Grant Agreement Number: 825225

Safe-DEED

www.safe-deed.eu

D2.4 User experiment report

Deliverable number	D2.4
Dissemination level	Public
Delivery date	29 November 2019
Status	Final
Author(s)	Mark de Reuver, Tobias Fiebig, Wirawan Agahari, Vidyottama Faujdar

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 825225.

Abstract

Safe-DEED develops technologies that enable safe and trusted data sharing. This deliverable presents the first evaluation of Safe-DEED technologies from a business user perspective. Specifically, multi-party computation is evaluated within the context of the use case in WP7. Three mock-ups are developed that visualize and communicate the workings of the multi-party computation algorithms. A small-scale evaluation is conducted through an experiment in an artificial setting. It is found that explanatory text describing the workings of the algorithms enhance the perceived security of users. Visualizations of input data and encryption enhance the trust in the system and the provider. A combination of the approaches may elicit the highest levels of trust and perceived security. The findings are an important first step to understand the economic implications of Safe-DEED technologies, and more specifically, the impact of the technologies on trust and willingness to use.

Changes Summary

Date	Author	Summary	Version
16.08.2019	Vidyottama Faujdar	First draft	0.1
03.10.2019	Mark de Reuver	Second draft	0.2
08.10.2019	Tobias Fiebig	Revisions Added related work on security visualization (Ch2)	0.3
09.10.2019	Wirawan Agahari	Revisions Added section about MPC (Ch2)	0.4
30.10.2019	Mark de Reuver	Version for internal review	0.5
18.11.2019	Mark de Reuver	Version after internal review	1.0

Table of Contents

Executive summary

The vision of Safe-DEED is to develop novel technologies that enable safe and trusted data sharing. The goal of task T2.3 is to evaluate whether Safe-DEED technologies indeed increase the trust of users in data sharing, which is crucial in building confidence in a data economy.

One of the prominent technologies that Safe-DEED develops is multi-party computation (MPC). MPC algorithms compute answers to questions without having to disclose the raw data. The main value proposition of MPC is that it hereby preserves confidentiality while exposing sensitive data. In this deliverable, we evaluate this value proposition, by examining how MPC affects the trust of users in exposing sensitive data.

We develop and evaluate a mock-up that visualizes MPC. A mock-up is a visual representation of a system, which is not yet a working prototype. The mock-up is being developed within the context of WP7 (i.e. revenue management in supply chains). Creating a mock-up for visualizing MPC is challenging. There are currently hardly any working implementations of MPC, so users are largely unaware. Further, the workings of the algorithm are not directly visible to users, which makes it difficult to visualize and communicate the technology in a mock-up.

The mockups are evaluated through three experiments, with nine participants in total. Feedback from users is collected through qualitative interviews after each experiment. Results indicate that users value trust, security and the extent of control that MPC offers them. Especially describing the functionality of MPC in a disclaimer promotes perceptions of security. Visual indicators of how the process works positively affect trust. Giving information about what other users are doing with the mock-up improves the perception of being in control of the process.

Our findings are tentative since they are based on small-scale experiments with a limited number of participants. A limitation is that participants were students who were asked to play the role of buyers in the WP7 use case, since exposing MPC to actual customers was considered premature.

With these limitations in mind, our findings pose important implications for the follow-up work on MPC in Safe-DEED and beyond. Our recommendation for development of MPC prototypes (as for instance takes place in WP5) is to include descriptions as well as visualizations of how the algorithms work. An implication for follow-up work in WP2 is to consider trust, security and perceived control since all three aspects are important in the context of MPC. Quantitative research can be undertaken to statistically test the extent that trust, security and perceived control influence customer acceptance.

1 Introduction

Safe-DEED technologies are designed to resolve the hurdles that stand in the way of a data-driven economy. The main goal of WP2 is to evaluate the economic and business value of Safe-DEED technologies. In T2.1, we evaluate whether and how Safe-DEED technologies affect security threats in data sharing and data marketplaces. In T2.2, we develop data-driven business models that become possible, once the Safe-DEED technologies are available. In T2.3, we quantify the economic impact by evaluating how the Safe-DEED technologies affect the willingness and trust of business users to share or expose data.

This deliverable is part of T2.3. Within the scope of this deliverable, we focus on one of the three main Safe-DEED technologies: multi-party computation (MPC). MPC is a cryptographic technique that allows distributed parties to compute a function without disclosing their private inputs jointly. MPC, as a tool for computing with confidential data, has an advantage in overcoming security and privacy concerns in situations where owners of different types of data would like to compute cooperatively. This allows the integration of resources while preserving confidentiality, and resultantly obtaining more valuable information(Zhao et al., 2019).

To the best of our knowledge, there is no scientific study done on the impact of MPC on trust and willingness to share or expose data by business users. Addressing this gap in research is challenging. MPC algorithms run in the back-end of information systems and are not directly visible to business users. How to visualize and communicate MPC to business users is thus a core issue that has to be resolved, before being able to study the impact on trust and willingness to share data. A complicating factor is that there are hardly any implementations of MPC in the industry today, which implies that there are no reference systems to build upon in research. Observing these gaps and challenges, the main research objective of this deliverable is to evaluate the impact of MPC on trust and willingness of business users to share data, by designing and evaluating alternative ways of visualizing and communicating MPC algorithms. While MPC could principally be used by any actor, we focus on business users here rather than consumers or government organizations.

In line with the Safe-DEED proposal, this first version of D2.4 evaluates only a mock-up within the context of a specific use case in Safe-DEED. We select the use case of WP7: data sharing for revenue management within a supply chain. Revenue management enables dynamic pricing by relocating capacity based on willingness to pay and has been widely applied in consumer service industries such as airline tickets and hotels. However, in business-to-business settings, revenue management has hardly been applied, due to concerns over confidentiality of willingness-to-pay information. This use case thus represents (1) a setting where business users are unwilling to share data; (2) a setting where MPC can overcome confidentiality concerns by allowing computations without having to share raw willingness-to-pay data.

The deliverable contributes to understanding the economic impact of Safe-DEED technologies, specifically MPC, within a specific use case setting. Through the developed mockups and their evaluation, we inform the further development of Safe-DEED prototypes in the other work-packages.

Given the nature of the research objective, we follow a design science research approach (Gregor and Hevner, 2013). First, we develop justificatory knowledge in Section 2, through a background of MPC and revenue management. Next, Section 3 presents the artefact, based on requirements and specifications. Section 4 discusses the method for evaluation, using an experimental design. Section 5 provides the results from a preliminary qualitative study with nine participants, followed by discussion, conclusions and next steps in Section 6.

2 Background

The background information serves as justificatory knowledge to inform the development of the artefact (Kuechler and Vaishnavi, 2012). Section 2.1 provides a brief theoretical basis on trust and security, which are the main variables of interest for our evaluation. Next, Section 2.2 provides a

background on MPC, which is the technology being studied. Finally, Section 2.3 provides an overview of revenue management, which is the context in which the mock-ups are being created.

2.1 Trust and security

Trust is often considered a multidimensional concept. Trust in this deliverable refers to the inter-organizational rather than interpersonal level, and includes both reliability and benevolence related aspects (Zaheer, McEvily and Perrone, 1998). Specifically, we define trust as "the willingness of a party (trustor) to be vulnerable to the actions of another party (trustee) based on the expectation that the other (trustee) will perform a particular action important to the trustor, irrespective of the ability to monitor or control that other party (trustee)" (Mayer, Davis and Schoorman, 1995).

Security control plays a role in building trust as it lowers the customers' risk in transacting (Gefen, Karahanna, & Straub, 2003). Especially in the context of the world wide web (WWW) and encryption, the visualization of security functionality has been instrumental. According to prior studies (Akhawe and AP Felt, 2013; Felt et al., 2014), the design and visualization of security features significantly influence users ability to make the right choices. Similarly, the design of security mechanisms, a permission system for installed applications in this case, may influence users' behaviour in adopting secure choices (Felt et al., 2012). The usability and trustability of interfaces is especially crucial in the context of encryption tools, as for example in the case of GNU/PGP (Pretty Good Privacy), which has been consistently found to be lacking (Whitten and Tygar, 1999; Ruoti, Andersen, Zappala and Seamons, 2015). As history, and more recent work on secure messaging clients shows, a usable and trust provoking interface are a necessity for the wide-spread adoption of security mechanisms (Unger et al., 2015).

For MPC specifically, Lapets et al. (2016) explain that, for a secure multi-party computation protocol to be trusted, it should be (1) comprehensible (i.e. simple enough to understand); (2) auditable (i.e. transparent with open-source code); (3) accessible (i.e. deployable without requiring effort such as setup or specialized software / hardware); (4) simple (i.e. requiring no technical know-how from users); (5) idempotent (i.e. allowing resubmitting the data); and (6) providing feedback (i.e. proactively warn users about spurious data). Privacy and confidentiality are crucial as no party should be able to get information from others (Bogetoft et al., 2006; Zhao et al., 2019). Anonymity is vital as nobody should be able to identify a bidder (Omote and Miyaji, 2001).

In the context of visualizing Secure MPC, the problem is that we have to find a visualization that is comprehensible, simple, and provides feedback, while ideally also improving the auditability (Lapets et al. 2016). Given the underlying technical concepts, auditability is difficult to combine with simplicity. However, given the issues with earlier work, e.g., PGP (Whitten and Tygar, 1999; Ruoti et al., 2015), and findings from Felt et al. (2012), comprehension and simplicity outweigh auditability: A complex, but auditable, system is no use if it is not used at all.

2.2 Secure MPC

Businesses are typically reluctant to share sensitive data with competitors, suppliers or clients, because of concerns over confidentiality, privacy and harming the competitive position. Trusted third parties or data platforms are one solution to overcome these issues but comes with challenges of trust and data governance (Lapets et al., 2016). A specific concern is that trusted third parties may resell or otherwise reuse the data in undesirable ways.

One way to tackle this problem is by implementing MPC.[1] It is a cryptographic technique where two or more parties perform a joint computation which results in a meaningful output, without disclosing the input provided by either party (Roman and Vu, 2019; Zhao et al., 2019). On completion of the protocol via a joint function *f* that obtains its inputs from the secret data of each party, users only find out the final output of the function *f* and their own input, without any information about the other

[1] See for details on MPC: Safe-DEED Deliverable 5.1 `Requirements for Secure Computations on Large Datasets with Multiple Onwers' (M6)

parties' inputs. To ascertain whether a protocol is secure, researchers have established various security requirements, as can be seen in Table 1.

Table 1 Security requirements for MPC protocol (Zhao et al., 2019)

Security requirements	**Description**
Privacy	no party should be able to get information apart from their own
Correctness	the output obtained by each party must be correct
Independence of input	the input of a corrupted party must be independent of the inputs of the honest parties
Guarantee of output	Corrupted parties should not be able to stop honest parties from receiving their own outputs
Fairness	Corrupted parties should obtain their outputs if and only if the honest parties obtain their own outputs

One illustration of MPC is a secure comparison function to determine which one of two millionaires is richest, without revealing the net worth to each other. This illustration often referred to as the millionaire's problem (Yao, 1982). An extension of this example includes a situation where patients want to access their clinical records. They are able to make a query to the medical database of DNA related diseases by using their own private DNA code, while not wanting the hospital and others to know about their DNA and their potential disease. Simultaneously, the hospital would not want to reveal its database to the patients. Other deployments of MPC in a real-life setting can be found in Table 2.

Table 2 Real-world deployment of MPC, summarized by Choi & Butler (2019)

Real-world deployment	**Description**	**References**
Auction-based pricing	A nationwide auction to determine a market-clearing price for sugar-beet contracts in Denmark	Bogetoft et al. (2009)
Tax fraud detection	Secret-sharing-based system to detect VAT fraud in Estonia	Bogdanov et al. (2015)
Satellite collision prevention	Computing collision probability without sharing private orbital information	Hemenway et al. (2016)

MPC offers an important value proposition in allowing safe and secure data sharing that complies with data protection regulations (Archer et al., 2018). However, despite its potential, there are several barriers to implementing MPC, as identified by Choi & Butler (2019). For instance, current data protection regulations discourage data sharing and collaboration among private actors, which creates little incentives to implement MPC. Also, the complexity of MPC makes it difficult for non-experts to understand how MPC works. Moreover, there is a concern over the correctness of the computation results because it is not possible to verify the input data. Finally, there are concerns over performance limitations and scalability.

2.3 Revenue management

Through revenue management, businesses apply dynamic pricing to optimize supply and demand, without having to enlarge their capacity (Wang and Bowie, 2009). A well-known example is how airline companies change the pricing of tickets based on their capacity and observed demand from customers. Revenue management is especially suitable in settings where customers have heterogeneous and time-variable demands (e.g. airline tickets have different value at different times of a day), while suppliers have a limited capacity to deliver (e.g. limited seating capacity in an airline) which is perishable (e.g. airline seats cannot be stocked). These characteristics imply that revenue management is especially suitable in service industries.

In manufacturing, revenue management is hardly applied yet, as the main difference is that production and consumption can be separated in time (Zatta, 2016). For this reason, orders can be sequenced, and revenue management should take into account the `lead time' (i.e. duration between order and delivery of goods). Specifically, customers may be willing to pay surplus prices for earlier deliveries.

Implementing revenue management in manufacturing is challenging. Manufacturers may want to keep their unallocated capacity and ability to deliver early confidential, since this is sensitive information. For similar reasons, customers may not want to disclose their willingness to pay for earlier deliveries. MPC may resolve these concerns by keeping such input data confidential, while only sharing the computed optimal delivery date and price. Yet, applications of MPC in revenue management are lacking in practice and in literature, as MPC has hardly been applied or visualized in any real-life context at all (Sousa, Antunes and Martins, 2018; Roman and Vu, 2019; Zhao et al., 2019).

3 Mockup development

3.1 Context of WP7 use case

The use case in WP7 focuses on revenue management in the semiconductor industry. This industry is suitable for exploring revenue management, as it is characterized by volatile market demand, long production times and high capital investment. For details on the use case, we refer to deliverables in WP7.

The use case assumes that there is one manufacturer and multiple buyers. Currently, volume-based pricing is used, which means that customers receive a quote based on the desired volume. The alternative pricing model, which we explore, is one in which the price depends on the lead time between placing the order and delivery.

While most buyers currently use Electronic Data Interchange (EDI) to place orders, it is anticipated that, in the future, buyers will use a common web-portal. On the portal, buyers can search for products while specifying details. When selecting a product, buyers see the price per unit, shipping charge and available stock. Buyers can adjust the volume of units, and immediately see the associated unit price, whereas higher volumes typically imply lower unit price. When placed, every quote is checked by the supplier internally.

3.2 Requirements

Revenue management requires input from buyers and suppliers. From the supplier, information is required on production cycle dates, capacity planning and end-user purchase patterns. From the buyers, the maximum price they are willing to pay for an earlier delivery date is needed. It is likely that neither party is willing to disclose these business-sensitive information in public.

The MPC protocol should compute the earlier due date and updated price based on unallocated production capacity and willingness to pay from buyers. Input for the computation are initial lead time price and volume matrix from manufacturers, as well as desired lead time, volume and price demands from buyers. After the computation, parties should only learn the new date and updated price, and all input data remains confidential. For the supplier, this implies that unallocated production capacity can

be monetized, without disclosing this information. For the buyers, the solution enables adapting volumes and delivery dates when their business conditions require so.

In this deliverable, we assume that MPC algorithms can compute securely, accurately and with reasonable latency, since these are the goals of WP5. Probably, the required computations can be done through algorithms such as homomorphic encryption, private set interaction or the general MPC protocols developed in the Safe-DEED FRESCO framework.

For our mockups, we build on ideas in auction design (Burmeister et al., 2002; Pla, López, Murillo and Maudet, 2014). We select a single-sided auction as there is only one supplier and multiple buyers and distributors. The auction is multi-attribute, because buyers bid through price, quantity as well as delivery dates. As the manufacturer would sell to several potential buyers who compete for the products, this would be a forward auction. Lastly, it is sealed-bid due to the desired confidentiality of buyers.

Requirements for the mockup are summarized in Table 1, listing functional requirements (FR) and non-functional requirements (NF).

Item	The mockup should …	Source
FR1	… allow buyers to insert desired lead time, price and volume matrix	This deliverable, Section 4
FR2	… simulate a single-sided, multi-attribute, forward, sealed-bid auction	Burmerister et al., (2002); Pla et al. (2014); Torrent-Fontbona et al. (2015)
FR3	… display outputs of the simulated auction to buyers	Bogetoft et al. (2006)
NF1	… not disclose information from other parties	Zhao et al. (2019); Bogetoft et al. (2006)
NF2	… not allow identifying bidders	Omote & Miyaji (2001)
NF3	… explain MPC in a way that is comprehensible for non-expert users	Lapets et al. (2016)
NF4	… be usable within a relatively narrow time frame by non-expert users whose technical expertise comprises basic web browsing clients	Lapets et al. (2016)

Note: FR = functional requirement; NF = non-functional requirement

Table 3: Mockup requirements

3.3 Specifications of three mock-ups

Mock-ups are designed such that they resemble the current web-portal in the WP7 use case. Firstly, buyers are prompted to input their multi-attribute bids. Once filled out and confirmed, buyers see an intermediate screen (visible for five seconds) indicating that the computation is in process. The next screen provides the outcome of the computation and asks buyers to accept or reject the new offer.

We design three mock-ups that visualize and communicate the use of MPC in revenue management. In each of the mock-ups, buyers can input their bids on the left-hand side of the screen. The mock-ups differ on how they visualize and communicate MPC to buyers, which is displayed on the right-hand side of the screen. As noted in Section 1, MPC is challenging to visualize and communicate, since the algorithms run in the back-end of the system. We draw upon literature on visualization of computer security to inspire the mock-up variations. All mock-ups are developed using the programming

language C#, on Visual Studio, in order to have the impression of realistically running websites to participants.

The first mock-up only provides text that explains the MPC solution, see Figure 1. The disclaimer text explains that the algorithm is secure, and that inputs will be kept strictly confidential. A padlock symbol is placed below the text, as such symbolic representations have been suggested to improve security perceptions of users (Whitten and Tygar, 1999; Kainda, Fléchais and Roscoe, 2010)(Kainda et al., 2010; Whitten & Tygar, 1999).

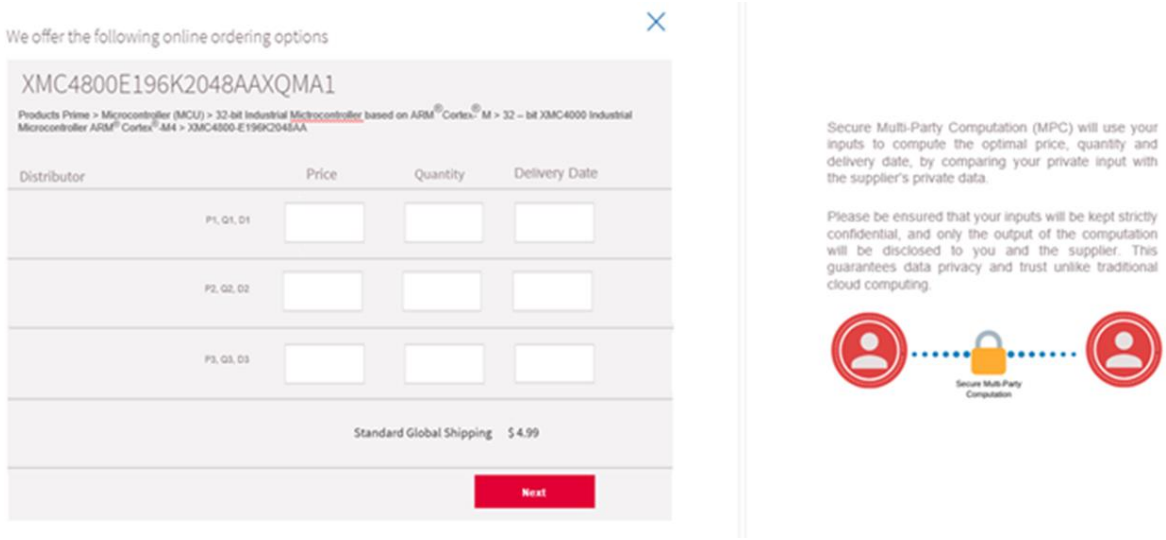

Figure 1: Mock-up with disclaimer text

The second mock-up visualizes the encryption of their input data, see Figure 2. We build upon suggestions in security literature that users are more confident in systems when encryption approaches are made transparent (Whitten & Tygar, 1999). After users provide input data on the left-hand side of the screen, a blurred matrix appears on the right-hand side, which visualizes that input data is being encrypted and kept secure. The right-hand side also shows that competing buyers are interested in the product, but that their input data is similarly kept confidential.

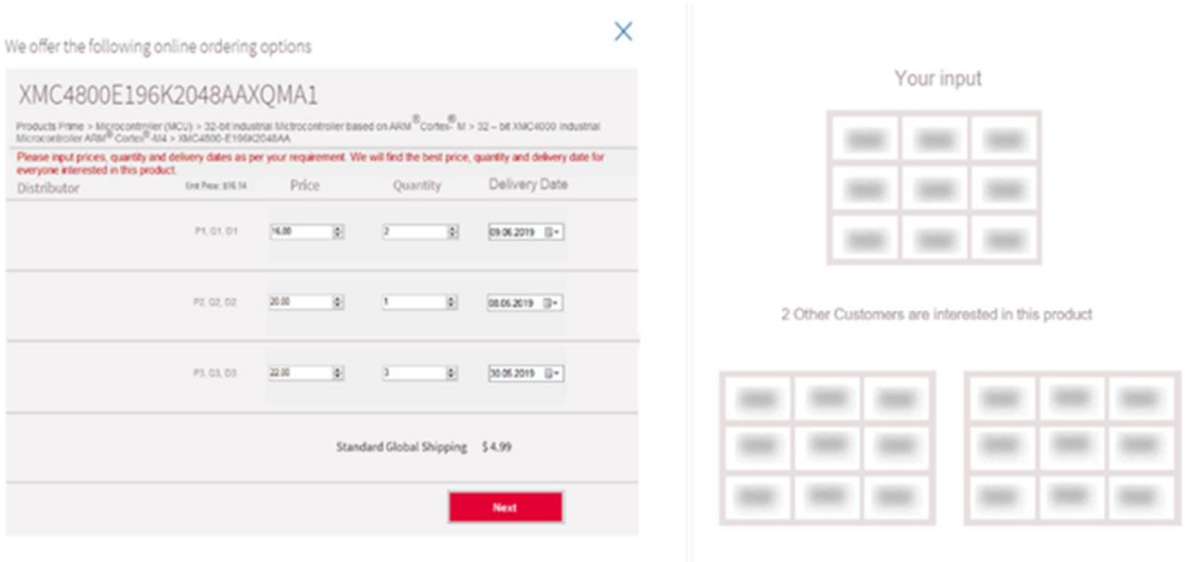

Figure 2: Mock-up with blurry input data

The third mock-up similarly visualizes that data from competing buyers is kept confidential, using a blurred matrix. However, it differs from the second mock-up by keeping the input data from the buyer transparent, see Figure 3. This suggests to buyers that their data will be encrypted, but not manipulated by the algorithm.

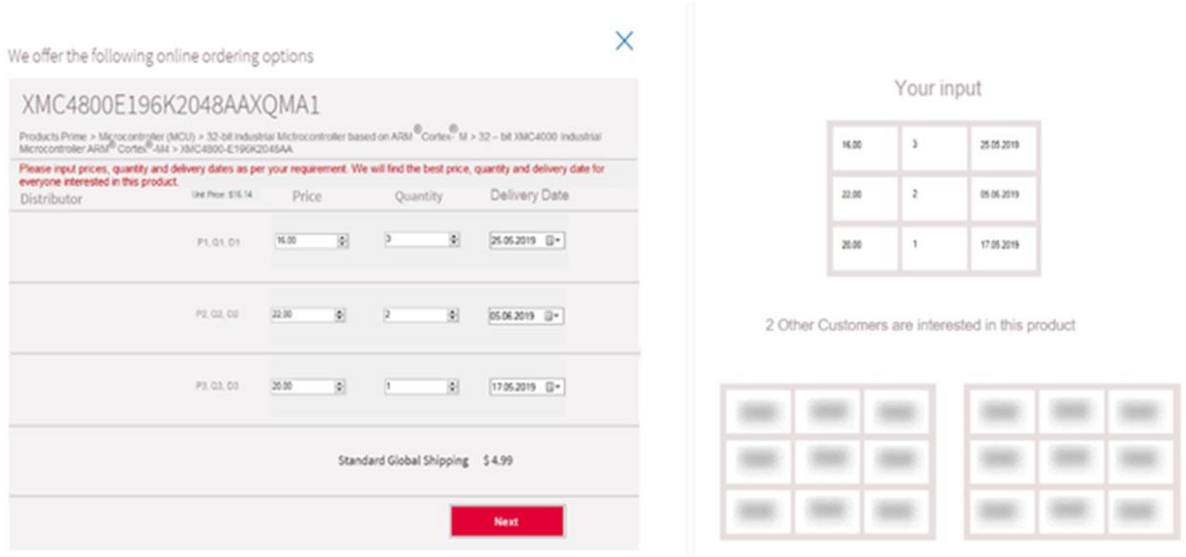

Figure 3: Mock-up with clear input data, blurred competitor data

4 Evaluation method

To evaluate the mock-ups, we conduct an experiments. Experiments allow systematically comparing different treatments (i.e. mock-ups) on variables of interest (i.e. trust and willingness-to-pay), while controlling for disturbing factors. In the experiment, participants use the mock-ups as if they were purchasing goods from the manufacturer.

4.1 Experimental design

The experiment assumes that participants are distributors of the goods produced by the manufacturer. Distributors are generally wholesalers who purchase and re-sell products at a mark-up price. For distributors, the confidentiality of purchase price is of high importance, since their customers should not learn about the level of their mark-ups. Hence, the need to keep input data confidential is high. Further, revenue management poses significant benefits to distributors, as they serve a large base of customers who have different and time-varying demands on delivery dates and order volumes. Distributors are also most likely to benefit from reduced bullwhip effects, which revenue management affords. In sum, distributors are considered to have significant benefits from revenue management, and significant concerns over confidentiality that MPC may resolve.

During the experiment, the concept of revenue management was first explained to participants, which took fifteen minutes on average. Next, participants received a demo of a mock-up that resembles the current buying web-portal, in order to explain the current situation. In the main part of the experiment, participants were asked to use the three mock-ups discussed in Section 3, which took about 30 minutes in total. They were asked to give realistic input regarding the desired price and delivery dates, resembling the order of magnitude in the demonstration. At the end of the experiment, participants were provided a template to provide feedback and rank the mock-ups. After that, semi-structured group interviews were held.

Experiments were done in groups of three participants, which resembles the actual situation of multiple buyers bidding for goods. While one of the three groups was less talkative than the others, each of the interview sessions produced sufficiently rich information. Semi-structured interviews were also done in groups, as we intended to create deeper and novel insights by having participants respond to each other. Typical downsides of group interviews, such as loss of privacy, were not considered relevant in this particular context. Participants were reminded to speak their thoughts freely to limit the risk of groupthink.

The semi-structured interviews started by asking for which of the mock-ups would generally be preferred. Next, the participants were asked to state their preference regarding which mock-up would keep their interest best, which was the easiest to use, and which was felt to be the most secure. Finally, questions were asked about their general perception on revenue management and what other comments they had. Interviews were recorded and transcribed, and notes were taken during and soon after each session.

4.2 Participants

Participants in the experiment are internship students working at the WP7 company, who are asked to play the role of buyers, specifically distributors. While students are generally less knowledgeable than actual buyers, the advantage is that they pose a more homogeneous population. Using actual buyers would introduce disturbing effects as they have heterogeneous experience levels and attitudes towards revenue management. Further, since most of the participants work in the revenue management department of the company, we consider them sufficiently knowledgeable of the setting.

Participants are recruited through emails and personal invitations. In total, nine students participate, which we consider suitable for an exploratory study. Prior to participating, students read and signed informed consent forms. As participants are internship students, we took effort to underline that participation was voluntary, and that they could withdraw at any point in time.

4.3 Analysis approach

Analysis of the qualitative interview transcripts and notes was done through a coding approach. In the first round of open coding, codes were assigned to relevant text fragments, with labels that closely resemble the wording of participants. The second round of coding was done to refine the coding and eliminate codes not related to the research objective. Next, the transcripts and codes were inserted into Atlas.TI 8.0, a software system for qualitative data analysis. In a round of axial coding, codes were grouped into higher-order categories. The categories were then combined into overarching themes. In addition to five pre-determined themes (usefulness, ease-of-use, trust, security and willingness to use), a sixth theme was identified relating to `perceived control', which refers to the feeling that participants feel in control of the situation while using the mock-ups.

5 Evaluation results

Of the six themes, especially trust, security and perceived control affect the preference for one of the mock-ups. The participants did not express a clear preference between the mock-ups when asked about usefulness, ease of use and willingness to use. Instead, these themes entailed only a general discussion about revenue management in terms of usefulness and barriers. This is partly to be expected since the mock-ups do not differ regarding functionality or affordances, but rather regarding the communicability and visibility of security and trustworthiness. Hence, the remainder of the analysis focuses on the themes of trust, security and perceived control only.

Table 2 provides an overview of the relevant codes within these themes. In the table, the Code frequency shows how often codes were mentioned in the sessions, whereas Preferred mock-up shows how often the codes were positively or negatively related to the mock-ups. how often codes within these themes were mentioned (middle set of columns) and how often these were positively or negatively related to the mock-ups.

Theme	Category	Code	Code frequency				Participants mentioning	Preferred mock-up		
			Session 1	Session 2	Session 3	Sum		Disclaimer	Blurry	Clear
Trust	Malicious intent of supplier	Unfair treatment	1	0	4	5	4	-1	-2	-2
		Dubious / malicious intent	0	1	6	7	4	-4	-3	-1
		Lack of trust in disclaimer supplier	0	1	4	5	3	-5	0	0
		Suspicious	1	0	4	5	3	-3	-2	0
	Transparency	Visuals induced confidence in encryption	3	4	3	10	6	0	9	1
		Transparency	0	4	5	9	5	-1	0	8
		Honest	0	1	2	3	3	0	0	3
		Assurance of inputs	0	1	6	7	4	0	-1	3
Security	Use of explicit textual information	Explicit information about security and data protection	2	5	6	13	7	12	-1	-1
		MPC information made it easy to follow	0	1	0	1	1	1	0	0
		MPC information felt safe	0	1	0	1	1	1	0	0

		Explanation of process	1	2	6	9	4	8	-1	-1
		Ability to research by yourself	2	0	0	2	1	2	0	0
	Insecurity in repetition	Unnecessary repetition reducing security	4	1	0	5	4	0	-1	-5
		Increased hacking possibility due to multiple data paths	1	0	0	1	1	0	-1	-1
		Disturbing	1	0	0	1	1	0	0	-1
		Tricky visuals without information	0	1	2	3	3	0	-3	-2
		Industrial espionage	1	0	0	1	1	0	0	-1
Perceived control	Competition induced stress	Competitors push bids	2	1	0	3	2	0	-3	-3
	Feeling of being in control	Feeling of being in control	2	1	0	3	3	0	2	2
		Competitor advancement	2	4	1	7	5	-1	4	4
		Assurance in market demand	0	0	1	1	1	0	1	1

Table 2: Frequency table for codes relating to trust, security and perceived control

Table 2 shows that, regarding the theme of trust, the Clear mock-up was clearly preferred, while the Disclaimer mock-up was not. In terms of security, the Disclaimer mock-up was favoured over both other interfaces. Regarding perceived control, both the Blurry and Clear mock-ups were slightly favoured. The remainder of this section discusses these findings in more detail.

5.1 Trust

Trust was generally the most frequently discussed of all six themes, especially in the third session. Participants discuss both aspects of trust that relate to the system and the supplier. The more transparent the mock-up, the less participants felt that the supplier has a malicious intent.

The Disclaimer mock-up generated distrust in the supplier, as participants felt they might be unfairly treated against competitors, that the supplier might have ulterior motives to use the data, lack of trust in the text, and overall feelings of suspicion. Specific concerns regarding unfair treatment were that data might still be shared with other distributors or that the algorithm may handle the data unfairly. Another concern was that, although the disclaimer signals that the supplier `tries to hard' to believe that no worry is needed, which works exactly the opposite way. Participants also questioned whether the text was accurate and credible. They found that the claim that data was handled securely was not supported by any visualization, as illustrated by the following quote:

S3-P3: "I got the feeling that I might be cheated now, you never know what's actually going on. … I personally wouldn't trust the text because anything can be done by a company even though it says something. But you don't actually know."

S3-P3: "We don't know actually if my input is going to be taken seriously or not"

The Blurry mock-up scored highly regarding trust in the system using confidential data. Participants found it assuring to see all inputs blurred, as they inferred that competitors would not see their inputs either, as illustrated by the following quote:

S1-P3: "Showing that the others are greyed out, gives me confidence that the others also see me as only greyed out"

On the other hand, the Blurry mock-up showed the input data on the right-hand side of the screen which prompted concerns over that raw data might still be disclosed or somehow used. The Blurry mock-up further led to suspicion that perhaps the correct input data were not being used, as users cannot verify these. The following conversation illustrates this:

S3-P1: "My data might be used under the hood somehow"

The Clear mock-up performed well regarding transparency as they could verify their bids after submitting, which makes participants trust the honesty of the system and the MPC algorithm. The lack of negative comments on the Clear mock-up suggests it is preferred the most regarding trust.

5.2 Security

Participants found the Disclaimer mock-up to be the most secure one as this mock-up give explicit information about security and data protection, which was mentioned by seven out of nine participants. Participants liked that they could read about the security of the system and confidentiality of their input. The explicit description allows them to understand the MPC algorithm as well as how outcomes are being calculated. Participants also liked that the information allowed them to research and verify the security claims being made.

The Blurry and Clear mock-ups made participants feel insecure as their input data was still being displayed. They found the display of input data unnecessary, expected increased chance of being hacked and generally felt disturbed due to the duplication of information. They were also unable to understand the visuals without supporting text.

5.3 Perceived control

Under this theme, discussions are grouped about how much participants feel in control of the situation while using the mock-ups.

The Blurry and Clear mock-ups made participants feel like they could monitor their competitors' results and could learn about the market demand for the products. This outcome elicited both positive and negative feelings. On the one hand, some participants felt less in control of the situation, since they experienced stress from competing bidders and felt that they had to increase their bids. On the other hand, other participants felt more in control as they could visualize what competitors were doing, and how the competitors were doing in the auction. They expressed that they missed this in the Disclaimer mock-up.

6 Discussion

We found that transparency, as in the Blurred and Clear mock-ups, elicit positive feelings of trust in the system. This finding is in line with existing literature that states users have more confidence in security when encryption is transparent (Whitten & Tygar, 1999; Turner, 2002). Textual explanations, as in the Disclaimer mock-up, contribute the most to positive feelings of security. This is in line with earlier work stating that information provisioning enables users to assess the security of a system (Yee, 2002) and that users search for indicators signalling security (Kainda et al, 2010). Importantly, however, such textual information reduces the trust in the supplier, since participants suspect the information may be incorrect or misleading.

An unexpected finding is that transparency over security, in this specific setting of an auction, induces stress for participants. This is because the a comprehensible visualization shows that competitors are bidding for products too, which reduces the feeling of being in control of the situation. Hence, in this specific context, transparency over the workings of the algorithm through visualization may lead to negative side-effects.

As part of the experimental design, we isolated the features of text and different forms of visuals, in order to understand their individual effects on trust and willingness to use. In terms of our three mock-ups, the Clear mock-up is generally preferred, building both trust and perceived control. However, in reality, mock-ups that combine transparency through visualizations and explanation through supporting text may perform best in building trust as well as perceived security. Follow-up research should focus on such mock-ups (or prototypes) and explore the interaction effects between the text and visualization features. A mock-up combining the features may look like the sketch in Figure 4.

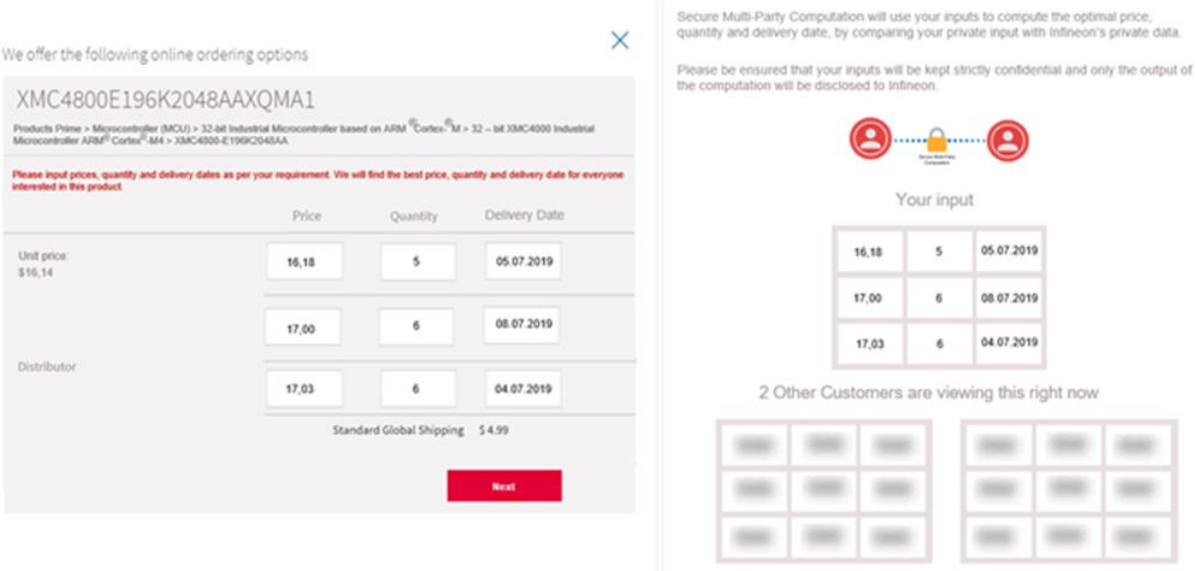

Figure 4: Sketch of mock-up combining text and visualization

A limitation is the artificial setting of the experiment, and the use of students as participants. To counter this, we briefed the participants extensively about the setting, their role and the background of revenue management. Yet, for instance, in reality, when buyers have an established relationship with their supplier, the feelings of distrust and suspicion may be lower than what we found in the experiment. Students from a technical background may also be more critical about security risks and misrepresentation of security procedures than business users. Another limitation is that although the nine participants elicited a wealth of feedback, adding more participants might lead to additional insights.

A second limitation is that the mock-ups are not working prototypes. In essence, we therefore evaluate visualizations and explanations of security, and not the security technologies themselves. Further, exchanging data on bids between the systems of participants might have led to a more realistic experience, and possibly amplified concerns over competition and lack of control of the situation. Exchanging of data might also have increased the perception of security, since participants could then verify that their input data is indeed not being disclosed to others.

7 Implications for Safe-DEED

A main contribution of the study is to set the basis for understanding how to visualize and communicate MPC, within a context of revenue management. We found that, if visualized and explained properly, the use of MPC can contribute to perceptions of trust, security and perceived control. This is important as trust and security in data sharing by business users are key conditions to unlock the data economy.

We infer from the study that Safe-DEED prototypes should likely both visualize and explain in the text the workings of the Safe-DEED technologies, in order to elicit trust and perceptions of security. In this way, this deliverable lays a basis for further development and experiments with MPC prototypes in the context of Safe-DEED.

From an ethical perspective, it is important to remark that the work in this deliverable evaluates the communicability and visualization of security technologies, and not the technologies themselves. This approach implies that results may be used to increase feelings of security for systems that are not necessarily secure. While this is not the case for Safe-DEED technologies, it is important to keep this in mind as a backdrop to the study results.

The work also shows that explaining and visualizing MPC is only a part of the puzzle. Trust in the supplier, competitive pressures and trust in the system as a whole are similarly important. In other words, the implementation of Safe-DEED technologies such as MPC should always be considered within the context of a specific use case application and business domain, which signifies the importance of the work in T2.2.

Next steps within T2.3 will likely include the evaluation of prototypes developed in the other work packages. The mock-ups developed for this deliverable were conducted within WP2 and WP7. Once working prototypes are available, the next step is to replicate and expand this preliminary study. Further, next steps will likely involve a wider range of application settings (i.e. beyond WP7 use case) and evaluations with larger and more diverse samples than students.

8 References

Akhawe, D. and AP Felt. (2013). “Alice in warningland: A large-scale field study of browser security warning effectiveness.” Presented at the 22nd {USENIX} Security Symposium.

Archer, D. W., D. Bogdanov, Y. Lindell, L. Kamm, K. Nielsen, J. I. Pagter, … R. N. Wright. (2018). “From Keys to Databases—Real-World Applications of Secure Multi-Party Computation.” *The Computer Journal*, *61*(12), 1749–1771.

Bogdanov, D., M. Jõemets, S. Siim and M. Vaht. (2015). “How the Estonian Tax and Customs Board Evaluated a Tax Fraud Detection System Based on Secure Multi-party Computation” (pp. 227–234). Presented at the International Conference on Financial Cryptography and Data Security, Springer, Berlin, Heidelberg.

Bogetoft, P., D. L. Christensen, I. Damgård, M. Geisler, T. Jakobsen, M. Krøigaard, … T. Toft. (2009). “Secure Multiparty Computation Goes Live.” In: R. Dingledine & P. Golle (Eds.), *Financial Cryptography and Data Security* (pp. 325–343). Springer Berlin Heidelberg.

Bogetoft, P., I. Damgård, T. Jakobsen, K. Nielsen, J. Pagter and T. Toft. (2006). “A Practical Implementation of Secure Auctions Based on Multiparty Integer Computation” (pp. 142–147). Presented at the International Conference on Financial Cryptography and Data Security, Springer, Berlin, Heidelberg.

Burmeister, B., T. Ihde, T. Kittsteiner, B. Moldovanu and J. Nikutta. (2002). “A practical approach to multi-attribute auctions” (pp. 670–674). Presented at the Proceedings. 13th International Workshop on Database and Expert Systems Applications, IEEE Comput. Soc.

Choi, J. I. and K. R. B. Butler. (2019). “Secure Multiparty Computation and Trusted Hardware: Examining Adoption Challenges and Opportunities.” *Security and Communication Networks*, *2019*, 1–28.

Felt, A. P., E. Ha, S. Egelman, A. Haney, E. Chin and D. Wagner. (2012). “Android permissions” (p. 1). Presented at the Proceedings of the Eighth Symposium on Usable Privacy and Security - SOUPS ’12, ACM Press.

Felt, A. P., R. W. Reeder, H. Almuhimedi, S. Consolvo, A. P. Felt, R. W. Reeder, … S. Consolvo. (2014). “Experimenting at scale with google chrome’s SSL warning” (pp. 2667–2670). Presented at the Proceedings of the 32nd annual ACM conference on Human factors in computing systems - CHI ’14, ACM Press.

Gregor, S. and A. Hevner. (2013). “Positioning and Presenting Design Science Research for Maximum Impact on JSTOR.” *MIS Quarterly*, *37*(2), 337–355.

Hemenway, B., S. Lu, R. Ostrovsky and W. Welser IV. (2016). “High-Precision Secure Computation of Satellite Collision Probabilities” (pp. 169–187). Presented at the International Conference on Security and Cryptography for Networks, Springer, Cham.

Kainda, R., I. Fléchais and A. W. Roscoe. (2010). “Security and Usability: Analysis and Evaluation” (pp. 275–282). Presented at the 2010 International Conference on Availability, Reliability and Security, IEEE.

Kuechler, W. and V. Vaishnavi. (2012). “A Framework for Theory Development in Design Science Research: Multiple Perspectives.” *Journal of the Association for Information Systems*, *16*(3), 395–423.

Lapets, A., N. Volgushev, A. Bestavros, F. Jansen and M. Varia. (2016). *Secure multi-party computation for analytics deployed as a lightweight web application*. Computer Science Department, Boston University. Retrieved from https://open.bu.edu/handle/2144/21786

Mayer, R. C., J. H. Davis and F. D. Schoorman. (1995). “An Integrative Model Of Organizational Trust.” *Academy of Management Review*, *20*(3), 709–734.

Omote, K. and A. Miyaji. (2001). “A Practical English Auction with One-Time Registration” (pp. 221–234). Presented at the Australasian Conference on Information Security and Privacy, Springer, Berlin, Heidelberg.

Pla, A., B. López, J. Murillo and N. Maudet. (2014). "Multi-attribute auctions with different types of attributes: Enacting properties in multi-attribute auctions." *Expert Systems with Applications*, *41*(10), 4829–4843.

Roman, D. and K. Vu. (2019). "Enabling Data Markets Using Smart Contracts and Multi-party Computation" (pp. 258–263). Presented at the International Conference on Business Information Systems, Springer, Cham.

Ruoti, S., J. Andersen, D. Zappala and K. Seamons. (2015). *Why Johnny Still, Still Can't Encrypt: Evaluating the Usability of a Modern PGP Client*.

Sousa, P. R., L. Antunes and R. Martins. (2018). "The Present and Future of Privacy-Preserving Computation in Fog Computing." In: *Fog Computing in the Internet of Things* (pp. 51–69). Cham: Springer International Publishing.

Torrent-Fontbona, F., A. Pla and B. López. (2015). "A New Perspective of Trust Through Multi-Attribute Auctions." Presented at the Workshops at the Twenty-Ninth AAAI Conference on Artificial Intelligence.

Unger, N., S. Dechand, J. Bonneau, S. Fahl, H. Perl, I. Goldberg and M. Smith. (2015). "SoK: Secure Messaging" (pp. 232–249). Presented at the 2015 IEEE Symposium on Security and Privacy, IEEE.

Wang, X. and D. Bowie. (2009). "Revenue management: the impact on business-to-business relationships." *Journal of Services Marketing*, *23*(1), 31–41.

Whitten, A. and J. D. Tygar. (1999). "Why Johnny can't encrypt: A usability evaluation of PGP 5.0." *Proceedings of the 8th USENIX Security Symposium*.

Yao, A. C. (1982). "Protocols for secure computations" (pp. 160–164). Presented at the 23rd Annual Symposium on Foundations of Computer Science (sfcs 1982), IEEE.

Zatta, D. (2016). *Revenue Management in Manufacturing*. Cham: Springer International Publishing.

Zhao, C., S. Zhao, M. Zhao, Z. Chen, C.-Z. Gao, H. Li and Y. Tan. (2019). "Secure Multi-Party Computation: Theory, practice and applications." *Information Sciences*, *476*, 357–372.