each insurer, a single customer was selected to make the request. A convenience sampling was used—volunteers were approached from acquaintances and the networks of the researchers. Now, if companies treat all their customers in the same way, this does not affect the reliability. Any one customer is as good as any other. However, if companies treat their customers very differently, then a response obtained from a particular company by a particular volunteer is not necessarily



representative. Unfortunately, the obvious remedy—having several volunteers approach each company—has an undesirable observer effect, where several identically worded requests within a short time-frame would likely skew responses. Due to this constraint we were unable to assess the consistency of the responses of the contacted companies.

A related limitation focuses on the *temporal reliability aspect*: as described in “Method”, the Swedish data collection was done earlier than the data collection in the other countries. While we do not believe this limitation to have a major effect on the outcomes, our results have to be interpreted in light of this.

There is also reliability aspect related to the particular type of insurance policy chosen: home insurance. As remarked in “The standardised approach”, this type of insurance policy, based on comparatively non-sensitive data, was deliberately chosen to facilitate the recruitment of volunteers. It cannot be known with certainty whether the insurers’ responses to article 15 requests would have been different if the requests concerned another type of insurance policy, such as health insurance. However, there is no reason to believe that responses would be *more* forthcoming, e.g., including mathematical formulæ in such a case. While it is possible that responses would be *less* forthcoming (e.g., containing fewer pricing categories in Table 2), the material obtained does not really support any hypotheses about particular differences to be expected.

Finally, there is an *inter-researcher reliability aspect*: none of the authors have the language skills to assess of all of the responses received, and in order to give the volunteers the maximum amount of privacy compatible with the research design, responses were not shared with the full group of authors but remained in the custody of the researchers who recruited the volunteers. Thus, the coding in Tables 2 and 3 is dependent on each researcher’s interpretation of the coding criteria. To mitigate this risk, several discussion meetings were conducted to reach consensus about the proper coding, but again, results have to be interpreted in light of this.

Despite these limitations, we believe that the validity of the study is good: RQ1 and RQ2 are about the disclosure practices of insurers in the EU when asked by customers about the logic behind automated decisions. This is precisely what was investigated: real customers made real requests to real companies. (It was *not* that researchers approached companies—or industry groups—in their capacity of researchers, or with merely hypothetical requests.) Exactly the object of study was studied.

Future work may consider an evaluation of GDPR right to meaningful information across different industries to identify differences in established practice. Although outside of the scope of our study, an interesting opportunity for future research is to evaluate the perceptions of customers to the answers received.

Conclusions

With respect to RQ1, address is the most common data type disclosed. Other common data types include living area, real estate data, age, indemnity limit, age of the insurance policy, and claims history.



With respect to RQ2, differences between companies do not seem to follow any simple pattern, e.g., of nationality. There is a tendency that the Danish companies disclose a few more variables and are a bit more forthcoming in explaining the logic of data use, and that Polish and Dutch companies often claim business confidentiality, but these observations must be interpreted with some caution, due to the reliability aspects already discussed. Other factors, such as size of company or type of ownership do not seem to have any qualitative effect either.

There are no responses that can be said with any confidence that they fail to comply with the GDPR, except for FI4, PL6 and SE3 where the requirement of ‘undue delay’ must reasonably be considered violated.

With respect to RQ3, we have shown that one complicating factor for companies and other actors is the fact that the requirement for meaningful information about algorithmic decisions is subject to differences in interpretation—both in terms of how the paragraph is phrased (in any language) and the fact that the different language versions seem to point to different interpretations. This creates uncertainty in the market, and until we have clarifications from either courts or legislators it will remain unclear.

For the most part, the responses only give partial explanations for what goes into the pricing algorithm and not one seems to contain an exhaustive explanation for all data points that are reasonably part of an automated decision regarding pricing in home insurance. As argued by Wilson and Keil, there can be acceptable explanations that are not exhaustive or ‘deep’. For the most part, we accept shallow explanations due to our prior knowledge and intuition of the causes and contexts that we see the explanations in. We can give explanatory value to statements that actually do not show the inner workings of a mechanism, and can be satisfied with that explanation. However, that does not mean that such an explanation is acceptable in the eyes of the law, or that it is right or just for a company to assume such prior knowledge when responding to requests for meaningful information.

Acknowledgements J. Dexe and U. Franke were supported by Länsförsäkringsgruppens Forsknings- & Utvecklingsfond, Agreement No. P4/18. K. Söderlund was supported by the Wallenberg AI, Autonomous Systems and Software Program-Humanities and Society (WASP-HS) funded by Marianne and Marcus Wallenberg Foundation and the Marcus and Amalia Wallenberg Foundation. J. Vaiste was supported by the Jenny and Antti Wihuri Foundation. K. Söderlund would like to thank the non-governmental organisation Panoptykon for assistance with finding the Polish volunteers for the study. The assistance of all the volunteers in the five countries is also gratefully acknowledged.

Funding Open access funding provided by RISE Research Institutes of Sweden.

Declarations

Conflict of interest On behalf of all authors, the corresponding author states that there is no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the



material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>

References

- Abdul, A., J. Vermeulen, D. Wang, B.Y. Lim, and M. Kankanhalli. 2018. Trends and trajectories for explainable, accountable and intelligible systems: An HCI research agenda. In *Proceedings of the 2018 CHI conference on human factors in computing systems*, pp 1–18. <https://doi.org/10.1145/3173574.3174156>.
- Alizadeh, F., T. Jakobi, A. Boden, G. Stevens, and J. Boldt. 2020. GDPR reality check—Claiming and investigating personally identifiable data from companies. In *2020 IEEE European symposium on security and privacy workshops (EuroS PW)*, 120–129. <https://doi.org/10.1109/EuroSPW51379.2020.00025>.
- Article 29 Data Protection Working Party. 2018. Guidelines on Automated individual decision-making and Profiling for the purposes of Regulation 2016/679. WP251, adopted on 6 February 2018.
- Bahşi, H., U. Franke, and E. Langfeldt Friberg. 2019. The cyber-insurance market in Norway. *Information and Computer Security* 28 (1): 54–670. <https://doi.org/10.1108/ICS-01-2019-0012>.
- Binns, R., M. Van Kleek, M. Veale, U. Lyngs, J. Zhao, and N. Shadbolt. 2018. 'It's reducing a human being to a percentage': Perceptions of justice in algorithmic decisions. In *Proceedings of the 2018 CHI conference on human factors in computing systems, ACM, CHI '18*, 1–14. <https://doi.org/10.1145/3173574.3173951>.
- Bottis, M., F. Panagopoulou-Koutnatzi, A. Michailaki, and M. Nikita. 2019. The right to access information under the GDPR. *International Journal of Technology Policy and Law* 3 (2): 131–142. <https://doi.org/10.1504/IJTPL.2019.104950>.
- Bradford, A. 2020. *The Brussels Effect: How the European Union rules the world*. Oxford: Oxford University Press. <https://doi.org/10.1093/oso/9780190088583.001.0001>.
- DAC Beachcroft. June 2016. *GDPR deep dive: Profiling in the insurance industry*. Technical report. DAC Beachcroft. <https://sites-dacb.vuturvev.com/110/3572/landing-pages/jade-rhiannon-gdpr-deep-dive--profiling-in-the-insurance-industry.asp>. Accessed 14 June 2021.
- Dellerman, D., P. Ebel, M. Söllner, and J.M. Leimeister. 2019. Hybrid intelligence. *Business and Information Systems Engineering* 61 (0): 637–643. <https://doi.org/10.1007/s12599-019-00595-2>.
- Dexe, J., J. Ledendal, and U. Franke. 2020. An empirical investigation of the right to explanation under GDPR in insurance. In *Trust, privacy and security in digital business. The 17th international conference on trust, privacy and security in digital business—TrustBus 2020*. Springer. https://doi.org/10.1007/978-3-030-58986-8_9.
- Dexe, J., U. Franke, and A. Rad. 2021. Transparency and insurance professionals: A study of Swedish insurance practice attitudes and future development. *The Geneva Papers on Risk and Insurance: Issues and Practice*. <https://doi.org/10.1057/s41288-021-00207-9>.
- Du, M., N. Liu, and X. Hu. 2019. Techniques for interpretable machine learning. *Communications of the ACM* 63 (1): 68–77. <https://doi.org/10.1145/3359786>.
- European Commission. 2017. *Attitudes towards the impact of digitisation and automation on daily life*. Special Eurobarometer 460, March 2017. <https://europa.eu/eurobarometer/surveys/detail/2160>. Accessed 14 June 2021.
- European Commission. 2020. *The Digital Economy & Society Index (DESI) 2020*. <https://ec.europa.eu/digital-single-market/en/desi>. Accessed 10 June 2021.
- Fan, M., L. Yu, S. Chen, H. Zhou, X. Luo, S. Li, Y. Liu, J. Liu, and T. Liu. 2020. An empirical evaluation of GDPR compliance violations in Android mHealth apps. In *2020 IEEE 31st international symposium on software reliability engineering (ISSRE)*, 253–264. <https://doi.org/10.1109/ISSRE5003.2020.00032>.
- GDPR. 2016. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the Protection of Natural Persons with Regard to the Processing of Personal Data and on the Free Movement of Such Data, and Repealing Directive 95/46/EC (General Data Protection Regulation).



- Official Journal of the European Union (OJ) L* 119 (4.5): 1–88. <http://data.europa.eu/eli/reg/2016/679/oj>.
- Guidotti, R., A. Monreale, S. Ruggieri, F. Turini, F. Giannotti, and D. Pedreschi. 2018. A survey of methods for explaining black box models. *ACM Computing Surveys* 51 (5): 1–42. <https://doi.org/10.1145/3236009>.
- Insurance Sweden. n.d. *Insurance in Sweden 2010–2019*. Technical report. Insurance Sweden. <https://www.svenskforsakring.se/globalassets/engelska/statistics/insurance-in-sweden-2010-2019.pdf>. Accessed 14 June 2021.
- Machuletz, D., and R. Böhme. 2020. Multiple purposes, multiple problems: A user study of consent dialogs after GDPR. *Proceedings on Privacy Enhancing Technologies* 2: 481–498.
- McKinsey & Company. 2020. *How nine digital front-runners can lead on AI in Europe*. Report. McKinsey & Company. <https://www.mckinsey.com/business-functions/mckinsey-digital/our-insights/how-nine-digital-front-runners-can-lead-on-ai-in-europe>. Accessed 14 June 2021.
- Meske, C., E. Bunde, J. Schneider, and M. Gersch. 2020. Explainable artificial intelligence: Objectives, stakeholders, and future research opportunities. *Information Systems Management* 0 (0): 1–11. <https://doi.org/10.1080/10580530.2020.1849465>.
- Momen, N., M. Hatamian, and L. Fritsch. 2019. Did app privacy improve after the GDPR? *IEEE Security and Privacy* 17 (6): 10–20. <https://doi.org/10.1109/MSEC.2019.2938445>.
- Nouwens, M., I. Liccardi, M. Veale, D. Karger, and L. Kagal. 2020. Dark patterns after the GDPR: Scraping consent pop-ups and demonstrating their influence. In *Proceedings of the 2020 CHI conference on human factors in computing systems*, 1–13. <https://doi.org/10.1145/3313831.3376321>.
- Rai, A. 2020. Explainable AI: From black box to glass box. *Journal of the Academy of Marketing Science* 48 (0): 137–141. <https://doi.org/10.1007/s11747-019-00710-5>.
- Sanchez-Rola, I., M. Dell’Amico, P. Kotzias, D. Balzarotti, L. Bilge, P.A. Vervier, and I. Santos. 2019. Can I opt out yet? GDPR and the global illusion of cookie control. In *Proceedings of the 2019 ACM Asia conference on computer and communications security*, 340–351. <https://doi.org/10.1145/3321705.3329806>.
- Scott, J. 2004. Ethics, governance, trust, transparency and customer relations. *The Geneva Papers on Risk and Insurance: Issues and Practice* 29 (1): 45–51.
- Selbst, A.D., and J. Powles. 2017. Meaningful information and the right to explanation. *International Data Privacy Law* 7 (4): 233–242.
- Sørum, H., and W. Presthus. 2020. Dude, where’s my data? The GDPR in practice, from a consumer’s point of view. *Information Technology and People* 34 (3): 912–929. <https://doi.org/10.1108/ITP-08-2019-0433>.
- Syrmoudis, E., S. Mager, S. Kuebler-Wachendorff, P. Pizzinini, J. Grossklags, and J. Kranz. 2021. Data portability between online services: An empirical analysis on the effectiveness of GDPR Article 20. *Proceedings on Privacy Enhancing Technologies* 3: 351–372.
- Temme, M. 2017. Algorithms and transparency in view of the new general data protection regulation. *European Data Protection Law Review* 3: 473.
- van Berkel, N., J. Goncalves, D. Russo, S. Hosio, and M.B. Skov. 2021. Effect of information presentation on fairness perceptions of machine learning predictors. In *Proceedings of the 2021 CHI conference on human factors in computing systems, CHI '21*. New York: Association for Computing Machinery. <https://doi.org/10.1145/3411764.3445365>.
- Wachter, S., B. Mittelstadt, and L. Floridi. 2017. Why a right to explanation of automated decision-making does not exist in the general data protection regulation. *International Data Privacy Law* 7 (2): 76–99.
- Wilson, R.A., and F. Keil. 1998. The shadows and shallows of explanation. *Minds and Machines* 8 (1): 137–159. <https://doi.org/10.1023/A:1008259020140>.

Publisher’s Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



About the authors

Jacob Dexe is a Researcher at Research Institutes of Sweden (RISE) and a PhD Student at the KTH School of Electrical Engineering and Computer Science. His research is in transparency in algorithmic decision-making.

Ulrik Franke PhD, is a Senior Researcher at RISE and an Affiliated Faculty Member at the KTH School of Electrical Engineering and Computer Science. His research interests include transparent algorithms, cyber situational awareness, cyber insurance, IT service availability and enterprise architecture.

Kasia Söderlund LLM, PhD Student in Technology and Society at LTH, Lund University, and a Member of the Wallenberg AI, Autonomous Systems and Software Program-Humanities and Society (WASP-HS). Her research focus is on the issues of transparency, explainability, fairness and trust in AI and autonomous systems.

Niels van Berkel PhD, is an Associate Professor at Aalborg University. His research areas include human-computer interaction, social computing, and ubiquitous computing.

Rikke Hagensby Jensen PhD, is an Assistant Professor at Aalborg University. Her research is centered around human-computer interaction, interaction design and how design of smart and ubiquitous systems may influence practice.

Nea Lepinkäinen LLM, PhD Candidate in the Faculty of Law, University of Turku. Her research falls in the field of law and society, and currently she studies social harm and crime following from the increasing use of artificial intelligence and automated decision-making systems.

Juho Vaiste is a Project Researcher and PhD Candidate at University of Turku. His research focuses on ethics of artificial intelligence and management of ethical and responsible technology development.

Authors and Affiliations

Jacob Dexe^{1,2}  · **Ulrik Franke**^{1,2} · **Kasia Söderlund**³ · **Niels van Berkel**⁴ · **Rikke Hagensby Jensen**⁴ · **Nea Lepinkäinen**⁵ · **Juho Vaiste**⁵

Ulrik Franke
ulrik.franke@ri.se

Kasia Söderlund
katarzyna.soderlund@lth.lu.se

Niels van Berkel
nielsvanberkel@cs.aau.dk

Rikke Hagensby Jensen
rjens@cs.aau.dk

Nea Lepinkäinen
nea.lepinkainen@utu.fi

Juho Vaiste
juho.vaiste@utu.fi

¹ RISE Research Institutes of Sweden, P.O. Box 1263, 164 29 Kista, Sweden

² KTH Royal Institute of Technology, 100 44 Stockholm, Sweden



-
- ³ Lund University, PO Box 117, 221 00 Lund, Sweden
 - ⁴ Aalborg University, Selma Lagerlöfs Vej 300, Aalborg East, 9220 Ålborg, Denmark
 - ⁵ University of Turku, Turun yliopisto, 20014 Turku, Finland

