

The Power of Conversation: Learners Become More Cautious Sharing Learning Data after a Group Discussion

Louis Longin^{1,*}, Deisy Montenegro Briceno², and Oleksandra Poquet^{2,3,4,+}

¹Faculty of Philosophy, Philosophy of Science and the Study of Religion, LMU Munich

²LEAPS Lab - LEarning Analytics and Practices in Systems, Department of Educational Sciences, School of Social Sciences and Technology, Technical University of Munich

³Munich Data Science Institute, Technical University of Munich

⁴Centre for Change and Complexity in Learning, University of South Australia

*these authors contributed equally to this work

*corresponding author: louis.longin@lmu.de

ABSTRACT

The ethical integration of the data generated by learners into educational practices is of great importance now that data-rich technologies are prevalent in education. Despite the common agreement that learners should have agency in deciding what to do with their data, existing ethical discussions focus on policies or algorithms, with limited attention to participatory learner practices. Participatory practices, particularly around informed consent, can support ethical and meaningful engagement with data sharing decisions. Using a novel experimental methodology, we explore the effect of group discussion on learner decisions to share their data. We found that learners become more cautious in sharing their data in and after a group discussion. The willingness to share is the lowest when these data are submitted to a government entity and for a collective benefit. Further network analysis of group discussions confirms the observed behavioural effects: participants consistently discussed different aspects of sharing learning data based on the context such as sharing process vs outcome-related learning data. The results suggest that educational data consent is contextual and that mechanisms for consent in educational technology may differ from those in moral judgement. The proposed method of interactive consent therefore not only contributes to theories explaining privacy and effective data collection, but also represents a new way of conceptualising and realising participatory informed consent.

Keywords: participatory informed consent, learning analytics, group discussion, data privacy, interactive groups

Introduction

The use of data collected by educational technologies to provide feedback to learners, known as learning analytics (Siemens, 2013), has been a cornerstone of educational innovation for over a decade. Educational technologies like online learning platforms and recently virtual assistants powered by large language models have not only transformed how students learn, but also how the learning progress can be tracked. By analysing the learning process alongside the learning outcomes, educators can improve not only educational material, and student experience (Brooks et al., 2021; Ferguson et al., 2016; Knoop-van Campen et al., 2023; Sclater, 2017) but also contribute to educational theory (Reimann, 2021). However, extending learning analytics does not come without risks. The collection and processing of large amounts of student data raises concerns about data security and ownership (Wollny et al., 2023; Gašević et al., 2022); whereas large language models can perpetuate and amplify existing biases and unfairness when deployed on learning data (Kasneci et al., 2023).

To address these concerns regulators developed policies while researchers created equitable privacy-preserving algorithms (Joksimovic et al., 2021; Liu and Khalil, 2023). However, less work has been done on developing interventions that involve learners into their data sharing decisions. Nonetheless, there is a growing need to develop participatory interventions that empower students to exercise agency over their data sharing decisions. Existing learning frameworks

are unequivocal about students' involvement in decisions about data (Drachler and Greller, 2016; Slade and Prinsloo, 2013; Pardo and Siemens, 2014). Students themselves also expect to have a choice - a tendency that holds both for European and Anglo-Saxon countries with different educational systems (Jones et al., 2020; Kumi-Yeboah et al., 2023; Sun et al., 2019; Wollny et al., 2023; Whitelock-Wainwright et al., 2021; Ifenthaler and Schumacher, 2019). Co-design initiatives are common to educational technology cycles (Buckingham Shum et al., 2024) but they do not address participatory decision-making around data. They involve students in participatory processes related to the tool design but rarely focus on continuous participation and agency over one's data-related decisions. By giving learners a voice in how their data are used, we can ensure that educational technologies are developed in a way that respects learners' rights and promotes their well-being.

Developing a learner-centred intervention for bottom-up participation in data-sharing and learning faces two main challenges: 1. overcoming traditional limitations of participatory processes, and 2. accounting for contextual variations where data sharing takes place.

In educational settings, participatory processes surrounding informed consent are complex to implement for two main reasons. First, the effectiveness of informed consent is controversial as individuals often do not engage with consent meaningfully. Researchers have long shown that people avoid reading privacy statements and user agreements (Coles-Kemp and Kani-Zabihi, 2010), even when the stakes are high (Cassileth et al., 1980; Byrne et al., 1988). Students have low information retrieval upon reading consent forms (Beardsley et al., 2020), miss critical information in the consent, even when prompted to pay more attention by guiding questions (Knepp, 2018). Students are also more likely to carelessly read consent when studying remotely (Pedersen et al., 2011). Second, in educational settings informed consent is challenged by power and information imbalances. When consent is requested by a teacher, it may evoke the power-relations underpinning classroom interactions (Clark et al., 2022). If students decide on their own, they may have insufficient information or limited decision-making skills to fully engage with this decision (David et al., 2001).

Besides the traditional imitations of participatory processes, accounting for contextual variations of data sharing decision represents another challenge. Current consent practices rely on strict privacy laws to protect data, rather than address contextual factors related to agentic engagement with consent. Scholarship has long questioned if top-down privacy regulations and rigid protocols put in place by the ethical boards are at all suited to address alleged situational differences shaping consent decisions (Solove, 2012). In learning analytics, situations when decisions need to be made around the use of data are nuanced and contextual (Kitto and Knight, 2019), challenging the notion of consent as 'disclosure stated for good' (Luger and Rodden, 2013). However, so far, little is known about how contextual differences shape learner perceptions of data sharing acceptability. In learning analytics, researchers predominantly focus on psychometric instruments measuring generic attitudes towards data sharing in university (Whitelock-Wainwright et al., 2019; Mutimukwe et al., 2022). Results of very few observational studies on situational perceptions of learning analytics (Ifenthaler and Schumacher, 2019; Korir et al., 2023) suggest that learning data decisions may be, to some extent, contextual.

Attempting to address both challenges, we setup an in-person experiment to find out which contextual factors influence the perceived acceptability of sharing learning data (see Figure 1). As consent to share learning data requires that individuals engage in a meaningful deliberation about the decision, we combined contextual learning practices with forms of collective decision making and group deliberation. We drew on the decision-making literature around the wisdom of inner crowd (Herzog and Hertwig, 2014a), wisdom of crowds and interactive groups (Bahrami et al., 2010; Dezechache et al., 2022; Navajas et al., 2018; Niella et al., 2016) to explore ways of active participatory consent in the learning context. The situational nature of data sharing decisions can be explored using the theory of privacy as contextual integrity (Nissenbaum, 2004, 2009). These characteristics include the information subject, sender, receiver, attributes of the information shared, and transmission principles, i.e. the norms governing how information flows from

data sender to data receiver (see also Martin and Nissenbaum (2017), Silber et al. (2022), Gerdon et al. (2020), Longin and Deroy (2024), Longin et al. (2023) for empirical validation studies around contextual integrity of data privacy).

We find a consistently lower acceptability in and after the group discussion compared to the initial, individual baseline rating (see Figure 1C). Surprisingly, we found no difference in acceptability of learning data types (see Figure 1D). Participants rated sharing outcome-related learning data (final performance measures like grades) as acceptable as process-related learning data (behaviour during learning process like sequence of clicks). Participants instead were sensitive to the data recipient and the purpose for which data was shared. In particular, participants found sharing learning data with private companies more acceptable than sharing with public governments, and preferred sharing data for the individual over sharing for the collective benefit. Additional epistemic network analysis using the content of the recorded group discussions validated these effects. Overall, we contribute a novel method to study student’s data sharing preferences promoting to advance the development of bottom-up participatory consent practices over additional top-down learning protocols.

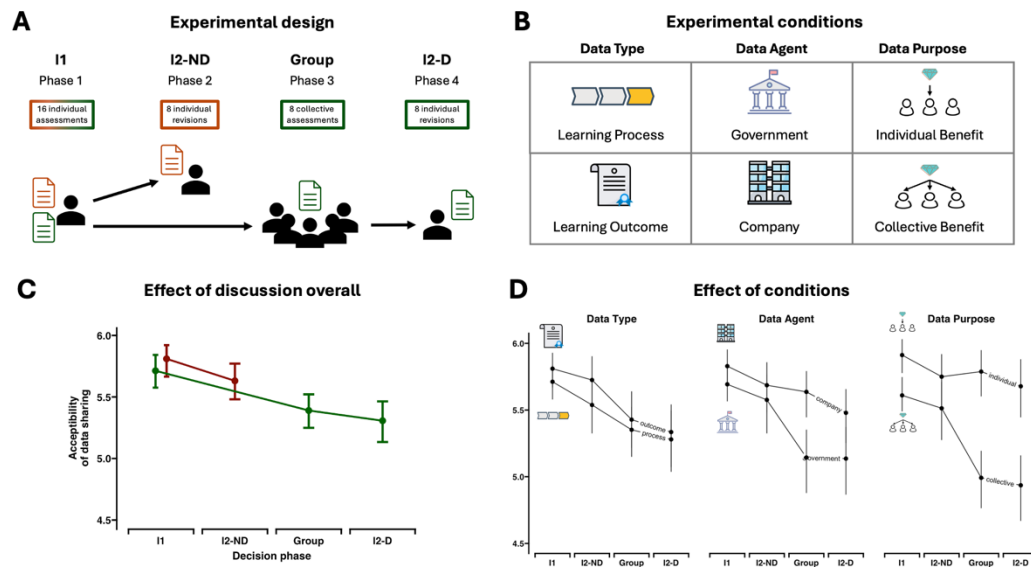


Figure 1. Experimental design and main effects. **A:** We tested participants’ perceived acceptability of sharing learning using contrasting vignettes in a multi-phase experimental design. **B:** Vignettes varied across three factors with two levels each: data type (process vs outcome), data agent (government vs company), and data purpose (individual vs collective benefit). **C:** Progressing through each phase, participants overall find data sharing less acceptable. **D:** Data Sharing with governments and for collective benefit are the main drivers for the decrease in data sharing acceptability in and after group discussion. **C&D:** Plotted are mean values and 95% confidence intervals obtained from resampling the collected data using the bias-corrected and accelerated (bca) bootstrap method (Canty, 2002).

Methods

Participants. We targeted and recruited 12 groups of five participants via the Munich Experimental Laboratory for Economic and Social Sciences (Melessa). We targeted uneven participant numbers per group to ensure a clear group consensus. Each group was gender balanced with either three males to two females or vice-versa resulting in six male-dominant and six female-dominant groups. Our sample size estimation followed similar experiments (Keshmirian et al., 2022; Myers and Kaplan, 1976) as well as previous pilot data. Each experimental session lasted approximately 60 minutes involving one group of five participants. There were no minimal educational requirements for participants except being fluent in speaking and reading English. Out of 60 participants, 27 participants reported a university entrance qualification (general or subject-related university entrance qualification/Abitur [high school or EOS]), while

20 participants held an additional bachelor degree. The mean age of participants was 23.92 ± 3.85 SD years old. Variance in age composition across groups was relatively low and stable (min SD = 0.45 at a mean age of 24.2 years old; max SD = 9.34 at a mean age of 26.2 years old).

Design. We used an in-person within-subject experimental design adapted from the wisdom of crowds and collective decision-making literature (Bahrami et al., 2010; Dezechache et al., 2022; Herzog and Hertwig, 2014a,b; Longin et al., 2023; Myers and Kaplan, 1976; Navajas et al., 2018; van Dolder and van den Assem, 2018; Longin and Deroy, 2024; Longin et al., 2023) to test the impact of contextual learning factors on the perceived acceptability of sharing learning data on an individual and a group-level. Applying the contrastive vignette technique (Burstin et al., 1980), we varied three main factors with two levels each: data type (process vs outcome learning data), data recipient (private company vs public government), and data collection purpose (individual vs collective benefit) (see Figure 1B). By combining one level of each factor with each other, we obtain eight possible vignettes and experimental conditions. Each vignette was designed with a unique background story to avoid any transfer effects. In addition, each vignette had two versions that differed only in their background story. While all versions and vignettes were rated individually in the first phase of the experiment, only one version was rated individually in phase 2. The other version was discussed and rated in phases 3 and 4 (see Figure 1A). The assignment of versions was randomised and counterbalanced. The experiment consists of a total of 16 vignettes mapped onto eight experimental conditions (see supplementary methods for a detailed overview of all vignette variations). This experimental design allows to compare the impact of each experimental factor on the perceived acceptability of sharing learning data within a set of diverse stimuli.

Measures. To measure the perceived acceptability of data-sharing, we asked: 'How acceptable is it to use data for this reason?'. Responses were collected on a 7-point Likert scale ranging from 1 (not at all acceptable) to 4 (neutral) and 7 (completely acceptable). In addition, with full consent of the participants, we recorded and transcribed the audio of the group discussions. Content analysis of transcripts was implemented using a coding framework derived from contextual integrity theory. Labelled data were analysed via ordered epistemic network analysis (Shaffer et al., 2016; Tan et al., 2022; see below for more information).

Procedure. Throughout the experiment, participants were seated in a U-shaped seating arrangement, facing each other and the wall where the experimental instructions and vignettes were projected for everyone to see. After having been shown experimental instructions and a practice scenario, participants started the main experiment. Presented with one vignette at a time, participants had one minute to rate each vignette in the first phase, half a minute in the second and fourth stage, and two and a half minutes in the third phase. Participants noted down their ratings on sheets of paper which were replaced with new ones after each experimental phase. After completing the fourth phase of the main experiment, participants were given a final sheet with demographic questions. The demographics question involved age, gender, highest completed educational degree, general trust in public authorities and private companies, general privacy concerns, sensitivity of learning data, a validated 11-item social conformity questionnaire (Mehrabian and Stefl, 1995), five items to capture the social experience during the group discussion (Aron et al., 1992; Sprecher, 2021; Sun et al., 2020), and an open-text box for feedback.

Discussion analysis. Group discussions were recorded, transcribed, and labelled using a coding scheme developed from the contextual integrity theory (see supplementary for details). To analyse differences between the presence of contextual integrity characteristics in group discussions, we applied ordered network analysis (Tan et al., 2022) to the coded transcript data using the epistemic network analysis web tool (ENA, version 1.7.0) (Marquart et al., 2018). This method bridges qualitative and quantitative data analysis (Shaffer, 2017; Shaffer et al., 2016; Siebert-Evenstone et al., 2017). It combines content analysis with network analysis to examine the thematic codes within each observation (text from an individual, a group, a time-bounded session) as a network. In such a network, content analysis codes are network nodes, and their co-occurrences are network ties. To compare multiple structures comprised of the same set of codes, ENA tool projects each network graph into the same space, using single value decomposition. This enables to describe all graphs in relation to each other and compare them statistically and visually. We applied this technique to compare 96 networks of contextual integrity themes at the level of a vignette conversation within each

group (eight conversations per group), i.e. 'group-vignette' thematic networks (Figure 3 A). We also statistically compared 18 networks of contextual integrity themes at the level of each vignette, with themes aggregated across all groups (Figure 3 B). These structures were compared statistically between experimental conditions and for vignettes rated with the highest and the lowest acceptability. Details of the content and epistemic network analysis are presented in the supplementary material.

Ethics. The study was approved by the School of Advanced Study, University of London Ethics committee. All participants provided informed consent.

Results

Acceptability of sharing learning data decreases in and after group discussion

Acceptability ratings. Participants rated all items individually in the first stage, half of them again individually in the second stage, discussed the other half openly in the group in the third stage, and rated the discussed items again individually in the fourth stage. To analyse a general effect of group discussion, we fitted a linear mixed model to predict acceptability ratings of sharing learning data with decision stages as a fixed effect. The experimental setup consists of four decision stages (see methods). The model included the unique participant ID and associated group ID as random effects (formula: $\text{rating} \sim \text{stage} + (1 | \text{group/participant})$).

Participants on average found sharing learning data less acceptable in and after a group discussion compared to their initial, individual ratings. The average rating on a 7-point scale for the group discussion in stage 3 is by 0.3 points significantly lower (95% CI [-0.45, -0.14], $p < .001$) compared to the initial ratings in stage 1. Similarly, the average rating in stage 4 is by 0.38 points significantly lower (95% CI [-0.53, -0.23], $p < .001$) than the initial ratings in stage 1. Notably, ratings in stage 2 were non-significantly different from those in stage 1 ($b = -0.14$, 95% CI [-0.30, 0.01], $p = 0.06$). The impact of group discussion becomes further evident when comparing individual ratings before (stage 2) and after (stage 4) group discussion. Here, the average rating was significantly lower by 0.23 points (95% CI [-0.46, 0], $p = 0.05$) in stage 4 compared to stage 2.

Sharing process vs outcome learning data makes no difference to data sharing acceptability

Acceptability ratings. We compared two kinds of learning data: outcome-related learning data (final performance measures like grades) and process-related learning data (behaviour during learning process like sequence of clicks). To find out whether participants' acceptance of sharing learning data depends on the presented data type, we fitted two linear mixed models to predict acceptability ratings of sharing learning data by *learning data types*: a general model only predicting a difference in learning data type (formula: $\text{rating} \sim \text{type} + (1 | \text{group/participant})$), and a more specific model including an interaction term with the decision phases (formula: $\text{rating} \sim \text{type} * \text{phase} + (1 | \text{group/participant})$). Both models included the unique participant ID and associated group ID as random effects.

Overall, we find that participants rated the sharing of process-related learning data significantly lower than outcome-related learning data ($b = -0.14$, 95% CI [-0.25, -0.02], $p = 0.017$; see Figure 2A1). However, when including decision phases as an interaction term in the second model, the general effect of data type breaks down (see Figure 2A2). Participants found the sharing of process learning data as acceptable as outcome-related learning data. Process-data ratings are on average 0.07 points non-significantly lower than outcome-data ratings in stage 1 (95% CI [-0.17, -0.03], $p = 0.11$), 0.11 points lower in stage 2 (95% CI [-0.25, -0.04], $p = 0.1$), 0.05 points lower in stage 3 (95% CI [-0.19, 0.09], $p = 0.41$), and 0.04 points lower in stage 4 (95% CI [-0.18, 0.1], $p = 0.55$).

Discussion structure. 'Group-vignette' thematic networks (Figure 3 A3) were compared across discussions focused on sharing process learning data versus outcome learning data. A Mann-Whitney test showed that they were statistically different along the x-axis ($\text{Mdn}=0.06$, $N=48$ $U=706.00$, $p=0.00$, $r=0.39$), explaining 6% of network structure variance. Vignette-level thematic networks (Figure 3 B3) were also statistically significantly different along the x-axis ($\text{Mdn}=-0.12$, $N=8$ $U=9.00$, $p=0.01$, $r=0.72$), explaining 14% of network structure variance.

Surprisingly, participants discussed the relationship between outcome data, who receives it and what benefit would that offer in ways where they questioned the utility of this common data from learning environments. In process data, the focus on transmission norms were more prominent. Yet, some participants were both positive and negative about these two data types, using other context information to make sense of their decision. With outcome data, they could be concerned that this outcome would be misused and known to others. With process data, they often failed to see its relationship to learning, e.g.: “I don’t know about the timing of clicks and sequence of language exercises (*data attribute, utility*). I think it depends on each person. It has nothing to do with [learning]”. The nature of discussions suggests that while LA researchers consider performance (outcome) data and learning (process) data as different types, learners potentially make sense of the data types in other ways, more so as “what are these data a proxy for”. If learning outcome data are a proxy for failure or aptitude, there may be a risk, despite these data perhaps being the most commonly used in educational applications.

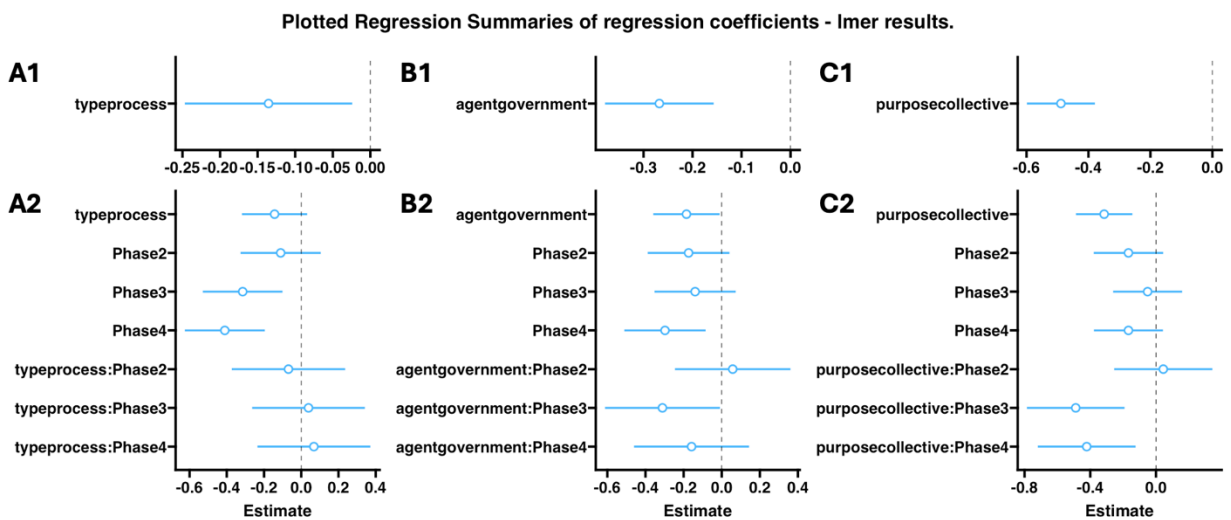


Figure 2. Plotted Regression Summaries of lmer regression coefficients. **A1:** Overall effect of learning data type: process-related learning data was rated on average lower than outcome-related learning data. **A2:** Effect of learning data type by decision phases: comparing data types by decision phases reveals no difference in data types. **B1:** Overall effect of data recipient/agent: sharing learning data with governments was rated less acceptable than with companies on average. **B2:** Effect of data recipient by decision phases: sharing learning data with governments was rated less acceptable overall as well as in phases one and three. **C1:** Overall effect of data sharing purpose: sharing learning data for a collective benefit was rated less acceptable than sharing for an individual benefit. **C2:** Effect of data purpose by decision phases: sharing learning data for a collective rather than individual benefit was rated less acceptable overall as well as in decision phases three and four.

Sharing learning data for an individual rather than collective benefit is more acceptable

Acceptability ratings. Learning data can either be shared for a collective or an individual benefit. Sharing for a collective benefit includes cases of developing global teaching best practices, whereas sharing for an individual benefit capture cases of personalised learning recommendations. We find that participants rated the data sharing acceptability for an individual benefit consistently higher than for a collective benefit. We fitted two linear mixed models to predict acceptability ratings of sharing learning data by *sharing purpose*: a general model only predicting a difference in sharing purpose (formula: rating ~purpose + (1 | group/participant)), and a more specific model including an interaction term with the decision phases (formula: rating ~purpose*stage + (1 | group/participant)). Both models included the unique participant ID and associated group ID as random effects.

Overall, participants rated sharing learning data for a collective benefit by 0.49 points as significantly less acceptable than sharing learning data for an individual benefit (95% CI [-0.60, -0.38], $p < .001$; see Figure 2C1). The

overall effect holds when including the decision phases as an interaction term within the more specific model (see Figure 2C2). Already in the first phase, participants rated sharing learning data for a collective benefit by 0.32 points as significantly less acceptable than sharing for an individual benefit (95% CI [-0.49, -0.14], $p < .001$). While the effect was non-significant in the second phase ($b = 0.04$, 95% CI [-0.25, 0.34], $p = 0.771$), it was significant again in the third ($b = -0.49$, 95% CI [-0.79, -0.19], $p = 0.001$) and fourth phase ($b = -0.42$, 95% CI [-0.72, -0.12], $p = 0.006$). The effect of the data purpose on the acceptability of sharing data is hence strongest in the group discussion. Notably, the general decrease of data sharing acceptability of stages three and four compared to stages one and two are non-significant in the larger model.

Discussion structure. 'Group-vignette' networks (Figure 3 A2) were compared across discussions focused on sharing data for individual and collective benefit. A Mann-Whitney test showed that they were statistically different along the x-axis (Mdn=0.08, $N=48$ $U=618.00$, $p=0.00$, $r=0.46$), explaining 8% of network structure variance. Vignette-level thematic networks (Figure 3 B2) were not statistically significantly different. In vignettes with data sharing for individual benefit participants focused on the benefit elements against transmission norms.

In contrast, in data sharing for collective benefit, they tended to focus on transmission norms governing the collective sharing of their data and sensitivities surrounding data they are submitting. One may speculate that their risk perceptions was less associated with the entity receiving the data, more so on the others who will have access to these data, despite it being de-identified and anonymised. As exemplified here: "I think anonymised is a problem because you are recorded, they hear your voice or maybe see you (*data attribute*), so not very anonymous (*transmission norm*)", with another participant responding: "Yeah, but it depends on how they use the data of yourself (*data attribute*), because anonymised to me means that the video of myself (*data attribute*) won't be shown anywhere else. But maybe the contents of what I'm doing and how I performed will be used for continuously improve the website." Exception to these were the data used to teach medical emergency skills where participants were open to collective data sharing, as the quotes shows: "It's surgery and medical students (*data subject*), so it's important. For sure, it's a seven."

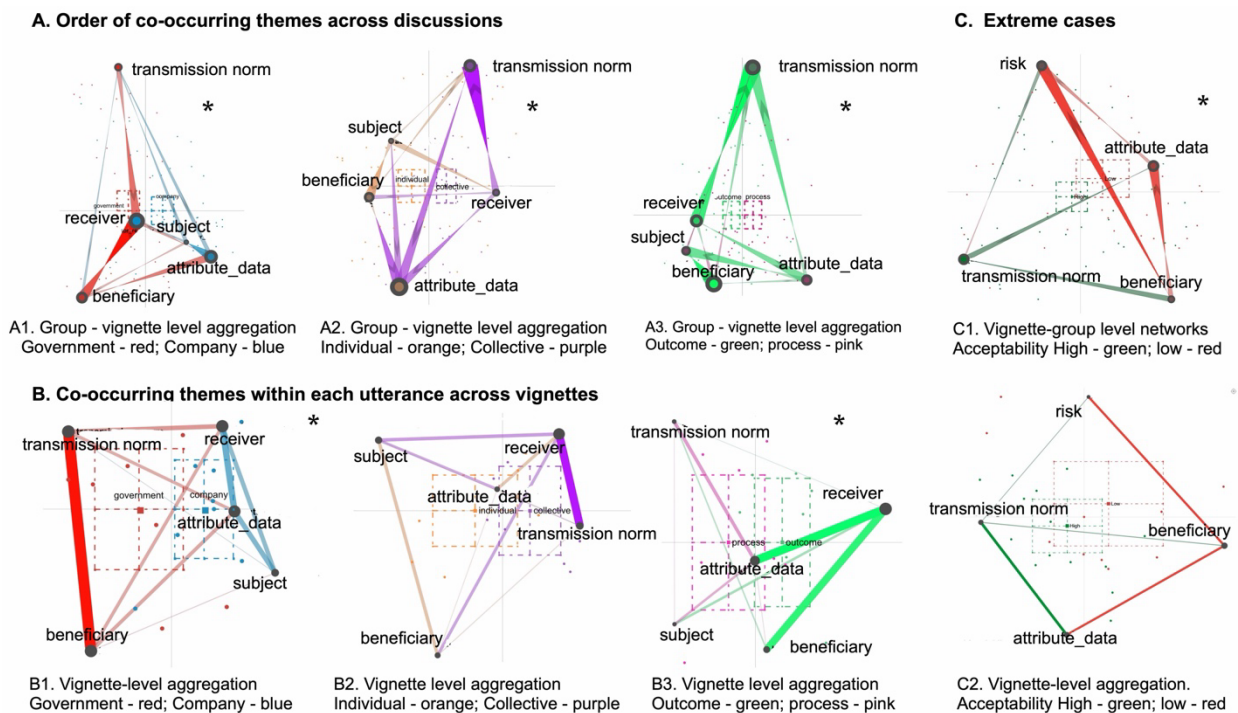


Figure 3. Subtracted plots of co-occurring discussion themes. Nodes represent contextual integrity themes assigned to each participant utterance; the size of a node represents if a theme auto-recurred. If two contextual integrity themes co-occurred in a

discussion, they share an edge projected in the two-dimensional space. Edges are directional, with thicker edges and error directions representing the sequence of co-occurrence. Subtraction plots only show edges that differed across two juxtaposed contextual integrity conditions. Plotted networks were averaged across all networks within the same condition. **A.** Subtracted ordered network plots per condition aggregated at the level of a group and vignette. **B.** Subtracted ordered network plots per condition aggregated at the level of vignette across groups. **C.** Sample of subtracted co-occurrence networks. (1) ordered at the level of the group and vignette for cases when vignettes were rated extremely low (n=26) and high (n = 58). (2) Unordered co-occurrence networks at the level of vignette across groups, for vignettes rated extremely low (n=11) and for vignettes rated extremely high (n=16).

Sharing learning data with companies is more acceptable than sharing data with governments

Acceptability ratings. Learning data can be shared with different institutions. Comparing participants' acceptability ratings of sharing data with private companies versus public governments revealed that participants found sharing learning data with companies more acceptable than sharing learning data with governments. We fitted two linear mixed models to predict acceptability ratings of sharing learning data by *data recipient*: a general model only predicting a difference in sharing purpose (formula: rating ~agent + (1 + group/participant)), and a more specific model including an interaction term with the decision phases (formula: rating ~agent*stage + (1 + group/participant)). Both models included the unique participant ID and associated group ID as random effects.

Overall, participants rated the acceptability of sharing learning data with the government significantly lower than sharing learning data with companies (b = -0.27, 95% CI [-0.38, -0.16], p < .001; see Figure 2B1). When including the decision phases as a regressor in the second model, the average decrease in acceptability for sharing learning data with governments rather than companies remains significant for the first (b = -0.18, 95% CI [-0.36, -0.01], p = 0.037; see Figure 2B2), and the third phase (b = -0.31, 95% CI [-0.61, -0.001], p = 0.043). However, the difference between sharing learning data with governments and companies was non-significant in the second (b = 0.06, 95% CI [-0.25, 0.36], p = 0.709) and fourth decision phases (b = -0.16, 95% CI [-0.46, 0.14], p = 0.304).

Discussion structure. A Mann-Whitney test showed that group-vignette networks compared between government versus company-oriented scenarios were statistically different along the x-axis (Mdn=0.07, N=48 U=629.00, p=0.00), explaining 7% of network structure variance. As per the Mann-Whitney test comparing vignette thematic networks (Fig. 3 B1), discussions of government versus company-oriented scenarios were statistically significantly different along the x-axis (Mdn=0.12, N=8 U=54.00, p=0.02, r=-0.69), explaining 29% of network structure variance.

When talking about data sharing with government, groups discussed government and regulations or intentions to process and use the data. The risk of a particular data type owned by the government was prominent, as demonstrated in this quote: "It's a public authority (data receiver) [offering] driving education for everyone in the world (beneficiary). How does that data get to everyone in the world (*transmission norm; beneficiary*)? Is it sold?". In contrast, the purpose of transmission norms was less prominent in discussions with company as data receiver. In company-related scenarios, participants focused on data types, reflecting on their utility, for example: "It's also like for personalized recommendations (*beneficiary*) so yeah it's really helpful although it's a company (*data receiver*). Timing of clicks and sequences is good to get recommendations and individual support (*attribute data*)".

Discussions differ between scenarios with high and low acceptability

Although thematic networks differed across conditions, the same themes co-occurred within multiple conditions. The results, therefore, are limited in clarifying the relationship between the differences in discussions and levels of acceptability. To identify if particular thematic patterns were associated with more or less acceptability, we constructed thematic networks at the vignette-group and vignette levels (Figure 3C1 and C2) only for scenarios rated extremely low (four or less) and extremely high (seven). Low acceptability of four was not only descriptive of the lower end of the distribution but also reflective of the participants' attitudes to these numbers as more negative, as gleaned from the transcripts. Differences across group-vignette networks were statistically significantly different on x-axis component, explaining 20% variability in structures representative of groups and vignettes (Mdn=0.00, N=58 U=530.50, p=0.03, r=0.3). Vignette networks were not statistically significantly different.

High acceptability scenarios were more likely to contain co-occurring themes of risk, transmission norm, and attribute data, whereas low acceptability scenarios were more likely to contain co-occurring themes of attribute data, risk, and beneficiary. Higher acceptability appeared in contexts interpreted as safe for this specific data to be transferred under these conditions and reasonable for the utility suggested. In contrast, if participants interpreted the context as a setting where the data they were asked to share could indicate something about them that they may not want others to know, they were less likely to find this acceptable. The same low acceptability would apply to contexts where the participants did not understand the benefit, even if evidence unknown to them would indicate that submitting these data may positively impact their outcomes. These suggest that context characteristics defined by contextual integrity theory may be interpreted in relation to the larger themes of risk and utility, which were dominant codes connected to other themes.

General trust levels predict data sharing attitudes

After completing the main experiment, participants answered demographic questions, including age, gender, highest completed educational degree, general trust in public authorities and private companies, general privacy concerns, sensitivity of learning data, a social conformity questionnaire, and five items to capture the social connectedness during the group discussion (Aron et al., 1992; Sprecher, 2021; Sun et al., 2020). To analyse demographics effects, we fitted multiple linear mixed models. We find that neither age, nor gender, nor education made a considerable impact on average data acceptability ratings (see supplementary results for a full breakdown of all demographic effects). Male participants rated the acceptability of data sharing non-significantly lower on average than females ($b = -0.34$, 95% CI [-0.71, 0.03], $p = 0.072$). Being older had no significant influence on the average data sharing ratings ($b = -0.03$, 95% CI [-0.09, 0.02], $p = 0.192$). Similarly, having obtained an university entrance qualification ($b = 1.18$, 95% CI [-0.28, 2.64], $p = 0.113$), a bachelor degree ($b = 1.05$, 95% CI [-0.46, 2.55], $p = 0.173$), or a master degree ($b = 0.43$, 95% CI [-1.07, 1.92], $p = 0.574$) had no significant influence on the average data sharing rating compared to the participants without a completed educational degree.

Examining general trust levels, however, revealed an expected effect: participants with higher trust levels in companies found data sharing with companies on average significantly more acceptable ($b = 0.14$, 95% CI [0.01, 0.28], $p = 0.034$), while participants with higher trust levels in governments found data sharing with governments on average significantly more acceptable ($b = 0.12$, 95% CI [0.005, 0.23], $p = 0.04$). Notably, general trust levels in companies reliably predicted only higher data sharing acceptability with companies, not with governments. Likewise, general trust levels in governments reliably predicted only higher data sharing acceptability with governments, not with companies. Both general trust levels correlate positively ($r = 0.55$, $t = 4.99$, $p < 0.001$). The effects of reported general trust ratings on data sharing acceptability are further validated by the effect of general privacy concerns on data sharing acceptability. We find that participants who were more concerned with privacy in general had on average significantly lower data sharing acceptability ratings ($b = -0.39$, 95% CI [-0.66, -0.13], $p = 0.003$).

Discussion

Educational technologies like online learning platforms or virtual assistants rely on students' learning data to offer personalised learning support. Some learning support may include the adaption of learning content to each student's individual progression; other may leverage learning data to inform educational policies and provide personalised scaffolds to enhance student learning outcomes. Despite potential advantages, risk management and full consent from learners are essential for data-driven educational technologies. As a part of this process, the design and implementation of the participatory consent processes that remain a work in progress need to be further developed.

In this paper, we have proposed a novel interactive methodology that takes a step towards rethinking consent practices by putting the learners first. Sharing learning data turns from a passive necessity into an active, participatory process. In fact, we developed a new way to conceptualise and realise active learner participation in the data-driven

learning process by combining work on the contextual integrity of data sharing with the dynamics of social decision-making. With an in-person experimental design, we tested our hypothesis of making learners more sensitive to the context of sharing learning data. In four sequential decision-making phases participants rated the acceptability of sharing learning data in different learning contexts. The learning context varied in terms of three experimental factors: learning data type (process vs outcome), data sharing purpose (individual vs collective benefit), and data recipient (private company vs public government).

We found out that (1) sharing learning data does depend on the context, as participants favoured sharing data for the individual over the collective benefit and sharing data with companies over governments while being agnostic to the learning data type; and that (2) group discussions were the main catalysts for this increase in data sharing context-sensitivity. It was only in and after the group discussion that participants distinguished between contextual data types consistently.

Unexpectedly, the type of learning data, whether process- or outcome-related, had no significant impact of the acceptability of data sharing. Regardless of the data type, participants become more cautious about sharing learning data after discussing possibly problematic aspects with others. The context-dependence of data sharing acceptance fits general trend of contextual drivers of data privacy observed in the empirical literature around contextual integrity. Silber et al. (2022) found that people were more open to share biomarker and medical records than sensor data in health, with universities being the most accepted data recipients. In contrast, Gerdon et al. (2020) showed that individuals were more likely to share health data for public benefit, as this tendency increased throughout the COVID-19 pandemic.

The analysis of the group discussions provides further insight into the role of social decision-making in shaping learners' attitudes towards sharing learning data. The results indicate that the group discussion was a crucial catalyst for learners to become more cautious when sharing learning data, as it was only in and after these discussions that participants consistently distinguished between contextual data types. A closer examination of the discussion dynamics reveals that the kind the data recipient and the purpose of data sharing influenced the focus of the conversations. When discussing data sharing with government, participants emphasized the risks associated with government ownership and regulation of data, whereas discussions centred on companies as data recipients focused on the utility of specific data types.

In scenarios where data sharing was intended for individual benefit, participants prioritised the benefits over transmission norms, whereas collective benefit scenarios led to a focus on transmission norms and sensitivities surrounding specific data. Notably, the exception to this pattern was the sharing of data for teaching medical emergency skills, where participants were more open to collective data sharing. The thematic analysis of the discussions reveals that high acceptability scenarios were characterized by the co-occurrence of themes related to risk, transmission norms, and attribute data, suggesting that learners were more willing to share data in contexts perceived as safe and where the utility of sharing data was clear. In contrast, low acceptability scenarios were marked by the co-occurrence of themes related to attribute data, risk, and beneficiary, indicating that learners were less likely to share data in contexts where the benefits were questionable or where the data could reveal sensitive information about the sender. These findings suggest that learners interpret the context of data sharing through the lens of risk and utility, which are dominant themes connected to other factors such as transmission norms and beneficiary.

In sum, this paper marks an important step in understanding participatory and ethical data sharing in learning situations and educational technologies as contextual and necessarily social. The analysis of data acceptability ratings and group discussions highlights the importance of social decision-making in shaping learners' attitudes towards data sharing and underscores the need for a contextual integrity framework that considers the complex interplay of factors influencing learners' willingness to share their data in an active and participatory way.

Due to the scope of the experiment, several important questions remain. The experiment did not consider university as the receiver of the learning data, nor it unpacked whether learners were sensitive to their data being available only

to fellow students or the instructor. However, these scenarios are common to learning data collection. Cultural differences would likely further shape whether these receivers would be identified as public or private, and these beneficiaries viewed as collective or individual benefit. Learner cautiousness in submitting data for collective benefit is particularly concerning. It remains unknown if it results from the lack of understanding of the benefit, resistance to contribute, or lack of trust in transmission norms. The insight that learners did not differentiate between the acceptability of learning data is surprising. Learning data are a lot more diverse than included in the vignette scenarios and deeper investigation if the lack of differentiation would hold is required. Another question is whether further studies should strive to understand if researcher-driven conceptualisations of learning data are at all aligned with learner understanding of differences in the data collected about them. Finally, the question of mechanisms driving the effect of the group remains relevant. Group discussion may be an interesting form of classroom-based consent, but it is not easy to scale. Future work should focus on further disentangling the mechanisms and understanding the biases embedded in the interactive consent processes to create a new generation of learner-centric consent practices.

References

- Aron, A., Aron, E. N., and Smollan, D. (1992). Inclusion of other in the self scale and the structure of interpersonal closeness. *Journal of personality and social psychology*, 63(4):596.
- Bahrami, B., Olsen, K., Latham, P. E., Roepstorff, A., Rees, G., and Frith, C. D. (2010). Optimally interacting minds. *Science*, 329(5995):1081–1085.
- Beardsley, M., Martínez Moreno, J., Vujovic, M., Santos, P., and Hernández-Leo, D. (2020). Enhancing consent forms to support participant decision making in multimodal learning data research. *British Journal of Educational Technology*, 51(5):1631–1652.
- Brooks, C., Quintana, R. M., Choi, H., Quintana, C., NeCamp, T., and Gardner, J. (2021). Towards culturally relevant personalization at scale: Experiments with data science learners. *International Journal of Artificial Intelligence in Education*, 31:516–537.
- Buckingham Shum, S., Martínez-Maldonado, R., Dimitriadis, Y., and Santos, P. (2024). Human-centred learning analytics: 2019–24.
- Burstein, K., Doughtie, E. B., and Raphaeli, A. (1980). Contrastive Vignette Technique: An indirect Methodology Designed to Address Reactive Social Attitude Measurement1. *Journal of Applied Social Psychology*, 10(2):147–165.
- Byrne, D., Napier, A., and Cuschieri, A. (1988). How informed is signed consent? *Br Med J (Clin Res Ed)*, 296(6625):839–840.
- Canty, A. J. (2002). Resampling Methods in R: The boot Package. *The Newsletter of the R Project*, 2.
- Cassileth, B. R., Zupkis, R. V., Sutton-Smith, K., and March, V. (1980). Informed consent—why are its goals imperfectly realized? *The New England journal of medicine*, 302(16):896–900.
- Clark, A. T., Ahmed, I., Metzger, S., Walker, E., and Wylie, R. (2022). Moving from co-design to co-research: Engaging youth participation in guided qualitative inquiry. *International Journal of Qualitative Methods*, 21:16094069221084793.
- Coles-Kemp, L. and Kani-Zabihi, E. (2010). On-line privacy and consent: a dialogue, not a monologue. In *Proceedings of the 2010 New Security Paradigms Workshop*, pages 95–106.
- David, M., Edwards, R., and Alldred, P. (2001). Children and school-based research: ‘informed consent’ or ‘educated consent’? *British educational research journal*, 27(3):347–365.
- Dezecache, G., Dockendorff, M., Ferreiro, D. N., Deroy, O., and Bahrami, B. (2022). Democratic forecast: Small groups predict the future better than individuals and crowds. *Journal of Experimental Psychology: Applied*, 28(3):525.
- Drachsler, H. and Greller, W. (2016). Privacy and analytics: it’s a delicate issue a checklist for trusted learning analytics. In *Proceedings of the sixth international conference on learning analytics & knowledge*, pages 89–98.
- El-Sabagh, H. A. (2021). Adaptive e-learning environment based on learning styles and its impact on development students’ engagement. *International Journal of Educational Technology in Higher Education*, 18(1), 53.
- Ferguson, R., Brasher, A., Clow, D., Cooper, A., Hillaire, G., Mittelmeier, J., Rienties, B., Ullmann, T., and Vuorikari, R. (2016). Research evidence on the use of learning analytics: Implications for education policy. "JRC Science for Policy Report, funded by the European Commission".
- Gašević, D., Tsai, Y.-S., and Drachsler, H. (2022). Learning analytics in higher education—stakeholders, strategy and scale.
- Gerdon, F., Nissenbaum, H., Bach, R. L., Kreuter, F., and Zins, S. (2020). Individual acceptance of using health data for private and public benefit: Changes during the covid-19 pandemic. *Harvard Data Science Review*.
- Herzog, S. M. and Hertwig, R. (2014a). Harnessing the wisdom of the inner crowd. *Trends in cognitive sciences*, 18(10):504–506.
- Herzog, S. M. and Hertwig, R. (2014b). Think twice and then: Combining or choosing in dialectical bootstrapping? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 40(1):218.
- Ifenthaler, D. and Schumacher, C. (2019). Releasing personal information within learning analytics systems. *Learning technologies for transforming large-scale teaching, learning, and assessment*, pages 3–18.
- Joksimovic, S., Marshall, R., Rakotoarivelo, T., Ladjal, D., Zhan, C., and Pardo, A. (2021). Privacy-driven learning analytics. In *Manage your own learning analytics: Implement a Rasch modelling approach*, pages 1–22. Springer.
- Jones, K. M., Asher, A., Goben, A., Perry, M. R., Salo, D., Briney, K. A., and Robertshaw, M. B. (2020). “we’re being tracked at all times”: Student perspectives of their privacy in relation to learning analytics in higher education. *Journal of the Association for Information Science and Technology*, 71(9):1044–1059.
- Kasneji, E., Seßler, K., Küchemann, S., Bannert, M., Dementieva, D., Fischer, F., Gasser, U., Groh, G., Günemann, S., Hüllermeier, E., et al. (2023). Chatgpt for good? on opportunities and challenges of large language models for education. *Learning and individual differences*, 103:102274.

- Keshmirian, A., Deroy, O., and Bahrami, B. (2022). Many heads are more utilitarian than one. *Cognition*, 220:104965.
- Kitto, K. and Knight, S. (2019). Practical ethics for building learning analytics. *British Journal of Educational Technology*, 50(6):2855–2870.
- Knepp, M. M. (2018). Using questions to improve informed consent form reading behavior in students. *Ethics & Behavior*, 28(7):560–577.
- Knoop-van Campen, C. A., Wise, A., and Molenaar, I. (2023). The equalizing effect of teacher dashboards on feedback in k-12 classrooms. *Interactive Learning Environments*, 31(6):3447–3463.
- Korir, M., Slade, S., Holmes, W., Héliot, Y., and Rienties, B. (2023). Investigating the dimensions of students' privacy concern in the collection, use and sharing of data for learning analytics. *Computers in human behavior reports*, 9:100262.
- Kumi-Yeboah, A., Kim, Y., Yankson, B., Aikins, S., and Dadson, Y. A. (2023). Diverse students' perspectives on privacy and technology integration in higher education. *British Journal of Educational Technology*, 54(6):1671–1692.
- Liu, Q. and Khalil, M. (2023). Understanding privacy and data protection issues in learning analytics using a systematic review. *British Journal of Educational Technology*, 54(6):1715–1747.
- Longin, L. and Deroy, O. (2024). Collective decisions on data privacy-follow-up. Pre-registration.
- Longin, L., Kreuter, F., and Deroy, O. (2023). Collective decisions on data privacy. Pre-registration.
- Luger, E. and Rodden, T. (2013). An informed view on consent for ubicomp. In *Proceedings of the 2013 ACM international joint conference on Pervasive and ubiquitous computing*, pages 529–538.
- Marquart, C., Hinojosa, C., Swiecki, Z., Eagan, B., and Shaffer, D. (2018). Epistemic network analysis (version 1.7. 0)[software]. Available from app. epistemicnetwork. org.
- Martin, K. and Nissenbaum, H. (2016/2017). Measuring Privacy: An Empirical Test Using Context to Expose Confounding Variables. *Columbia Science and Technology Law Review*, 18(1):176–218.
- Mehrabian, A. and Steffl, C. A. (1995). BASIC TEMPERAMENT COMPONENTS OF LONELINESS, SHYNESS, AND CONFORMITY. *Social Behavior and Personality: an international journal*, 23(3):253–263.
- Mutumukwe, C., Viberg, O., Oberg, L.-M., and Cerratto-Pargman, T. (2022). Students' privacy concerns in learning analytics: Model development. *British Journal of Educational Technology*, 53(4):932–951.
- Myers, D. G. and Kaplan, M. F. (1976). Group-Induced Polarization in Simulated Juries. *Personality and Social Psychology Bulletin*, 2(1):63–66.
- Navajas, J., Niella, T., Garbulsky, G., Bahrami, B., and Sigman, M. (2018). Aggregated knowledge from a small number of debates outperforms the wisdom of large crowds. *Nature Human Behaviour*, 2(2):126–132.
- Niella, T., Stier-Moses, N., and Sigman, M. (2016). Nudging Cooperation in a Crowd Experiment. *PLOS ONE*, 11(1):e0147125.
- Nissenbaum, H. (2004). Privacy as contextual integrity. *Wash. L. Rev.*, 79:119.
- Nissenbaum, H. (2009). Privacy in context: Technology, policy, and the integrity of social life. In *Privacy in Context*. Stanford University Press.
- Pardo, A. and Siemens, G. (2014). Ethical and privacy principles for learning analytics. *British journal of educational technology*, 45(3):438–450.
- Pedersen, E. R., Neighbors, C., Tidwell, J., and Lostutter, T. W. (2011). Do undergraduate student research participants read psychological research consent forms? examining memory effects, condition effects, and individual differences. *Ethics & Behavior*, 21(4):332–350.
- Reimann, P. (2021). Methodological progress in the study of self-regulated learning enables theory advancement. *Learning and Instruction*, 72:101269.
- Romero, C., & Ventura, S. (2020). Educational data mining and learning analytics: An updated survey. *Wiley interdisciplinary reviews: Data mining and knowledge discovery*, 10(3), e1355.
- Sclater, N. (2017). *Learning analytics explained*. Routledge.
- Shaffer, D. W. (2017). *Quantitative ethnography*. Lulu. com.
- Shaffer, D. W., Collier, W., and Ruis, A. R. (2016). A tutorial on epistemic network analysis: Analyzing the structure of connections in cognitive, social, and interaction data. *Journal of Learning Analytics*, 3(3):9–45.

- Siebert-Evenstone, A. L., Irgens, G. A., Collier, W., Swiecki, Z., Ruis, A. R., and Shaffer, D. W. (2017). In search of conversational grain size: Modeling semantic structure using moving stanza windows. *Journal of Learning Analytics*, 4(3):123–139.
- Siemens, G. (2013). Learning analytics: The emergence of a discipline. *American Behavioral Scientist*, 57(10):1380–1400.
- Silber, H., Gerdon, F., Bach, R., Kern, C., Keusch, F., and Kreuter, F. (2022). A preregistered vignette experiment on determinants of health data sharing behavior: Willingness to donate sensor data, medical records, and biomarkers. *Politics and the Life Sciences*, 41(2):161–181.
- Slade, S. and Prinsloo, P. (2013). Learning analytics: Ethical issues and dilemmas. *American Behavioral Scientist*, 57(10):1510–1529.
- Solove, D. J. (2012). Introduction: Privacy self-management and the consent dilemma. *Harv. L. Rev.*, 126:1880.
- Sprecher, S. (2021). Closeness and other affiliative outcomes generated from the fast friends procedure: A comparison with a small-talk task and unstructured self-disclosure and the moderating role of mode of communication. *Journal of Social and Personal Relationships*, 38(5):1452–1471.
- Sun, J., Harris, K., and Vazire, S. (2020). Is well-being associated with the quantity and quality of social interactions? *Journal of personality and social psychology*, 119(6):1478.
- Sun, K., Mhaidli, A. H., Watel, S., Brooks, C. A., and Schaub, F. (2019). It's my data! tensions among stakeholders of a learning analytics dashboard. In *Proceedings of the 2019 CHI conference on human factors in computing systems*, pages 1–14.
- Tan, Y., Ruis, A. R., Marquart, C., Cai, Z., Knowles, M. A., and Shaffer, D. W. (2022). Ordered network analysis. In *International Conference on Quantitative Ethnography*, pages 101–116. Springer.
- van Alten, D. C., Phielix, C., Janssen, J., & Kester, L. (2020). Self-regulated learning support in flipped learning videos enhances learning outcomes. *Computers & Education*, 158, 104000.
- van Dolder, D. and van den Assem, M. J. (2018). The wisdom of the inner crowd in three large natural experiments. *Nature Human Behaviour*, 2(1):21–26.
- Whitelock-Wainwright, A., Gašević, D., Tejeiro, R., Tsai, Y.-S., and Bennett, K. (2019). The student expectations of learning analytics questionnaire. *Journal of Computer Assisted Learning*, 35(5):633–666.
- Whitelock-Wainwright, A., Tsai, Y.-S., Drachsler, H., Scheffel, M., and Gašević, D. (2021). An exploratory latent class analysis of student expectations towards learning analytics services. *The Internet and Higher Education*, 51:100818.
- Wollny, S., Di Mitri, D., Jivet, I., Muñoz-Merino, P., Scheffel, M., Schneider, J., Tsai, Y.-S., Whitelock-Wainwright, A., Gašević, D., and Drachsler, H. (2023). Students' expectations of learning analytics across europe. *Journal of computer assisted learning*, 39(4):1325–1338.

Supplementary Material:

The Power of Conversation: Learners Become More Cautious Sharing Learning Data after a Group Discussion

Louis Longin, Deisy Montenegro Briceno, and Oleksandra Poquet

Supplementary Methods

Detailed vignettes

We used a 2 x 2 x 2 with-subject, in-person experimental design. We varied three factors with two levels each. This included variations in: 1. learning data type (process vs outcome), 2. data recipient (private company vs public government), and 3. data sharing purpose (individual vs collective benefit). Each combination of experimental factors yields eight possible experimental conditions. To compare the effect of group discussion, we further created two versions for each condition yielding 16 vignettes total. One version that would be discussed and the other would only be rated only individually (randomised and counterbalanced).

Below are all the vignettes used for the experiment.

Type (Outcome) x Agent (Company) x Purpose (Individual) - Version 1: LegoCodes, a private company, offers an online game for learning how to code. In the game, students connect Lego-like code blocks to run a computer program. Mistakes and correct tries are recorded to identify how well students can program to provide students with information about their programming level and suggest learning materials matching their knowledge. With the full consent of their users, LegoCodes collects detailed records of student performance to provide personalised recommendations that support individual learners. The data are secure, anonymised, and protected against misuse.

Type (Outcome) x Agent (Company) x Purpose (Individual) - Version 2: MindMentor, a private company, prepares students for university admission. As students practice exam questions online, their grades are continuously recorded to identify how well they are doing compared to the expected level of knowledge. Student performance is used to tailor motivational messages and suggest follow-up materials adapted to student knowledge. With the full consent of their students, MindMentor collects detailed records of student performance to provide personalised recommendations that support individual learners. The data are secure, anonymised, and protected against misuse.

Type (Outcome) x Agent (Government) x Purpose (Individual) -Version 1: Stache State, a public authority, offers virtual cooking lessons for future parents. These programs use virtual reality where future parents learn cooking basics as they assemble different meals. Cooking successes and mistakes are recorded to identify the level of cooking skills and suggest new recipes that adjust in difficulty. With the full consent of their users, Stache State University continuously records student performance to provide personalised recommendations that support individual learners. The data are secure, anonymised, and protected against misuse.

Type (Outcome) x Agent (Government) x Purpose (Individual) Version 2: The town of Pupum, a public authority, produces online games for practicing fire drills and medical emergencies. By engaging in an online game, trainees complete safety challenges that vary in difficulty. The game records the mistakes that trainees make to suggest areas for improvement. With the full consent of their users, the town of Pupum continuously records learner performance to

provide personalised recommendations that support individual learners. The data are secure, anonymised, and protected against misuse.

Type (Outcome) x Agent (Company) x Purpose (Collective) - Version 1: BrushStroke, a private company, teaches adults how to write Chinese characters. In an online app, learners use a virtual brush to write characters. The app records their mistakes and successes to identify individual learning difficulties. With the full consent of the users, BrushStroke collects detailed records of learner performance to share best practices that improve teaching Chinese for everyone in the world. The data are secure, anonymised, and protected against misuse.

Type (Outcome) x Agent (Company) x Purpose (Collective) - Version 2: EmpatoMeter, a private company, offers verbal presentation training to university students. During a training session, the students talk to a virtual audience, prompting different presentation scenarios. Throughout each session an evaluation system continuously assesses the student's communication skills. Student errors and scores within each scenario are recorded and used to evaluate the presentation quality. With the full consent of their students, EmpatoMeter continuously records student performance to share best practices that improve presentation training for everyone in the world. The data are secure, anonymised, and protected against misuse.

Type (Outcome) x Agent (Government) x Purpose (Collective) Version 1: The town of Hogum, a public authority, offers a virtual driving simulation for driving beginners. The simulation presents learners with extreme weather scenarios ranging from probable events like snowstorms to rare events like flooded streets. The successes and failures of the participants are recorded to show learners how well they do in relation to the level expected to succeed on exam. With the full consent of their users, the town of Hogum records detailed learner performance to share best practices that improve driving education for everyone in the world. The data are secure, anonymised, and protected against misuse.

Type (Outcome) x Agent (Government) x Purpose (Collective) Version 2: The city of Ulmberg, a public authority, built a driving simulation to teach beginners international traffic rules. Driving beginners are placed in different countries and must reach a particular target by driving their virtual car. Participant's successes and failures in complying with the local traffic laws are recorded. With the full consent of their users, Ulmberg continuously collects learner performance to share best practices that improve driving education for everyone in the world. The data are secure, anonymised, and protected against misuse.

Type (Process) x Agent (Company) x Purpose (Individual) - Version 1: SpeakUp, a private company, manages an online app for learning German. The app adapts the difficulty of its learning materials to its users based on what they do when learning online. Learning activity of each learner is tracked, including the timing of clicks and sequences of language exercises. With the full consent of their users, SpeakUp continuously records detailed learner activity to provide personalised recommendations that support individual learners. The data are secure, anonymised, and protected against misuse.

Type (Process) x Agent (Company) x Purpose (Individual) - Version 2: MathGuru, a private company, creates online games for learning math. Users can select and play games based on their individual progress. The game platform records which games are selected and for how long they are played. Based on the activity data, students receive recommendations on how to improve their study habits. With the full consent of their users, MathGuru continuously records detailed learner activity to provide personalised recommendations that support individual learners. The data are secure, anonymised, and protected against misuse.

Type (Process) x Agent (Government) x Purpose (Individual) - Version 1: Stanleyville, a public authority, offers virtual-reality classrooms to simulate medical procedures. Within the platform, medical students can experience virtual surgery classes first-hand. The platform records all of students' activity to adapt the virtual experience to their learning profiles. With the full consent of their students, Stanley College continuously records detailed student activity to provide personalised recommendations that support individual learners. The data are secure, anonymised, and protected against misuse.

Type (Process) x Agent (Government) x Purpose (Individual) Version 2: The city of Mappenbruken, a public authority, hosts augmented reality classrooms for engineering students. While attending the face-to-face lecture, students are presented with 3D illustrations of the relevant engineering parts that support learning. Student activity when they interact with the illustration is recorded, including the timing and sequence of clicks, and used to adapt the illustrations to the individual needs. With the full consent of their students, the city of Mappenbruken continuously records detailed student activity to provide personalised recommendations that support individual learners. The data are secure, anonymised, and protected against misuse.

Type (Process) x Agent (Company) x Purpose (Collective) - Version 1: ScienceHub, a private company, offers a virtual platform for re-creating famous chemistry experiments. The platform presents students with an experimental problem, minimal instructions, and chemical materials and tools. The platform records everything students do, including the timing and sequence of their clicks to adapt the task difficulty to the individual students. With the full consent of their students, ScienceHub continuously records detailed student activity to share best practices that improve teaching of science for everyone in the world. The data are secure, anonymised, and protected against misuse.

Type (Process) x Agent (Company) x Purpose (Collective) - Version 2: PhysicsLabs, a private company, runs a virtual platform for teaching atomic particles. Teachers create experimental challenges for students to solve. using a virtual lab. The platform records everything students do including the timing and sequence of their clicks. These activity data are then used to adapt the task difficulty to each individual. With the full consent of their students, PhysicsLabs continuously records detailed student activity to develop best practices for teaching science for everyone in the world. The data are secure, anonymised, and protected against misuse.

Type (Process) x Agent (Government) x Purpose (Collective) Version 1: The town of Grunberg, a public authority, offers online courses about how to run a small business. In the courses, students watch tutorials and complete assignments. The course platform records students' online activity, including the timing of clicks and the sequence of selected tutorials. Student activity data are used to recommend students other videos related to their learning needs. With the full consent of their students, the Town of Grunberg continuously records detailed student activity to share best practices that improve business education for everyone in the world. The data are secure, anonymised, and protected against misuse.

Type (Process) x Agent (Government) x Purpose (Collective) Version 2: The city of Asheville, a public authority, offers self-paced online training on writing job applications. The online training involves written assignments and video tutorials. To measure students' progress, the city of Asheville records students' activity, including the timing of all clicks and the sequence of selected assignments. With the full consent of their users, the city of Asheville continuously records detailed learner behaviour to share best practices that improve job preparation for everyone in the world. The data are secure, anonymised, and protected against misuse.

Supplementary Results

Descriptive Summary Statistics of conditions by phases

Below you can find a summary of average ratings by decision stages (phases) and experimental conditions. The experimental conditions include data type (outcome vs process learning data), data purpose (sharing for a private/individual vs public/collective benefit), and data agent (sharing learning data with a private company vs a public government). For each of the eight possible condition combinations, we calculated mean, median, standard deviation (sd), number of observations (obs), upper and lower bounds for 95% confidence levels. Confidence levels were calculated with the `mean_cl_normal()` function from the `Hmisc` package.

Phase		type	purpose	agent	Mean	Median	SD	Lower	Upper	obs
1	1	outcome	individual	company	5.97	7.00	1.40	5.71	6.22	120
2	1	outcome	individual	government	5.76	6.00	1.65	5.46	6.06	120
3	1	outcome	collective	company	5.62	6.00	1.58	5.34	5.91	120
4	1	outcome	collective	government	5.74	6.00	1.52	5.47	6.02	120
5	1	process	individual	company	6.00	6.00	1.28	5.77	6.23	120
6	1	process	individual	government	5.72	6.00	1.61	5.43	6.01	119
7	1	process	collective	company	5.58	6.00	1.64	5.29	5.88	120
8	1	process	collective	government	5.22	5.00	1.66	4.92	5.53	120
9	2	outcome	individual	company	5.88	6.00	1.27	5.55	6.21	59
10	2	outcome	individual	government	5.53	6.00	1.89	5.03	6.02	59
11	2	outcome	collective	company	5.51	6.00	1.65	5.08	5.94	59
12	2	outcome	collective	government	5.69	7.00	1.73	5.24	6.15	59
13	2	process	individual	company	5.76	6.00	1.48	5.38	6.15	59
14	2	process	individual	government	5.56	6.00	1.75	5.10	6.02	59
15	2	process	collective	company	5.29	6.00	1.68	4.85	5.73	59
16	2	process	collective	government	5.15	5.00	1.84	4.67	5.63	59
17	3	outcome	individual	company	5.85	6.00	1.10	5.57	6.13	60
18	3	outcome	individual	government	5.38	6.00	1.94	4.88	5.88	60
19	3	outcome	collective	company	5.43	6.00	1.83	4.96	5.91	60
20	3	outcome	collective	government	5.17	5.00	1.59	4.76	5.58	60
21	3	process	individual	company	6.17	6.00	0.83	5.95	6.38	60
22	3	process	individual	government	5.83	6.00	1.29	5.50	6.17	60
23	3	process	collective	company	5.17	5.00	1.09	4.88	5.45	60
24	3	process	collective	government	4.25	3.50	1.94	3.75	4.75	60
25	4	outcome	individual	company	5.78	6.00	1.25	5.46	6.11	60
26	4	outcome	individual	government	5.35	6.00	2.10	4.81	5.89	60
27	4	outcome	collective	company	5.18	6.00	1.94	4.68	5.68	60
28	4	outcome	collective	government	5.12	6.00	1.93	4.62	5.62	59
29	4	process	individual	company	5.88	6.00	1.25	5.56	6.21	60
30	4	process	individual	government	5.75	6.00	1.63	5.33	6.17	60
31	4	process	collective	company	5.12	5.00	1.60	4.70	5.54	59
32	4	process	collective	government	4.38	4.50	2.06	3.85	4.92	60

Table 1: Descriptive statistics by phases and experimental conditions.

Based on the data above, we have visualised the effect of decision-making phases on each experimental condition (see Figure 1).

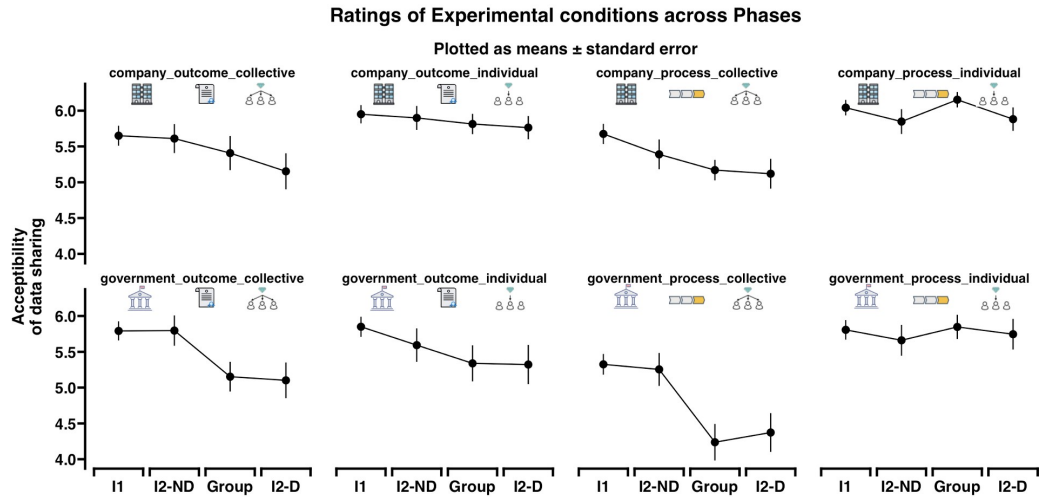


Figure 1: Average acceptability of data sharing ratings by experimental conditions and phases. Plotted are means and standard errors based on descriptive data (see table 1).

We also explored a comparative breakdown of data acceptability decisions per group and decision-making phase (see Figure 2). Ratings ranged from '1 not acceptable at all' to '4 neutral' to '7 completely acceptable'. We can see that the most common choice for data sharing acceptability decisions on average was '7 completely acceptable', which was consistent across groups and phases. Notably, some groups (groups 5, and somewhat group 4) were not critical of data sharing at all; an effect was only amplified during the discussion phase 3 and following individual decision-making phase 4.

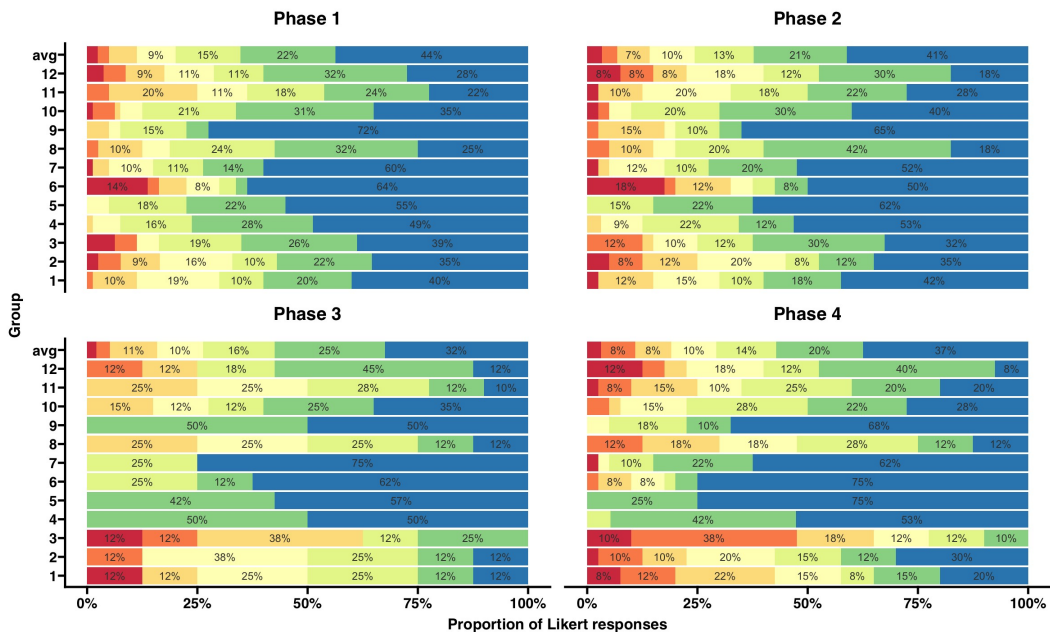


Figure 2: Detailed breakdown of data sharing decisions by phases and group decisions.

Behavioural results

For higher interpretability and ease of access, we choose to report linear/Gaussian regression model results over cumulative link model results. While Gaussian models are in principle unsuited for Likert regressions, they often make no difference in practice, at least when the responses are not concentrated near either end of the Likert scale.

We observe that the patterns of significance are the same as the original linear/Gaussian model reported above (see 3). As logit regression coefficients are not as straightforwardly interpretable as the linear coefficients in the main models, and it makes little difference to our general conclusions, we choose to report the more parsimonious model in the main text. Nonetheless, for the supplementary material, we included a visual comparison of main effect estimates of both linear and cumulative link models.

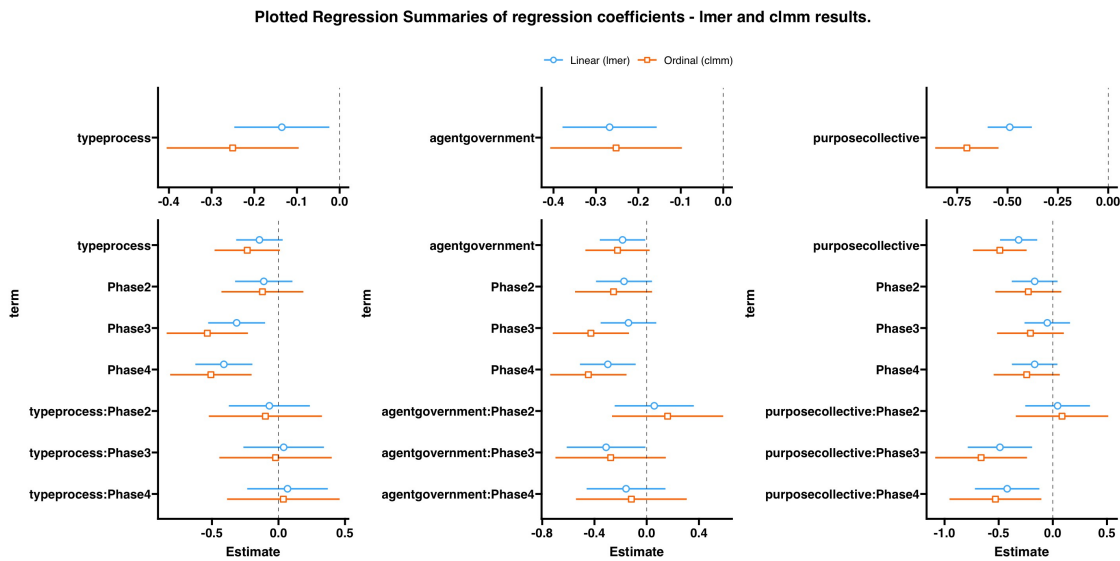


Figure 3: Supplement to Main Figure 2. Comparison of mixed linear (lmer) with cumulative linked models (clmm).

Additional robustness check

As an additional robustness for a possible dependence of model regressors, we fitted two models: a linear mixed model to predict data acceptability ratings based on a four-way interaction (ratings ~Phase x agent x type x purpose), and a linear mixed model to predict data acceptability ratings based on a two-way interaction but a larger scope (ratings ~Phase x (agent + type + purpose)). The first model yielded no notable effects. The second model confirmed the effects reported in the main text.

For the second model, we fitted a linear mixed model (estimated using REML and nlptwrap optimizer) to predict Answer with type, purpose, agent and Phase (formula: Answer ~ (type + purpose + agent) * Phase). The model included p_g as random effects (formula: list(~1 | p_g:Group, ~1 | Group)). The model's total explanatory power is substantial (conditional R2 = 0.34) and the part related to the fixed effects alone (marginal R2) is of 0.05. The model's intercept, corresponding to type = outcome, purpose = individual, agent = company and Phase = 1, is at 6.02 (95% CI [5.62, 6.42], t(2370) = 29.61, p < .001). Within this model:

- The effect of type [process] is statistically non-significant and negative (beta = -0.14, 95% CI [-0.31, 0.03], t(2370) = -1.64, p = 0.100; Std. beta = -0.09, 95% CI [-0.19, 0.02])
- The effect of purpose [collective] is statistically significant and negative (beta = -0.32, 95% CI [-0.49, -0.14], t(2370) = -3.62, p < .001; Std. beta = -0.19, 95% CI [-0.30, -0.09])
- The effect of agent [government] is statistically significant and negative (beta = -0.18, 95% CI [-0.36, -0.01], t(2370) = -2.12, p = 0.034; Std. beta = -0.11, 95% CI [-0.22, -8.55e-03])
- The effect of Phase [2] is statistically non-significant and negative (beta = -0.16, 95% CI [-0.46, 0.14], t(2370) = -1.07, p = 0.286; Std. beta = -0.10, 95% CI [-0.28, 0.08])
- The effect of Phase [3] is statistically non-significant and positive (beta = 0.09, 95% CI [-0.21, 0.38], t(2370) = 0.57, p = 0.571; Std. beta = 0.05, 95% CI [-0.13, 0.23])
- The effect of Phase [4] is statistically non-significant and negative (beta = -0.12, 95% CI [-0.42, 0.17], t(2370) = -0.82, p = 0.415; Std. beta = -0.07, 95% CI [-0.25, 0.10])
- The effect of type [process] × Phase [2] is statistically non-significant and negative (beta = -0.07, 95% CI [-0.37, 0.23], t(2370) = -0.45, p = 0.650; Std. beta = -0.04, 95% CI [-0.22, 0.14])
- The effect of type [process] × Phase [3] is statistically non-significant and positive (beta = 0.04, 95% CI [-0.26, 0.33], t(2370) = 0.26, p = 0.796; Std. beta = 0.02, 95% CI [-0.16, 0.20])
- The effect of type [process] × Phase [4] is statistically non-significant and positive (beta = 0.07, 95% CI [-0.23, 0.37], t(2370) = 0.46, p = 0.647; Std. beta = 0.04, 95% CI [-0.14, 0.22])
- The effect of purpose [collective] × Phase [2] is statistically non-significant and positive (beta = 0.04, 95% CI [-0.25, 0.34], t(2370) = 0.29, p = 0.771; Std. beta = 0.03, 95% CI [-0.15, 0.21])
- The effect of purpose [collective] × Phase [3] is statistically significant and negative (beta = -0.49, 95% CI [-0.78, -0.19], t(2370) = -3.24, p = 0.001; Std. beta = -0.30, 95% CI [-0.48, -0.12])
- The effect of purpose [collective] × Phase [4] is statistically significant and negative (beta = -0.42, 95% CI [-0.72, -0.13], t(2370) = -2.80, p = 0.005; Std. beta = -0.26, 95% CI [-0.44, -0.08])
- The effect of agent [government] × Phase [2] is statistically non-significant and positive (beta = 0.06, 95% CI [-0.24, 0.35], t(2370) = 0.38, p = 0.704; Std. beta = 0.04, 95% CI [-0.15, 0.22])
- The effect of agent [government] × Phase [3] is statistically significant and negative (beta = -0.31, 95% CI [-0.61, -0.02], t(2370) = -2.06, p = 0.039; Std. beta = -0.19, 95% CI [-0.37, -9.49e-03])
- The effect of agent [government] × Phase [4] is statistically non-significant and negative (beta = -0.16, 95% CI [-0.45, 0.14], t(2370) = -1.05, p = 0.295; Std. beta = -0.10, 95% CI [-0.28, 0.08])

Standardised parameters were obtained by fitting the model on a standardised version of the dataset. 95% Confidence Intervals (CIs) and p-values were computed using a Wald t-distribution approximation.

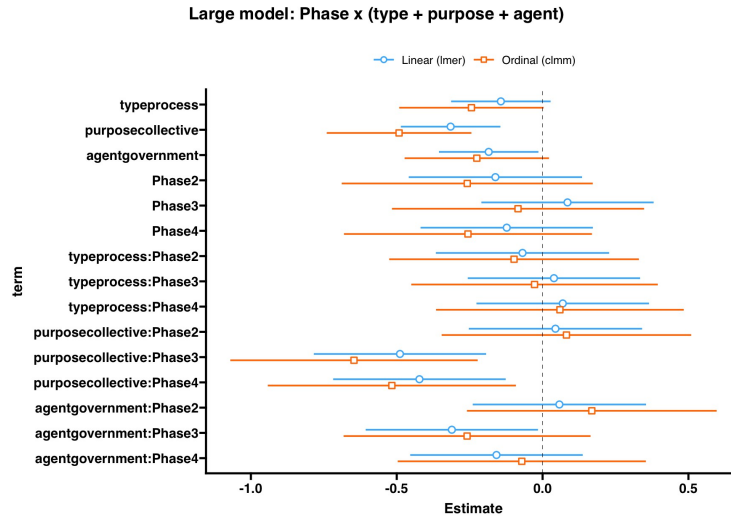


Figure 4: Additional robustness check modelling decision phases with all three experimental conditions as an interaction effect. Comparison of lmer and clmm model results using the plot_summs function from the jtools package Long (2022).

Demographics

Building on the reported demographic effects in the main text, here we expand on the full suite of demographic effects. In our final demographics questionnaire, we measured age, gender, highest completed educational degree, general trust in public authorities and private companies, general privacy concerns, sensitivity of learning data, a validated 11-item social conformity questionnaire, 5-items to capture social experience during the group discussion, and an open-text box for feedback.

Gender

We fitted a linear mixed model (estimated using REML and nloptwrap optimiser) to predict Answer with Gender (formula: Answer ~Gender). The model included p_g as random effects (formula: list(~1 | p_g:Group, ~1 | Group)). The model's total explanatory power is substantial (conditional R2 = 0.30) and the part related to the fixed effects alone (marginal R2) is of 0.01. The model's intercept, corresponding to Gender = female, is at 5.71 (95% CI [5.30, 6.12], t(2384) = 27.44, p < .001). Within this model:

- The effect of Gender [male] is statistically non-significant and negative (beta = -0.34, 95% CI [-0.71, 0.03], t(2384) = -1.80, p = 0.072; Std. beta = -0.21, 95% CI [-0.43, 0.02])

Standardised parameters were obtained by fitting the model on a standardised version of the dataset. 95% Confidence Intervals (CIs) and p-values were computed using a Wald t-distribution approximation.

A similar cumulative link mixed model (clmm) from the ordinal package was used. Results were comparable with the liner model (see Figure 5).

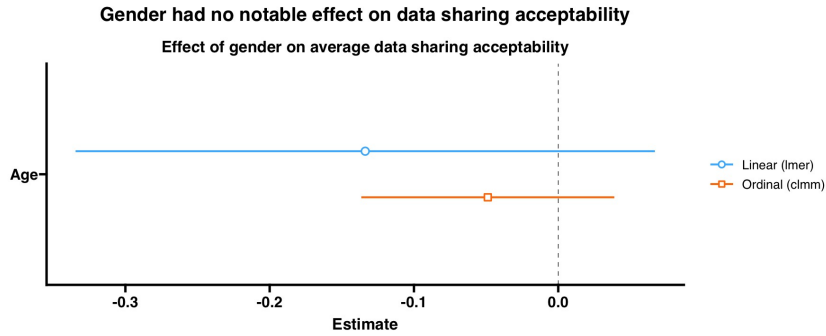


Figure 5: No average *gender* effect on data sharing acceptability ratings. Comparison of lmer and clmm model results using the plot_summs function from the jtools package Long (2022).

Education

We fitted a linear mixed model (estimated using REML and nlptwrap optimiser) to predict Answer with Education (formula: Answer ~Education). The model included p_g as random effects (formula: list(~1 | p_g:Group, ~1 | Group)). The model's total explanatory power is substantial (conditional R2 = 0.30) and the part related to the fixed effects alone (marginal R2) is of 0.03. The model's intercept, corresponding to Education = no_degree, is at 4.57 (95% CI [3.09, 6.05], t(2382) = 6.06, p < .001). Within this model:

- The effect of Education [a_levels] is statistically non-significant and positive (beta = 1.18, 95% CI [-0.28, 2.64], t(2382) = 1.59, p = 0.113; Std. beta = 0.72, 95% CI [-0.17, 1.60])
- The effect of Education [bachelor] is statistically non-significant and positive (beta = 1.05, 95% CI [-0.46, 2.55], t(2382) = 1.36, p = 0.173; Std. beta = 0.64, 95% CI [-0.28, 1.55])
- The effect of Education [master] is statistically non-significant and positive (beta = 0.43, 95% CI [-1.07, 1.92], t(2382) = 0.56, p = 0.574; Std. beta = 0.26, 95% CI [-0.65, 1.17])

Standardised parameters were obtained by fitting the model on a standardised version of the dataset. 95% Confidence Intervals (CIs) and p-values were computed using a Wald t-distribution approximation.

A similar cumulative link mixed model (clmm) from the ordinal package was used. Results were comparable with the liner model (see 6).

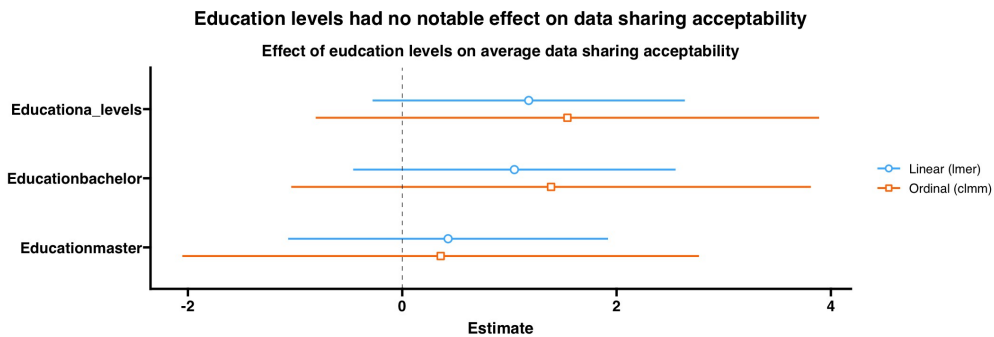


Figure 6: No average *education* effect on data sharing acceptability ratings. Comparison of lmer and clmm model results.

Trust Data Agents

We fitted a linear mixed model (estimated using REML and nloptwrap optimiser) to predict Answer with Trust_Private and Trust_Public (formula: Answer ~Trust_Private + Trust_Public). The model included p_g as random effects (formula: list(~1 | p_g Group, ~1 | Group)). The model's total explanatory power is substantial (conditional R2 = 0.38) and the part related to the fixed effects alone (marginal R2) is of 0.04. The model's intercept, corresponding to Trust_Private = 0 and Trust_Public = 0, is at 4.97 (95% CI [4.26, 5.68], t(1189) = 13.79, p < .001). Within this model:

- The effect of Trust Private is statistically significant and positive (beta = 0.14, 95% CI [0.01, 0.28], t(1189) = 2.12, p = 0.034; Std. beta = 0.19, 95% CI [0.01, 0.36])
- The effect of Trust Public is statistically non-significant and positive (beta = 5.89e-03, 95% CI [-0.10, 0.11], t(1189) = 0.11, p = 0.915; Std. beta = 9.29e-03, 95% CI [-0.16, 0.18])

Standardised parameters were obtained by fitting the model on a standardised version of the dataset. 95% Confidence Intervals (CIs) and p-values were computed using a Wald t-distribution approximation.

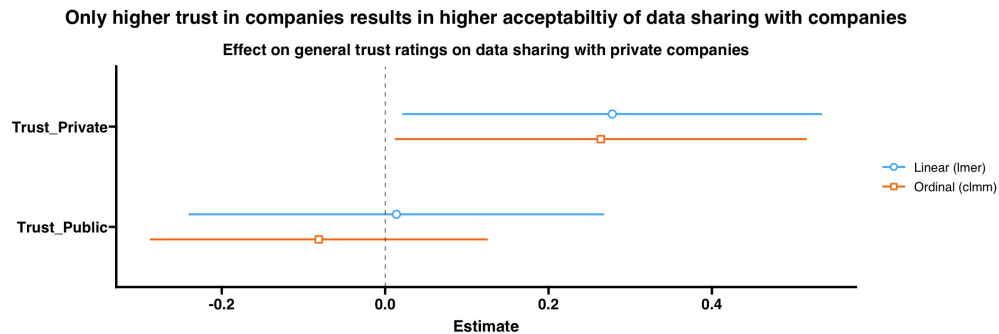


Figure 7: Effect of *general trust* ratings on data sharing acceptability rating with private companies. Comparison of lmer and clmm model results using the plot_summs function from the jtools package Long (2022).

We fitted a linear mixed model (estimated using REML and nloptwrap optimiser) to predict Answer with Trust_Private and Trust_Public (formula: Answer ~Trust_Private + Trust_Public). The model included p_g as random effects (formula: list(~1 | p_g:Group, ~1 | Group)). The model's total explanatory power is substantial (conditional R2 = 0.34) and the part related to the fixed effects alone (marginal R2) is of 0.02. The model's intercept, corresponding to Trust_Private = 0 and Trust_Public = 0, is at 4.88 (95% CI [4.08, 5.68], t(1188) = 12.01, p < .001). Within this model:

- The effect of Trust Private is statistically non-significant and negative (beta = -0.04, 95% CI [-0.18, 0.09], t(1188) = -0.63, p = 0.531; Std. beta = -0.05, 95% CI [-0.20, 0.10])
- The effect of Trust Public is statistically significant and positive (beta = 0.12, 95% CI [5.75e-03, 0.23], t(1188) = 2.06, p = 0.040; Std. beta = 0.16, 95% CI [7.55e-03, 0.31])

Standardised parameters were obtained by fitting the model on a standardised version of the dataset. 95% Confidence Intervals (CIs) and p-values were computed using a Wald t-distribution approximation.




Figure 8: Effect of *general trust* ratings on data sharing acceptability rating with public companies. Comparison of lmer and clmm model results.

As an additional robustness check, we compared the relationship between general trust ratings in governments and companies. As expected, general trust ratings with governments correlate with general trust ratings with private companies (see Figure 9).

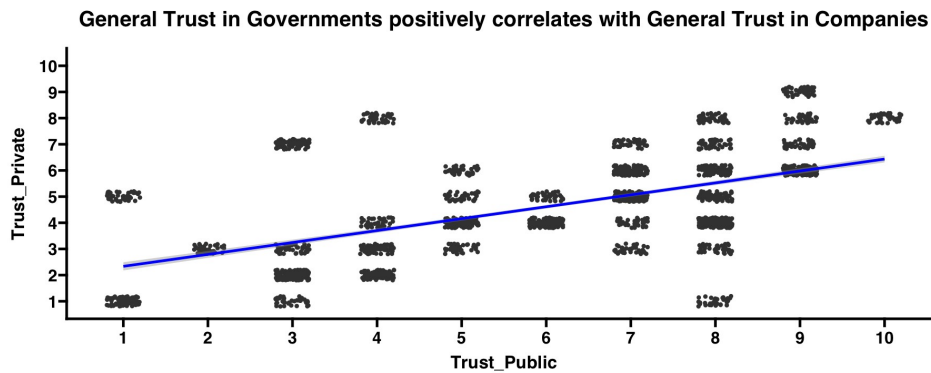


Figure 9: The higher general trust in companies the higher the general trust in governments. Comparison of lmer and clmm model results using the plot_summs function from the jtools package Long (2022).

General Privacy concerns

We fitted a linear mixed model (estimated using REML and nloptwrap optimiser) to predict Answer with Privacy (formula: Answer ~ Privacy). The model included p_g as random effects (formula: list(~1 | p_g:Group, ~1 | Group)). The model's total explanatory power is substantial (conditional R² = 0.29) and the part related to the fixed effects alone (marginal R²) is of 0.03. The model's intercept, corresponding to Privacy = 0, is at 6.56 (95% CI [5.80, 7.32], t(2384) = 16.87, p < .001). Within this model:

- The effect of Privacy is statistically significant and negative (beta = -0.39, 95% CI [-0.66, -0.13], t(2384) = -2.95, p = 0.003; Std. beta = -0.17, 95% CI [-0.28, -0.06])

Standardised parameters were obtained by fitting the model on a standardised version of the dataset. 95% Confidence Intervals (CIs) and p-values were computed using a Wald t-distribution approximation.

The higher the general privacy concern the lower the average data sharing acceptability




Figure 10: Effect of general concerns of data privacy on average data sharing acceptability. Left: Comparison of lmer and clmm model results using the plot_summs function from the jtools package Long (2022). Right: Individual acceptability ratings by general concern of data privacy.

General student behaviour concerns

We fitted a linear mixed model (estimated using REML and nloptwrap optimiser) to predict Answer with Student_behaviour (formula: Answer ~Student_behaviour). The model included p_g as random effects (formula: list(~1 | p_g:Group, ~1 | Group)). The model’s total explanatory power is substantial (conditional R2 = 0.30) and the part related to the fixed effects alone (marginal R2) is of 2.90e-05. The model’s intercept, corresponding to Student_behaviour = 0, is at 5.57 (95% CI [4.71, 6.43], t(2384) = 12.68, p < .001). Within this model:

- The effect of Student behaviour is statistically non-significant and negative (beta = -0.01, 95% CI [-0.27, 0.25], t(2384) = -0.08, p = 0.933; Std. beta = -5.42e-03, 95% CI [-0.13, 0.12])

Standardised parameters were obtained by fitting the model on a standardised version of the dataset. 95% Confidence Intervals (CIs) and p-values were computed using a Wald t-distribution approximation.

Concerns around student behavioural data have no effect on average data sharing acceptability




Figure 11: Effect of general concerns of processing student data on average data sharing acceptability. Left: Comparison of lmer and clmm model results using the plot_summs function from the jtools package Long (2022). Right: Individual acceptability ratings by general concern of processing student data.

Social conformity

We used a validated basic temperament scale (see Mehrabian and Stefl (1995)) to test possible effects of social conformity on the influence of group discussion on individual ratings.

We fitted a linear mixed model (estimated using REML and nloptwrap optimiser) to predict Answer with social_conformity (formula: Answer ~social_conformity). The model included p_g as random effects (formula: list(~1 | p_g:Group, ~1 | Group)). The model's total explanatory power is substantial (conditional R2 = 0.30) and the part related to the fixed effects alone (marginal R2) is of 2.43e-03. The model's intercept, corresponding to social_conformity = 0, is at 5.93 (95% CI [4.88, 6.98], t(2344) = 11.04, p < .001). Within this model:

- The effect of social conformity is statistically non-significant and negative (beta = -0.01, 95% CI [-0.04, 0.02], t(2344) = -0.78, p = 0.434; Std. beta = -0.05, 95% CI [-0.18, 0.08])

Standardised parameters were obtained by fitting the model on a standardised version of the dataset. 95% Confidence Intervals (CIs) and p-values were computed using a Wald t-distribution approximation.

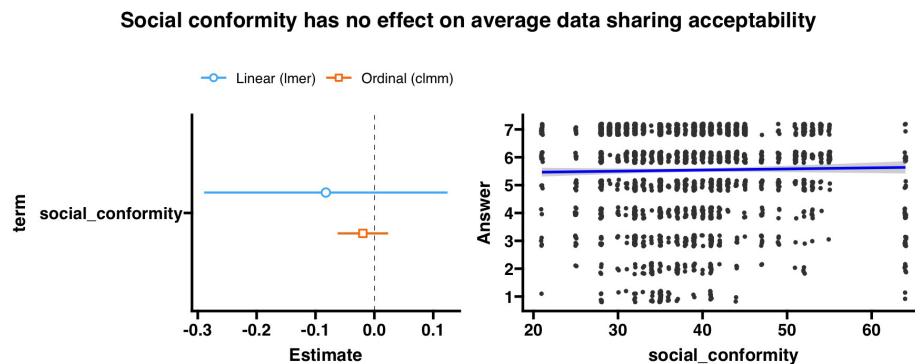


Figure 12: Social conformity has no effect on average data sharing acceptability. Left: Comparison of lmer and clmm model results using the plot_summs function from the jtools package Long (2022). Right: Individual acceptability ratings by social conformity scores.

Social connectedness

We fitted a linear mixed model (estimated using REML and nloptwrap optimiser) to predict Answer with connectedness (formula: Answer ~connectedness). The model included p_g as random effects (formula: list(~1 | p_g:Group, ~1 | Group)). The model's total explanatory power is substantial (conditional R2 = 0.32) and the part related to the fixed effects alone (marginal R2) is of 3.33e-04. The model's intercept, corresponding to connectedness = 0, is at 5.71 (95% CI [4.26, 7.16], t(2184) = 7.72, p < .001). Within this model:

- The effect of connectedness is statistically non-significant and negative (beta = -7.78e-03, 95% CI [-0.06, 0.05], t(2184) = -0.27, p = 0.785; Std. beta = -0.02, 95% CI [-0.15, 0.11])

Standardised parameters were obtained by fitting the model on a standardised version of the dataset. 95% Confidence Intervals (CIs) and p-values were computed using a Wald t-distribution approximation.

Social connectedness has no effect on average data sharing acceptability




Figure 13: Social connectedness has no effect on average data sharing acceptability. Left: Comparison of lmer and clmm model results using the plot_summs function from the jtools package Long (2022). Right: Individual acceptability ratings by social connectedness scores.

Transcript Analysis

Content analysis of transcripts

Group discussions were recorded and transcribed. A coding scheme was developed to capture contextual integrity theory characteristics taken into account when participants discussed the acceptability of data sharing. The coding scheme comprised seven distinct codes. Five codes were derived from the contextual integrity framework: actor_subject, actor_receiver, attribute_data, transmission_norm, and beneficiary. Two additional codes, utility and risk were incorporated to capture that often participants considered the context presented to them by discussing it against risk and utility. Transcripts of all discussions representing twelve groups were coded by two researchers. First, researchers coded the same two transcripts using a preliminary coding scheme. Upon discussing the results, resolving the discrepancies, and adapting the coding scheme, both researchers coded two more group transcripts. Results were again compared and discussed. The remainder eight transcripts were coded separately by the two researchers, as sufficient agreement was reached. Another code comparison discussion was conducted during this process.

Epistemic network analysis of transcripts: Method

To analyse the differences between the contextual integrity themes emergent in group discussions, we applied ordered network analysis (ONA) Tan et al. (2022) to the coded transcript data using the epistemic network analysis web tool (version 1.7.0) Marquart, Hinojosa, Swiecki, Eagan, and Shaffer (2018). The method gained prominence in educational settings due to its capacity to bridge qualitative and quantitative data analysis Shaffer (2017) and is well-described Shaffer, Collier, and Ruis (2016); Siebert-Evenstone et al. (2017); Tan et al. (2022). Briefly, epistemic network analysis Shaffer et al. (2016) combines content analysis with network analysis. Codes derived from content analysis are represented as network nodes and code co-occurrence within a selected segment represents a network tie. An observation of interest, e.g. speech by individual, a group, or similar, can be represented as a graph. Since the number of codes remains consistent across all observations, the graphs can be built for as many observations of interest as relevant, to be further compared across each other after each graph structure has been transformed. To this end, ENA web tool projects each graph representative of an observation of interest into the same space. Using single value decomposition, ENA estimates a position for each graph in relation to each other. This enables statistical and visual comparisons across multiple observations using graph positions within the space. Qualitative comparisons of the structure of observations averaged around a chosen attribute are possible, in cases of reasonable goodness of fit statistics.

Category	Description	Codes	Examples
Actors	Individuals or entities involved in the information-sharing process.	actor_subject actor_receiver	"It's not acceptable because it forces students to accept it if they want to join the course." "I'd rather have the private company ... For me, it makes [the acceptability] a five or six, something between..."
Attributes	Type of data being shared (e.g., personal information, behavioural data, logs, timestamps).	attribute_data	"This time, the exact timing of clicks and so on is not tracked so..."
Transmission Principle	Norms and considerations around the flow of information (e.g., data secured, anonymised, full consent).	transmission_norm	"But actually they're only providing the information anonymised... So I think there's no problem with it."
Purpose	Mentions who benefits from the data shared (collective vs. personalised). Concerns about risks of sharing data or risk data being used to control people. Mentions the purpose of the tool/app in the scenario.	beneficiary risk utility	"I'm also not sure about sharing to everyone in the world." "But why do they even collect it in the first place? Like why does the town control?" "I don't see how the data collected from one individual person could help the others in this case."

Table 2: Coding framework based on Nissenbaum's contextual integrity theory

Data Segmentation: Utterance, Dyadic, and Discussion Levels

ENA enables to segment text data in various ways for further aggregation of code co-occurrence as network ties. The capture of code co-occurrence is possible at the level of each speaker's utterance, in a moving window of various sizes, as well as across the totality of text. To gain insight into the sensitivity of network structures representative each group's discussion of a vignette, we constructed ordered epistemic networks for all conditions, at the level of individual utterance representative of what each speaker says at a time; at the level of a dyad (a moving window of 2) representative of what a dyad discusses at a time, and at the level of the entire discussion. Utility and risk codes were dominant in numbers across all networks and connected to all other codes. They were, therefore, removed from the analysis, following methodological guidelines for tackling dominant codes Ferreira Mello and Gašević (2019).

Figure 14 depicts ordered epistemic networks for each of these segments. For both Receiver Models and Data Type Models, the analyses of code co-occurrence at the level of utterance offer more heterogeneity, and the structures are comparable from the level of a dyad and up to the level of a whole conversation. This differs for the Purpose Model. Regardless of the level of analysis, we observe stability in the prominent patterns around the ties representing higher co-occurring codes per condition and prominence of nodes. Significant differences between networks on the x-axis are also observed across levels of analysis. Based on this analysis, main paper reports on ordered whole-level discussion networks constructed at the level of group-vignette, where the abundance of data observations enables cross-groups comparisons. Main results presented in These analyses are supplementary to Figure 3A in the main manuscript.

Receiver Models (top row, Figure 14)

	Infinite Stanza		Moving Window 2		Whole Conversation	
	x-axis	y-axis	x-axis	y-axis	x-axis	y-axis
Variance	8%	15.8%	7.3%	14.2%	8%	16%
Pearson corr (gof)	0.81	0.96	0.75	0.95	0.82	0.96
Spearman corr (gof)	0.82	0.96	0.74	0.95	0.83	0.96
Mann-Whitney test	p=0.00	ns	p=0.00	ns	p=0	ns
Cohen's d	r=0.48	-	r=0.45	-	r=0.48	-

Table 3: Comparison of Variance, Goodness of fit, and Statistical Comparisons across Different Models

Infinite Stanza Model. Along the X axis (MR1), a Mann-Whitney test showed that Government (Mdn=-0.10, N=48) was statistically significantly different at the alpha=0.05 level from Company (Mdn=0.11, N=48 U=604.00, p=0.00, r=0.48). Along the Y axis (SVD2), a Mann-Whitney test showed that Government (Mdn=-0.03, N=48) was not statistically significantly different at the alpha=0.05 level from Company (Mdn=0.00, N=48 U=1136.00, p=0.91, r=0.01).

Moving Stanza 2 Model. Along the X axis (MR1), a Mann-Whitney test showed that Government (Mdn=-0.08, N=48) was statistically significantly different at the alpha=0.05 level from Company (Mdn=0.07, N=48 U=629.00, p=0.00, r=0.45). Along the Y axis (SVD2), a Mann-Whitney test showed that Government (Mdn=-0.09, N=48) was not statistically significantly different at the alpha=0.05 level from Company (Mdn=-0.03, N=48 U=1075.00, p=0.57, r=0.07).

Whole Discussion Model. Along the X axis (MR1), a Mann-Whitney test showed that Government (Mdn=-0.23, N=8) was statistically significantly different at the alpha=0.05 level from Company (Mdn=0.19, N=8 U=8.00, p=0.01, r=0.75). Along the Y axis (SVD2), a Mann-Whitney test showed that Government (Mdn=0.00, N=8) was not statistically significantly different at the alpha=0.05 level from Company (Mdn=0.09, N=8 U=36.00, p=0.72, r=-0.12).

Purpose Models (middle row, Figure 14)

	Infinite Stanza		Moving Window 2		Whole Conversation	
	x-axis	y-axis	x-axis	y-axis	x-axis	y-axis
Variance	8%	13.5%	8%	13.4%	8%	13.6%
Pearson corr (gof)	0.78	0.94	0.75	0.94	0.77	0.94
Spearman corr (gof)	0.78	0.95	0.72	0.95	0.77	0.94
Mann-Whitney test	p=0.00	ns	p=0.00	ns	p=0.00	ns
Cohen's d	r=0.46	-	r=0.43	-	r=0.46	-

Table 4: Comparison of Variance, Goodness of fit, and Statistical Comparisons across Different Models

Infinite Stanza Model. Along the X axis (MR1), a Mann-Whitney test showed that Individual (Mdn=-0.11, N=48) was statistically significantly different at the alpha=0.05 level from Collective (Mdn=0.07, N=48 U=620.00, p=0.00, r=0.46). Along the Y axis (SVD2), a Mann-Whitney test showed that Individual (Mdn=0.00, N=48) was not statistically significantly different at the alpha=0.05 level from Collective (Mdn=0.00, N=48 U=1172.00, p=0.89, r=-0.02).

Moving Stanza 2. Along the X axis (MR1), a Mann-Whitney test showed that Individual (Mdn=-0.03, N=48) was statistically significantly different at the alpha=0.05 level from Collective (Mdn=0.07, N=48 U=652.00, p=0.00, r=0.43). Along the Y axis (SVD2), a Mann-Whitney test showed that Individual (Mdn=0.00, N=48) was not statistically significantly different at the alpha=0.05 level from Collective (Mdn=0.00, N=48 U=1155.00, p=0.99, r=0.00).

Whole Discussion Model. Along the X axis (MR1), a Mann-Whitney test showed that Individual (Mdn=-0.11, N=48) was statistically significantly different at the alpha=0.05 level from Collective (Mdn=0.08, N=48 U=618.00, p=0.00, r=0.46). Along the Y axis (SVD2), a Mann-Whitney test showed that Individual (Mdn=0.00, N=48) was not statistically significantly different at the alpha=0.05 level from Collective (Mdn=-0.02, N=48 U=1175.00, p=0.87, r=-0.02).

Data Type Models (bottom row, Figure 14)

	Infinite Stanza		Moving Window 2		Whole Conversation	
	x-axis	y-axis	x-axis	y-axis	x-axis	y-axis
Variance	5.5%	15.3%	6%	14%	6.4%	14.4%
Pearson corr (gof)	0.66	0.96	0.68	0.95	0.68	0.95
Spearman corr (gof)	0.65	0.96	0.65	0.95	0.65	0.95
Mann-Whitney test	p=0.00	ns	p=0.00	ns	p=0.00	ns
Cohen's d	r=0.37	-	r=0.39	-	r=0.39	-

Table 5: Comparison of Variance, Goodness of fit, and Statistical Comparisons across Different Models

Infinite Stanza Model. Along the X axis (MR1), a Mann-Whitney test showed that Outcome (Mdn=-0.03, N=48) was statistically significantly different at the alpha=0.05 level from Process (Mdn=0.09, N=48 U=721.00, p=0.00, r=0.37). Along the Y axis (SVD2), a Mann-Whitney test showed that Outcome (Mdn=-0.01, N=48) was not statistically significantly different at the alpha=0.05 level from Process (Mdn=-0.01, N=48 U=1152.50, p=1.00, r=0.00).

Moving Stanza 2 Model. Along the X axis (MR1), a Mann-Whitney test showed that Outcome (Mdn=-0.05, N=48) was statistically significantly different at the alpha=0.05 level from Process (Mdn=0.06, N=48 U=706.00, p=0.00, r=0.39). Along the Y axis (SVD2), a Mann-Whitney test showed that Outcome (Mdn=-0.01, N=48) was not statistically significantly different at the alpha=0.05 level from Process (Mdn=-0.11, N=48 U=1169.50, p=0.90, r=-0.02).

Whole Discussion Model. Along the X axis (MR1), a Mann-Whitney test showed that Outcome (Mdn=-0.05, N=48) was statistically significantly different at the alpha=0.05 level from Process (Mdn=0.06, N=48 U=706.00, p=0.00, r=0.39). Along the Y axis (SVD2), a Mann-Whitney test showed that Outcome (Mdn=-0.01, N=48) was not statistically significantly different at the alpha=0.05 level from Process (Mdn=-0.11, N=48 U=1169.50, p=0.90, r=-0.02).

Epistemic network analysis of transcripts: Statistical Comparisons between "Group-Vignette" Units of Analysis

	Receiver Model		Purpose Model		Data Type Model	
	x-axis	y-axis	x-axis	y-axis	x-axis	y-axis
Variance	7%	14%	8%	13.6%	6.4%	14.4%
Pearson corr (gof)	0.75	0.95	0.77	0.94	0.68	0.95
Spearman corr (gof)	0.74	0.96	0.77	0.94	0.65	0.95
Mann-Whitney test	p=0	ns	p=0.00	ns	p=0.00	ns
Cohen's d	r=0.45	-	r=0.46	-	r=0.39	-

Table 6: Comparison of Variance, Goodness of fit, and Statistical Comparisons across Different Models

These analyses are supplementary to Figure 3A in the main manuscript. Main analyses were conducted around ordered epistemic networks analysed at the unit of 'group and vignette', at the level of whole discussion, with with

dominant codes removed. Across all three factors - receiver, purpose, and data type, x-axis in these networks was statistically different for networks representing different conditions within a factor.

Receiver Model. Settings: Ordered Network, Group-Vignette Unit of Analysis, Whole Discussion. Along the X axis (MR1), a Mann-Whitney test showed that Government (Mdn=-0.08, N=48) was statistically significantly different at the alpha=0.05 level from Company (Mdn=0.07, N=48 U=629.00, p=0.00, r=0.45). Along the Y axis (SVD2), a Mann-Whitney test showed that Government (Mdn=-0.09, N=48) was not statistically significantly different at the alpha=0.05 level from Company (Mdn=-0.03, N=48 U=1075.00, p=0.57, r=0.07).

Purpose Model. Settings: Ordered Network, Group-Vignette Unit of Analysis, Whole Discussion. Along the X axis (MR1), a Mann-Whitney test showed that Individual (Mdn=-0.11, N=48) was statistically significantly different at the alpha=0.05 level from Collective (Mdn=0.08, N=48 U=618.00, p=0.00, r=0.46). Along the Y axis (SVD2), a Mann-Whitney test showed that Individual (Mdn=0.00, N=48) was not statistically significantly different at the alpha=0.05 level from Collective (Mdn=-0.02, N=48 U=1175.00, p=0.87, r=-0.02).

Data Type Model. Settings: Ordered Network, Group-Vignette Unit of Analysis, Whole Discussion. Along the X axis (MR1), a Mann-Whitney test showed that Outcome (Mdn=-0.05, N=48) was statistically significantly different at the alpha=0.05 level from Process (Mdn=0.06, N=48 U=706.00, p=0.00, r=0.39). Along the Y axis (SVD2), a Mann-Whitney test showed that Outcome (Mdn=0.01, N=48) was not statistically significantly different at the alpha=0.05 level from Process (Mdn=-0.11, N=48 U=1169.50, p=0.90, r=-0.02).

Epistemic network analysis of transcripts: Statistical Comparisons between "Vignettes Across Groups" Units of Analysis

	Receiver Model		Purpose Model		Data Type Model	
	x-axis	y-axis	x-axis	y-axis	x-axis	y-axis
Variance	29	39	35	29	14	38
Pearson corr (gof)	1	1	0.99	0.99	0.97	1
Spearman corr (gof)	1	1	0.99	0.99	0.97	1
Mann-Whitney test	p=0.02	ns	ns	ns	p=0.01	ns
Cohen's d	r=0.69	-	-	-	r=0.72	-

Table 7: Comparison of Variance, Goodness of fit, and Statistical Comparisons across Different Models

These analyses are supplementary to Figure 3B in the main manuscript. Main analyses were conducted around standard epistemic networks analysed at the unit of 'vignette across groups', at the level of individual utterance (infinite stanza), with with dominant codes and central codes removed. The choice of utterance was made to increase variability of the structures, given the fewer observations points in vignette-to-vignette comparisons. These structures represent what participants said individually rather than collectively. Networks in Receiver model and Data type model differ significantly in structures along the x-axis.

Receiver Model. Settings: Standard Network, Vignette Unit of Analysis, Infinite Stanza. Along the X axis (MR1), a Mann-Whitney test showed that Company (Mdn=-0.17, N=8) was statistically significantly different at the alpha=0.05 level from Government (Mdn=0.12, N=8 U=54.00, p=0.02, r=-0.69). Along the Y axis (SVD2), a Mann-Whitney test showed that Company (Mdn=0.02, N=8) was not statistically significantly different at the alpha=0.05 level from Government (Mdn=-0.11, N=8 U=34.00, p=0.88, r=-0.06). *Purpose Model. Settings: Standard Network, Vignette Unit of Analysis, Infinite Stanza.* Along the X axis (MR1), a Mann-Whitney test showed that Individual (Mdn=-0.15, N=8) was not statistically significantly different at the alpha=0.05 level from Collective (Mdn=0.08, N=8 U=15.00, p=0.08, r=0.53).

Along the Y axis (SVD2), a Mann-Whitney test showed that Individual (Mdn=0.02, N=8) was not statistically significantly different at the alpha=0.05 level from Collective (Mdn=0.07, N=8 U=34.00, p=0.88, r=-0.06).

Data Type Model. Settings: Standard Network, Vignette Unit of Analysis, Infinite Stanza. Along the X axis (MR1), a Mann-Whitney test showed that Outcome (Mdn=0.13, N=8) was statistically significantly different at the alpha=0.05 level from Process (Mdn=-0.12, N=8 U=9.00, p=0.01, r=0.72). Along the Y axis (SVD2), a Mann-Whitney test showed that Outcome (Mdn=0.10, N=8) was not statistically significantly different at the alpha=0.05 level from Process (Mdn=-0.06, N=8 U=33.00, p=0.96, r=-0.03).

Thematic Structures and Vignette Acceptability: Extreme Cases

All discussions where in Phase 3, the groups reached consensus that fell between 5 and 6 were removed. Discussions where consent was four or lower, representative of the bottom 25% of the data was coded as 'Low acceptability' and discussions representative of the top of the distribution was coded as 'High acceptability'. Decrease of the number of observations changed the set of analysed structures. These are the only models that include 'risk' code as it has different connectivity within these scenarios. Codes such as actor_subject and actor_receiver were removed as they were closely located and central to all discussions.

These analyses are supplementary to Figure 3C in the main manuscript. In Figure 3C1 in the main manuscript, analyses were conducted around standard epistemic networks analysed at the unit of 'vignette and groups', at the level of individual utterance (infinite stanza), with with dominant codes and central codes removed.

In Figure 3C2 in the main manuscript, analyses were conducted around standard epistemic networks analysed at the unit of 'vignette across groups', at the level of individual utterance (infinite stanza), with with dominant codes and central codes removed.

Only vignette-group level structures have significant statistical difference, yet vignette level structures have lower number of observations and maintain the same patterns.

	Group-Vignette Model		Vignette Model	
	x-axis	y-axis	x-axis	y-axis
Variance	20	15	36	20
Pearson corr (gof)	0.95	0.93	0.98	0.98
Spearman corr (gof)	0.96	0.94	0.92	0.91
Mann-Whitney test	p=0.03	ns	ns	ns
Cohen's d	r=0.30	-	-	-

Table 8: Comparison of Variance, Goodness of fit, and Statistical Comparisons across Different Vignette Levels

Acceptability Model 1. Settings: Ordered network, Infinite stanza, Group-Vignette Comparisons, Heavy and central codes removed. Along the X axis (SVD1), a Mann-Whitney test showed that Low (Mdn=0.16, N=26) was statistically significantly different at the alpha=0.05 level from High (Mdn=0.00, N=58 U=530.50, p=0.03, r=0.30). Along the Y axis (SVD2), a Mann-Whitney test showed that Low (Mdn=0.00, N=26) was not statistically significantly different at the alpha=0.05 level from High (Mdn=0.00, N=58 U=581.50, p=0.09, r=0.23).

Acceptability Model 2. Settings: Standard network, Infinite stanza, Vignette Comparisons, Heavy and central codes removed. Along the X axis (SVD1), a Mann-Whitney test showed that High (Mdn=-0.18, N=16) was not statistically significantly different at the alpha=0.05 level from Low (Mdn=0.09, N=11 U=62.00, p=0.21, r=0.30). Along the Y axis

(SVD2), a Mann-Whitney test showed that High (Mdn=-0.06, N=16) was not statistically significantly different at the alpha=0.05 level from Low (Mdn=0, N=11 U=102.00, p=0.51, r=-0.16).

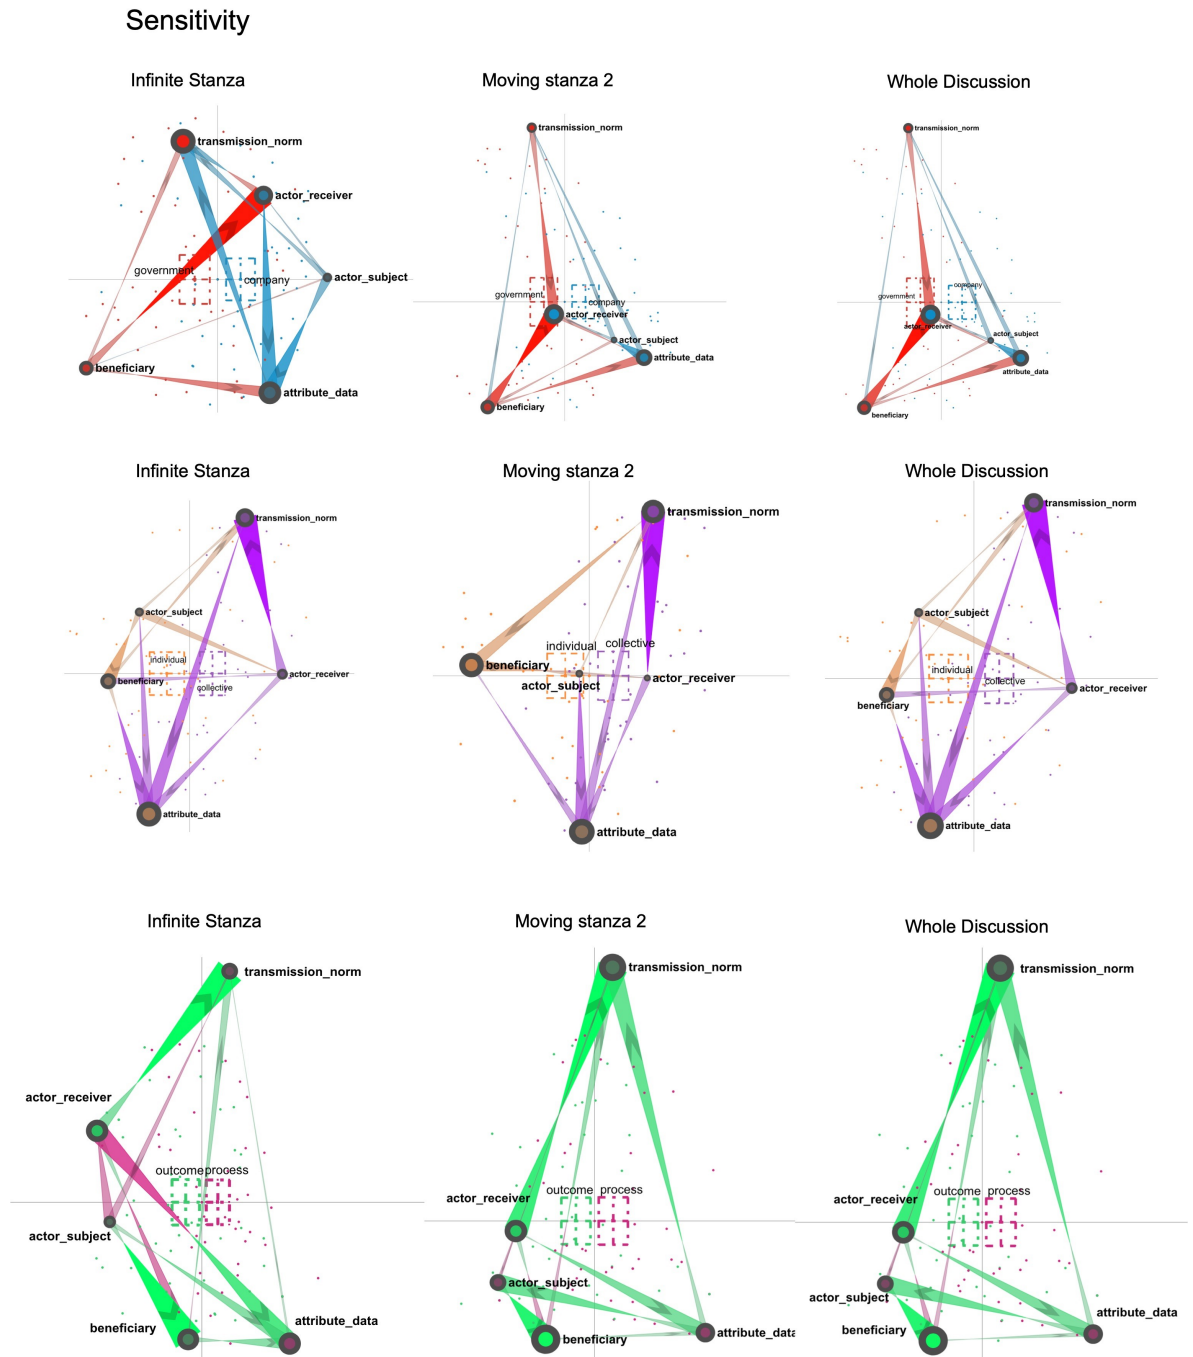


Figure 14: Sensitivity analyses of theme co-occurrence at varying levels of aggregation: utterance, dyad, entire discussion

References

- Ferreira Mello, R., & Gašević, D. (2019). What is the effect of a dominant code in an epistemic network analysis? In *Advances in quantitative ethnography: First international conference, icqe 2019, madison, wi, usa, october 20–22, 2019, proceedings 1* (pp. 66–76).
- Long, J. A. (2022). *Jtools: Analysis and Presentation of Social Scientific Data*.
- Marquart, C., Hinojosa, C., Swiecki, Z., Eagan, B., & Shaffer, D. (2018). Epistemic network analysis (version 1.7.0)[software]. Available from app.epistemicnetwork.org.
- Mehrabian, A., & Stefl, C. A. (1995, January). BASIC TEMPERAMENT COMPONENTS OF LONELINESS, SHYNESS, AND CONFORMITY. *Social Behavior and Personality: an international journal*, 23(3), 253–263.
- Shaffer, D. W. (2017). *Quantitative ethnography*. Lulu.com.
- Shaffer, D. W., Collier, W., & Ruis, A. R. (2016). A tutorial on epistemic network analysis: Analyzing the structure of connections in cognitive, social, and interaction data. *Journal of Learning Analytics*, 3(3), 9–45.
- Siebert-Evenstone, A. L., Irgens, G. A., Collier, W., Swiecki, Z., Ruis, A. R., & Shaffer, D. W. (2017). In search of conversational grain size: Modeling semantic structure using moving stanza windows. *Journal of Learning Analytics*, 4(3), 123–139.
- Tan, Y., Ruis, A. R., Marquart, C., Cai, Z., Knowles, M. A., & Shaffer, D. W. (2022). Ordered network analysis. In *International conference on quantitative ethnography* (pp. 101–116).