



ARTICLE



<https://doi.org/10.1057/s41599-025-04776-1>

OPEN

Data Slots: trade-offs between privacy concerns and benefits of data-driven solutions

Martina Mazzarello^{1,5}, Fàbio Duarte^{1,5}, Simone Mora^{1,2,5}, Cate Heine^{1,3} & Carlo Ratti^{1,4}

Cities are increasingly adopting data-driven solutions derived from diverse digital media interactions—from geolocated social media posts and self-tracking apps to CCTV surveillance and transit cards. We define a data-driven solution as an application, service, or device that leverages big data. However, concerns persist about the potential trade-off between the benefits of these solutions and individual privacy. To assess people's opinions on these trade-offs, we designed a physical and digital card game, Data Slots, which we played with participants around the world. In this game, cards embody data possibilities, enabling players to trade cards, develop their data-driven ideas for solutions, assess other players' proposals regarding benefits and privacy concerns, and strategically invest in their preferred solution. Here we describe the results based on more than two thousand times Data Slots have been played, in 79 countries. We show that perceived privacy concerns as well as benefits are not intrinsic values of specific solutions or datasets, but rather they are combinatorial, situational, transactional, and contextual. By understanding the complex interdependencies that shape public attitudes, policymakers, developers, and stakeholders can refine their approaches to prioritize privacy while harnessing the advantages of data-driven technologies.

¹Senseable City Lab, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge 02139 MA, USA. ²Department of Computer Science, Norwegian University of Science and Technology, Postboks 8900, Trondheim 7491, Norway. ³Centre for Advanced Spatial Analysis, University College London, Gower Street, London WC1E 6BT, England. ⁴ABC Department, Politecnico di Milano, Milan, Italy. ⁵These authors contributed equally: Martina Mazzarello, Fàbio Duarte, Simone Mora. ✉email: fduarte@mit.edu; moras@mit.edu

Introduction

Cities worldwide are declaring impressive goals for making policy decisions based on data, ranging from improving traffic flow to providing better public health services. Most of these solutions rely on data collected through various sources, including sensors deployed in public spaces, personal metadata collected by telecommunication operators, social media platforms, and self-tracking devices. However, the pursuit of data-driven solutions raises significant privacy concerns, prompting a global dialog on the ethical balance between accessing and using potentially invasive data and the benefits derived from these projects.

Privacy concerns are prominent worldwide: more than 70% of Nigerian, Egyptian, and Indian internet users are disquieted about their online privacy; only German users' concerns are lower than 30% (Petrosyan, 2019). In Europe, citizens have been voicing their worries about unintended effects and malicious uses of data and artificial intelligence (Commission, 2020); and in the United States, the White House has opened to public discussion the Blueprint for an AI Bill of Rights, intended to address the challenges posed by technology, data, and automated systems (House, 2024).

On the city scale, governments are taking actions to ensure that data is collected, stored, and applied ethically, protecting residents' privacy (Andrew Collinge, 2024). Amsterdam and New York City have created data privacy offices, and cities in North, Central, and South America have joined forces to establish the City Data Alliance to improve data-driven governance (Network, 2022). Companies that harvest a lot of data frequently used in urban studies, after several backlashes such as Sidewalk Lab's Quayside Smart City Project (Baldwin, 2024), have not only improved their privacy terms, but also created initiatives to promote the responsible use of their data, such as Facebook's Data for Good, Google's Community Mobility, and Spectus Social Impact program.

Yet, all these initiatives still rely on powerful companies or government bodies collecting and curating data. Because in their hands lies the power of implementing solutions at scale, they are *de facto* deciding on behalf of the public the trade-offs between privacy concerns and benefits of solutions that will be available worldwide.

In this paper, we define "trade-off" as "willingness to release personal data in exchange for perceived benefits", based on references such as Lutz et al. (2018), Ashworth and Free (2006), Sheehan and Hoy (2000). We define "privacy concerns" as concerns around the "perceived vulnerability and perceived ability to control submitted personal information when using the Internet" (Dinev and Hart, 2004), and more specifically, as "the disclosure and sharing of personal data", which make "users vulnerable to the potential loss of control over the spread and use of these data" (Lutz et al., 2018). Lutz et al. (2018), based on multiple references, also give a broad definition of the benefits of the digital economy, including "bonding and solidarity (...) financial profit, synergies (...) status improvement (...) increased environmental sustainability". The bottom line of Lutz et al.'s (2018) argument, as is ours, is that it is not the role of the researchers to define which are the benefits or the risks, because they are different for each user, and might be objective or subjective. This is also important in Acquisti et al. (2015) review paper, with their emphasis on the analysis that privacy concerns are malleable as we further develop in our discussion.

Scholars have also been voicing their apprehension from multiple perspectives. They range from the denouncement of data colonialism (Thatcher et al., 2016), which reflects the commodification and power asymmetry between data producers (us, who intentionally or unintentionally generate data) and data owners

(those who incorporate this data in their technologies and products), to philosophical reflections about whether AI systems should be embedded with moral values or, instead, should be allowed to develop some level of moral autonomy (Serafimova, 2020) and algorithmically develop some capacity "for moral reasoning, judgment, and decision making" (Coeckelbergh, 2020). Other perspectives address the social consequences of data obfuscation (Sareen et al., 2020). It occurs when people and social phenomena are not socially legitimized through data and become invisible, or when, on the premise of achieving social safety against terrorism or a pandemic, governments use personal data for permanent governmental surveillance even after threats dissipate (Tan et al., 2022). Researchers have also pointed to the disjointedness between people's choices to be used as building blocks in machine ethics and ethicists' approaches (Awad et al., 2018), and the policy implications of balancing people's moral intuitions with 'experts' intuitions' and ethical theories (Savulescu et al., 2019).

In this article, we propose to measure how multiple stakeholders, from academics to community residents to government officials, from diverse demographic groups in terms of gender and age, and people with different cultural backgrounds around the world, perceive the trade-off between privacy concerns and the benefits of data-driven solutions. We aim at understanding the role of different datasets in driving people's perceptions of benefits and privacy concerns and what factors may affect those perceptions.

Some authors propose to quantify the value of privacy. Il-Horn Hann and Png (2007) surveyed 268 participants from the United States and Singapore to see the precise financial premium (or cost), which ranged between \$30.49 and \$44.62 when facing trade-off situations between loose privacy policies regarding the handling and use of personal information with financial gains and convenience provided by these services; or Acquisti et al. (2015) and Acquisti et al. (2013), who found that users who had an anonymous gift card of \$10 would be less willing to exchange them with a privacy-invasive gift card of \$12 than the opposite. Although these approaches are interesting, the benefits and costs are given by the researchers, in these cases, in financial terms. However, as discussed by the same author (Acquisti et al., 2015) and Steinfield et al. (2008), trade-offs between releasing private information and the benefits provided by these companies are often, and more importantly, subjective in nature. Thus, our approach was that it should not be us, the researchers, to define the benefits or risks, but they would emerge from each bottom-up proposal, defined by each player and evaluated by their peers. This became even more important after reading one of papers suggested by the reviewer, which discusses that one of the main reasons for the phenomena of 'privacy fatigue' or 'privacy apathy', is "how a lack of agency and a sense of futility can impact privacy" (Draper et al., 2024).

An important point that has been raised by scholars is the discussion on individual and collective privacy. Gilliom (2011) sets the tone that privacy studies have been "hyper individualistic, spatial, legalistic, blind to discrimination, and, in the end, simply too narrow to catch the richness of the surveillance experience." Marwick (2022), reviewing the programs of conferences on privacy, discusses how privacy scholars should take into account that privacy is unequally distributed (with the usual unbalance favoring White, male, and wealthy people). And closer to the scope of our work, Galič (2022) makes the argument that it is not because some smart city frameworks do not collect any personal or private information (a 'privacy-preserving' approach) that this could not pose risks and profiling of specific groups.

We define a data-driven solution as an app, a service, or a device that operates by leveraging big data. We collect opinions from more than two thousand people through a card game developed by the authors, called *Data Slots*, that has been implemented in physical (played in-person) and digital (played remotely) versions. In the game, players have a deck identifying twelve different datasets to select and swap among players, in the in-person game, otherwise randomly chosen in the digital game; and propose data-driven solutions for one of three different scenarios: work, home, and public spaces. Rather than traditional polls or surveys in which experts, city managers, or developers propose solutions and residents only voice their opinion, in *Data Slots* we eschew top-down approaches, and instead allow participants to come up with their ideas, assess other players' proposals in terms of benefits and privacy concerns; and decide which ideas they would invest in. The game provides an environment to elicit players' perceptions, preferences, and fears, facilitating reflection and creativity, and allowing researchers to measure them as the play unfolds. Secondly, it also provides a medium to educate about data possibilities and to develop soft skills such as design thinking.

Here, we discuss the empirical evidence about how values people attribute to data vary in combinatorial, situational, transactional, and contextual ways and how these results can inform ethics about data-driven solutions.

Methods

Game design. Interactive tools and gamification have been used to reflect and foster ethical uses of data. The consulting firm IDEO developed the AI Ethics Cards as a tool to help guide an ethically responsible, culturally considerate, and humanistic approach to designing with data; at Microsoft's Ethics & Society organization created the Responsible Innovation Practices toolkit to include moral considerations and the socio-technical implications their proposals; the UN Global Pulse developed a Risk, Harms and Benefits Assessment Tool, a data privacy, ethics and

data protection compliance mechanism designed to help identify and minimize the risks of harms and maximize the positive impacts of data innovation projects; for a review see Wong et al. (2023). Finally, the highly successful Moral Machine (Awad et al., 2018) surveyed millions of opinions around the world to assess people's perception of the ethics of autonomous vehicles' driving decisions.

Our *Data Slots* is a card game. It has an in-person and a digital version, which have the same cards, the same rules, and, for the most part, the same game dynamics. Both versions have three scenario cards, two identical decks of twelve data cards, and ten investor chips per player. Cards and chips are either physical or digital, as shown in Fig. 1. The scenario cards are: *home*, *work*, and *public spaces*. The twelve data cards are: *personal profile*, *health data*, *dietary habits*, *electronic transactions*, *social networks*, *human mobility*, *animal mobility*, *vehicle mobility*, *utility data*, *environmental data*, *public infrastructure*, and *greenery*. Each card provides a title and a simple descriptive text (see SI for a detailed description of the scenario and data cards). These twelve data categories represented by the cards in *Data Slots* were selected through a systematic process to ensure relevance to real-world applications, theoretical grounding, and playability. The selection reflects a careful synthesis of practical utility, academic research, and public engagement. First, the categories were chosen to represent datasets commonly used in urban policy, technological innovation, and public services. These include, for instance, mobility data for traffic optimization, health data for wellness initiatives, and greenery data for environmental assessments. The relevance of these categories is further supported by their frequent citation in prominent reports such as the European Commission's White Paper on Artificial Intelligence (Kilian, 2020), which discusses the ethical applications of AI in cities. Second, the card typology is rooted in academic frameworks that explore data typologies and societal impacts. For example, discussions of privacy and utility trade-offs, such as those outlined in Thatcher et al.'s *Data Colonialism through Accumulation by Dispossession*, Thatcher et al. (2016) informed the

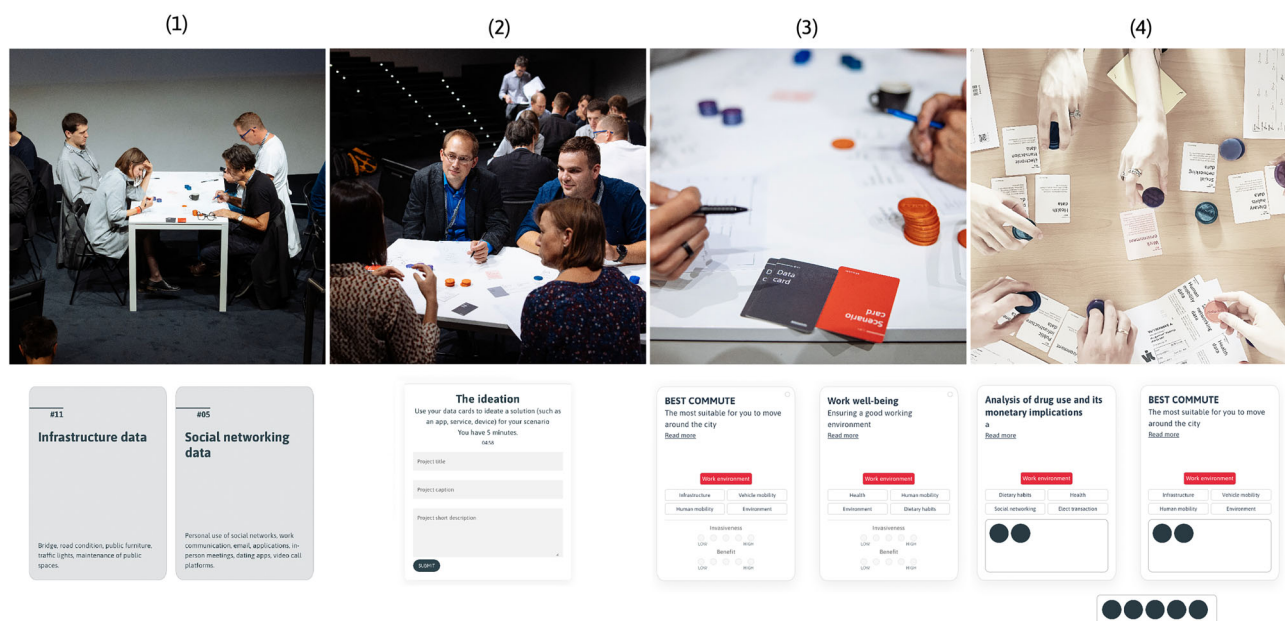


Fig. 1 Four phases of the game. (1) Card selection—players decide which cards to keep and which to transact among each others, for a given scenario, Ideation—players develop a data-driven solutions taking into account the selected cards (2), Assessment—players' mutually rate ideas against benefits and invasiveness criteria (3), Investment—tokens are exchanged among players to rank ideas for their business potential (4). The digital version of the game is depicted in the second row (digital cards and tokens).

inclusion of datasets like electronic transactions and personal profiles. Similarly, the ethical implications of specific data categories, as highlighted in studies like Awad et al.'s Moral Machine Experiment, Awad et al. (2018) underscore the significance of mobility data in assessing public perceptions of privacy and benefits. Finally, the cards were designed to balance accessibility with representativeness, making them comprehensible and relatable across diverse cultural and societal contexts. For instance, categories were abstracted into broader yet tangible concepts, such as "Health Data" rather than granular biometric details. This approach ensures the inclusion of underrepresented datasets like animal mobility or greenery data while maintaining the familiarity of categories like social network data. In the in-person game, participants play in turns, in teams of four players aided by a game master. In the digital version, available at [website affiliated to research institute omitted], participants play individually and interact asynchronously. In both versions, players start by randomly selecting a scenario card and receive a set of three data cards. The game has two rounds of data cards' swapping and one of retrieving one card from the card bank, with small differences between the in-person and digital versions, explained below. Following, the game can be divided into four phases.

The first phase is the card selection. It starts with each player receiving three unique cards and proceeds with three-card transactions, so in the end, each player has four cards. The second phase is the ideation, in which players, using the cards for inspiration, come up with idea briefs to be presented to other players synchronously in the in-person game and asynchronously in the digital game. In the third phase, players assess the benefits and privacy concerns of their own and other players' proposals, using a scale from 1 to 5. In the fourth and final phase, players must invest the 10 chips they have received in other players' proposals. The four phases of the games for the physical and digital versions are depicted in Fig. 1. For a detailed description of game mechanics, see SI.

The main difference between the in-person and digital versions of the game is that in the former, the four players know which set of cards each player has and which card transactions they make during the game. In the latter, since it is a single-person game, the player only sees the sets with which they can transact cards. However, since in the in-person game the transactions follow a specific order (see SI), the fact that the players know each other's set of cards should not influence their decisions about which cards to transact. In the evaluation and investments phases, only ideas generated by players who were given the same scenario are selected by the system, which mimics the physical game where all players share the same scenario card picked at the beginning of the session.

During in-person game sessions the game rules and mechanics were presented to participants from the start, by one of this paper's author. Participants were also invited to ask questions to clarify any doubts regarding game rules or the purpose of the game. In the online version of the game the same information is provided through a video tutorial.

Data Slots was designed as a data collection tool aimed at removing the implicit power imbalance of surveys or activities in which the proponent either poses the questions or guides the activities. The results from playing Data Slots in multiple countries are presented in the remaining of the paper.

Data collection. This study was approved by the Institutional Review Board (IRB) at [research institute omitted]. Researchers did not collect any personal information from participants. The sample size was not predetermined using statistical methods. In

the in-person version of the game, participants had direct contact with researchers; in the digital version, researchers were blinded to any feature of the participants.

The in-person version of Data Slots includes 700 plays, in which all players completed the four phases. In the digital version, 1493 players finished the card selection phase, 656 the ideation phase, 611 the privacy and benefits assessment, and 313 the investment phase. We can speculate that the drop in plays, in the digital version, between phases is due to the participants terminating their game sessions (e.g., closing their browsers) or internet connection issues. Although participants could leave the in-person game anytime with no penalty whatsoever, no one left a game session once started.

In-person play. In the in-person version, the game is played by groups of four people, with physical decks of cards and chips. The game was conceived to be played in two rounds, lasting approximately 25 min each. Players sign a consent form and receive one *scoring sheet* per round, in which they include their declared gender, age, and how they identify themselves as a public official, scholar, or resident. The scoring sheet is provided in SI. Each player receives 10 colored chips (orange, blue, green, purple), which will be used in the last phase of the game—the investment. Each of them starts with a set of 3 cards from one deck, randomly distributed. From a second identical deck of 12 data cards, 4 are randomly selected and placed in the center of the table. Data Slots is a fully open game, meaning all players know which cards other players have. Players sit around the table following the same clockwise direction based on the color of their chips: orange, green, blue, and purple.

The in-person version of Data Slots has been played so far in 17 countries, between March 2022 and October 2023. It was first played in Amsterdam, The Netherlands. As intended for the in-person version, 85.8% of players played two rounds. Italy has the highest number of plays (96 plays, or 13.7% of the total), and India has the lowest (4 plays, or 0.6% of the total). Most of the in-person players are between the ages of 18 and 35, 59.5%/39.7% are self-declared male/female, respectively, and in terms of roles, 69.7% are scholars, 7.7% are city officials, and 22.6% are residents (citizens not involved in city administration or research). The game was translated into six languages (Italian, Basque, Spanish, Arabic, Korean, and Japanese). The game was translated only if the context required it to make participants comfortable in sharing their ideas.

Digital play. In the investment phase of the digital version (phase four in Fig. 1), three randomly selected ideas based on the same scenario, retrieved from a bank of ideas, are displayed to the player. When the digital version of Data Slots was launched, the bank of ideas had thirty ideas per scenario. These ideas were selected from in-person plays, transcribed and translated into the 12 languages. In sequence, each new digital idea generated was added to the bank of ideas. As of April 2024, the bank had 656 ideas.

The digital version was launched in March 2023, in 12 languages (Arabic, Chinese, English, French, German, Hindi, Italian, Japanese, Korean, Portuguese, Spanish, and Russian). We don't collect players' IP addresses, but players self-declare the country in which they were playing, their age, and, optionally, their gender. Mexico has the highest (19.5%) number of online players, followed by Italy (18.8%) and Indonesia (10.7%). As in the physical version, players are predominantly young, between the ages of 16 and 35 (78.6% of players). Countries in which Data Slots has been played, in their in-person and digital versions, are highlighted in Fig. 2, and participant counts per country are provided in SI.

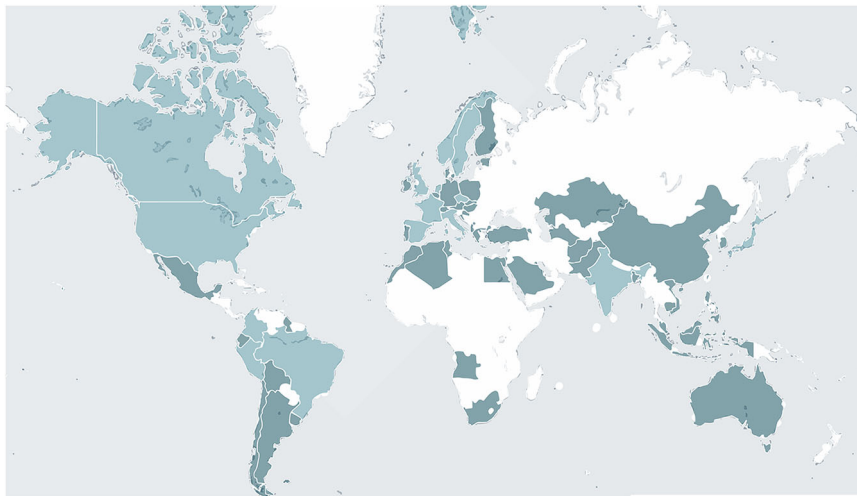


Fig. 2 Map showing countries where Data Slots were played worldwide. Light blue shading represents countries where both physical and digital versions were played, while dark blue indicates countries where only the digital version was played.

Data analysis. Data was analyzed in Python 3.7.3. Correlations were calculated using scipy version 1.6.1, and all regressions are Ordinary Least Squares regressions run using statsmodels 0.13.5. Error bars in probability plots represent a 95% confidence interval around the expected value. Error bars in regression plots represent a 95% confidence interval, or 1.96 standard errors from the point estimate.

Results

The value of data: cards locked, discarded, selected. The game starts with each player receiving three random (digital or physical) data cards and a disclosure of a scenario (work, home, public space) for which players will come up with a data-driven solution.

While work and home spaces are self-explanatory, we define “public space” as physical realms shared by potentially unknown individuals (we are not specifically including social media, though, in our bottom-up approach, players could include social media in their proposals).

The first action requires players to lock one card, which cannot be transacted throughout the game. Therefore, the locked card signals the dataset that is perceived as the most relevant for the player. The game proceeds through two turns of card transactions between players.

Among the cards randomly distributed to the players (the list is provided in SI), the human mobility data card is most often locked, 43% of the time. Human mobility is followed closely by health, utility, environmental, and personal profile data cards, without statistically significant differences in the locking probabilities among these cards. There are variations in locking probability across scenarios, though, indicating that the data categories are valued differently in different contexts: human mobility is most likely to be locked in the public space scenario, utility data is most likely to be locked in the home scenario, and health data is most likely to be locked in the work scenario. On the other hand, the animal mobility data card is by far locked the least across all scenarios—if a player is dealt the animal mobility data card, they choose to lock it only 10.7% of the time (9.1% in the physical game and 11.8% in the digital game). In general, the locking probabilities of each card in the digital and physical games are highly correlated with one another (Fig. 3, Pearson’s correlation coefficient = 0.914, p -value = 0). SI includes the tables for all points discussed below.

We can observe that participants play the digital and in-person games similarly: both groups appear to value data cards in similar

ways, as evidenced by their likelihood to lock and discard cards (Fig. 3). Therefore, we combine both, in each phase of the game, to present the results. However, as highlighted by second-discard choices (Fig. 3-right), some differences in game dynamics give special insights into how people see the trade-offs between the benefits and privacy concerns of data-driven solutions; to be discussed separately.

While the locked card signals the most valued card to the player, the first card they dispose of in a transaction represents the least important card (among the three given cards). With more than 2000 plays between digital and physical versions and the distribution of 12 cards per round, each card had a very similar probability of being among the three cards distributed to the first player to dispose of a card. We found that, conditional on it being in their hand, players are most likely to dispose of the animal mobility data card in all scenarios: 68.7% of the cases for the home scenario, 76.6% for the work scenario, and 67.4% for public spaces.

The player who first discards a card needs to pick another from a set of three, randomly placed on the table. Again, considering more than 2000 plays and the 12 cards in each set, the cards have a very similar probability of being among those open on the table. If the locked card is arguably the most important card the first player has in hand, the card transacted with the discarded card (the least important) is also relevant to the player’s strategy. The cards most often transacted in exchange for the least important are health data for the home scenario (23% of cases), environmental data for the work scenario (24% of cases), and human mobility data for public spaces (24% of cases).

Considering the locked, first discarded, and first picked card (transacted for the first discarded card), we found that the most valued datasets are health and human mobility data, and the least valued dataset is animal mobility data, with significant variation across scenarios (Fig. 4). For example, utility and health data are more valued in the home scenario; health, personal profile, and social networks data are more valued in the work scenario; and human mobility, environmental, and infrastructure data are more valued in the public space scenario.

Drivers of benefits and privacy concerns. To analyze whether and which cards drive players’ assessment of the benefits and privacy concerns of different proposals in specific scenarios, our sample includes the 700 in-person plays and the 611 digital plays that have completed the assessment phase.




Fig. 3 Cards being locked, discarded first, and discarded second in physical and digital plays. Panels **A–C** represent probabilities from games played digitally are plotted on the x-axis ($N = 887$), and probabilities from games played in person ($N = 675$) are plotted on the y-axis. The gray line represents the line $x = y$, around which data points are expected to fall if there is a high correlation between how cards are valued in the physical and digital games. Pearson's correlation coefficients are displayed in the upper left corner of each panel.

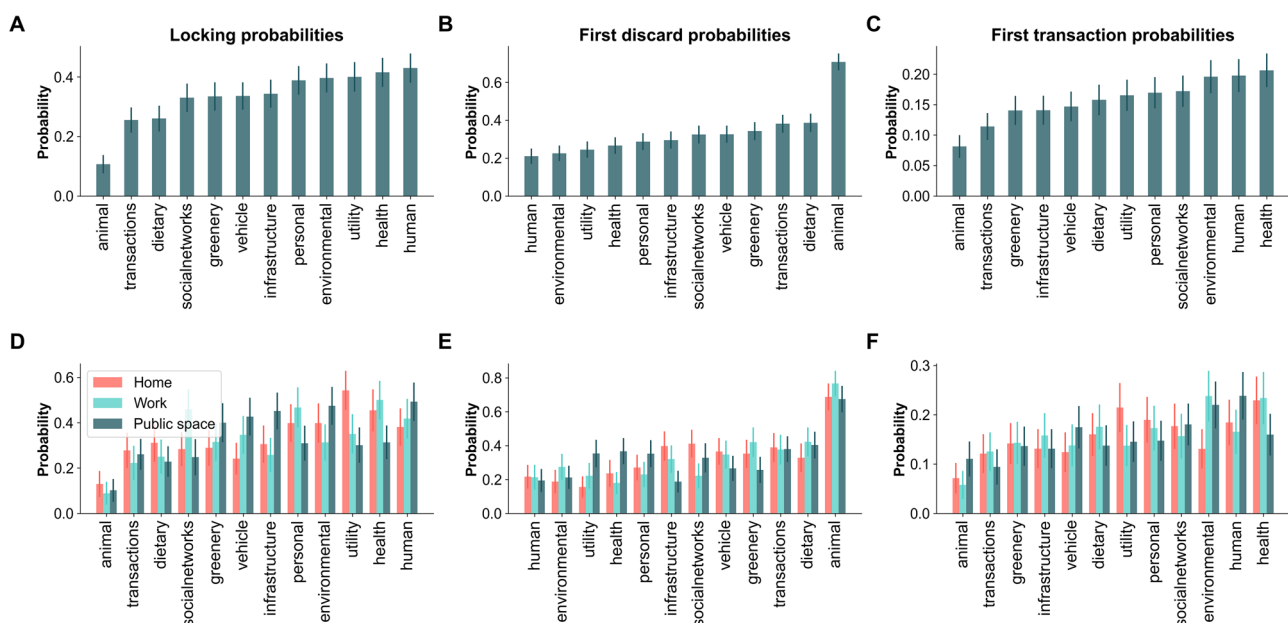


Fig. 4 Distribution of probabilities of each card being locked, discarded first, and discarded second according with home, work and public space scenarios. Probabilities of each card being locked, discarded first, and discarded second, for both in-person and digital games ($N = 1562$) are represented in panels **A–C**. Panels **D–F** represent the same probabilities, broken down by games played under the home, work, and public space scenarios. Error bars represent a 95% confidence interval around the sample mean ($\pm 1.96\sigma/\sqrt{n}$).

First, we analyzed whether (i) any set of cards would drive ideas deemed to impact players' perception of privacy invasiveness and benefits, and (ii) whether there was any difference based on work, home, and public spaces. We fit ordinary least squares (OLS) regressions with binary variables indicating whether each card is in a player's final set of cards as independent variables and the average invasiveness or benefit score given to that player by their competitors as the dependent variables. The regression coefficients of this model, representing the weight each card has in driving the assessment in terms of privacy concerns and benefits of data-driven proposals, are provided in SI. Results of OLS are summarized in Fig. 5.

As seen in Panels A and B of Fig. 5, we found that individual cards are not generally associated with higher or lower benefit or investment ratings (see Panel B of Fig. 5). This may indicate that the perceived benefit of the ideas comes less from the types of data used and more from the significance of the ideas and how the datasets are combined.

More interesting than which single cards were driving players' assessments of privacy concerns and benefits of data-driven proposals was the finding that the combination of data cards matters. The cards that were more present in proposals with the highest ratings in terms of benefits were also in some of the lowest-rated proposals. Likewise, data cards, more often present in proposals deemed the most invasive, were also present in the final deck of proposals considered the least invasive. Thus, one of the key findings is that the combination of data cards matters more than particular data cards.

Figure 6 shows the distribution of invasiveness, benefit, and investment scores for final proposals containing each individual card (row 1), the ten most popular combinations of two cards (row 2), and the ten most popular combinations of three cards (row 3). The small effect, seen previously in Fig. 5, that individual cards have on perceived value of proposals is reiterated here in the wide distribution of scores—distributions are wide and overlapping; there are proposals containing the most popular

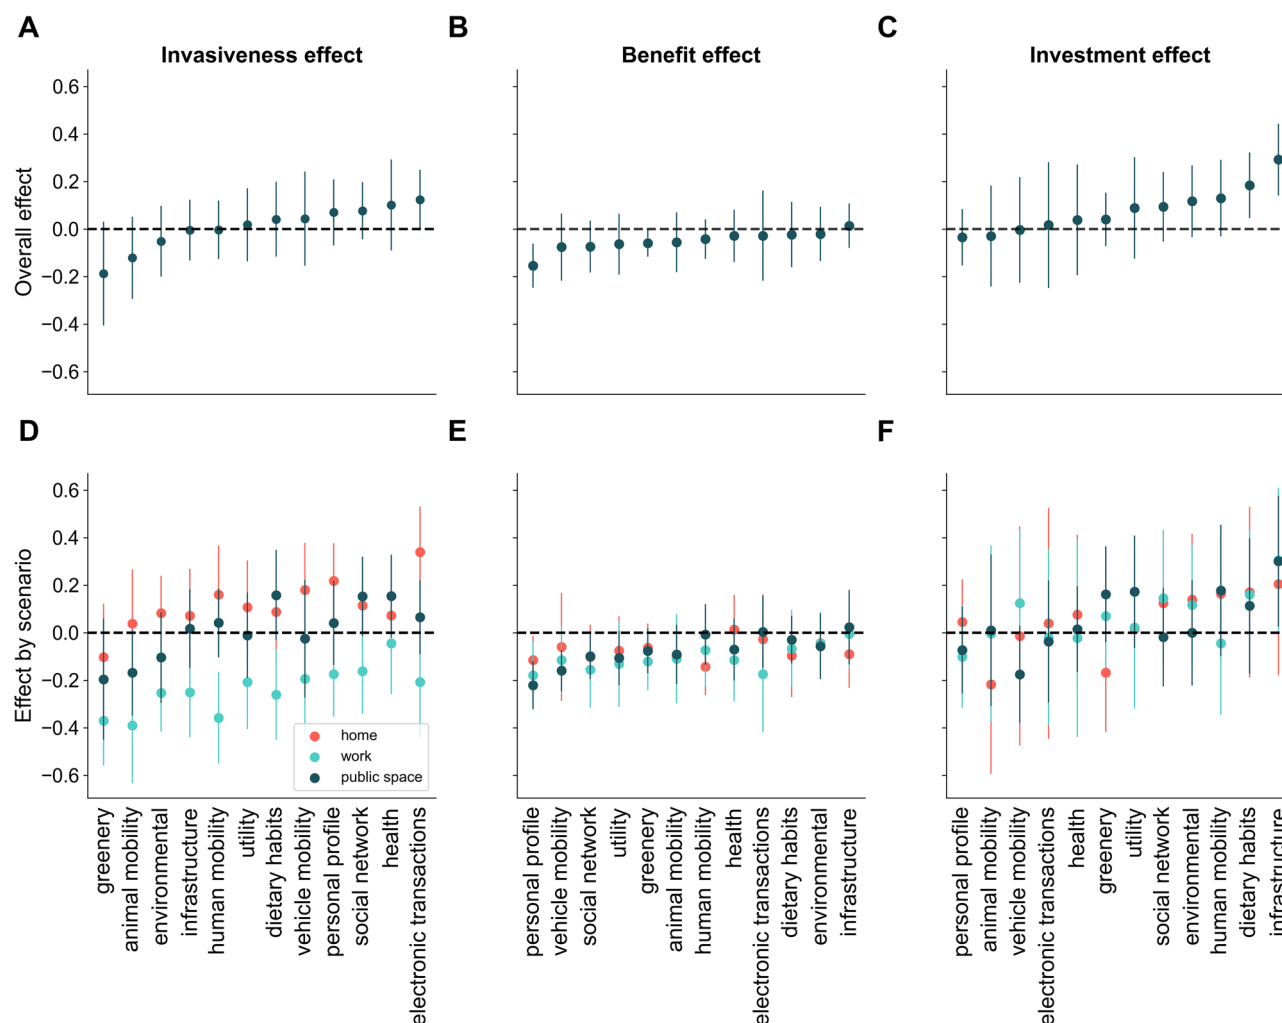


Fig. 5 Impact of each card on perceived invasiveness, benefit and investment for home, work and public space scenarios. Coefficients from an OLS regression where the dependent variable is invasiveness rating (A), benefit rating (B), and total investment by other players (C), and the independent variables are dummy variables representing whether a given card is in the final card set. Panels D–F represent the same coefficients broken down by home, work, and public space scenarios. Error bars represent 95% confidence intervals (1.96 standard deviations from the mean). The total sample is $N = 991$ observations for which there are no missing invasiveness, benefit, or investment scores.

card (human mobility) which receive very high and very low scores, and the same is true for the least popular card (animal mobility). More variation appears when we look at combinations of three cards; for example, human mobility data appears to be seen as more invasive on average when combined with health and utility data (blue) than when combined with infrastructure and greenery data (pink) or environment and greenery data (brown). While small sample sizes at the level of 3-card combinations make these trends more descriptive than statistically robust, they still point to the idea that combinations of cards mean more with respect to perceived invasiveness and benefit than individual cards.

As seen in comparing Panels A and D of Fig. 5, the relationships between datasets used and ideas' invasiveness ratings vary by scenario. Specifically, certain cards are seen as more invasive when in the context of the home scenario; for example, electronic transactions data and personal profile data are associated with higher invasiveness scores when played in the home scenario. Conversely, animal mobility, infrastructure, and human mobility data are seen as less invasive when played in the work scenario. This implies that perspectives on the invasiveness of projects using certain types of data are highly contextual and depend on the scenario in which that data is used.

An important feature of Data Slots is how people transact cards: which they lock, which they dispose of, and which they pick from the table or other players. We mentioned before that human mobility data is the most often locked card, and that we consider the locked card the most valuable card by the players. Also, we assume that players aim to have their proposals with the highest benefits rating, the lowest privacy concerns, and to receive the most investments. Thus, we investigated whether the cards more often locked were also the ones driving the benefits and privacy invasiveness ratings. This question is difficult to answer in terms of benefits ratings, as we didn't find strong relationships between specific cards and benefits ratings (Panel B in Fig. 5). In terms of invasiveness, under the previously mentioned assumption that players optimize for high benefits and low-invasiveness, we might expect that cards associated with high invasiveness ratings may be locked less often and transacted more often. We find that this is not necessarily the case. Health and electronic transaction data are associated with high invasiveness scores; however, while health data is relatively likely to be locked, electronic transaction data is relatively unlikely to be locked. This may indicate that while both cards are deemed invasive, health data are “worth it” in some sense—their perceived value outweighs their perceived invasiveness—while the same isn't true




Fig. 6 Distribution of investment (column 1), benefit (column 2), and invasiveness (column 3) scores for $N = 1777$ final proposals containing each individual card (row 1), the ten most popular combinations of two cards (row 2), and the ten most popular combinations of three cards (row 3). Boxes represent the 25th percentile, median, and 75th percentile; whiskers represent 1.5 interquartile ranges from the nearest edge.

for electronic transaction data. On the one hand, we hypothesized that the locked cards would drive proposals with the highest benefits and lowest privacy concerns. On the other hand, we hypothesized that the first card to be discarded would have the opposite effect: the player would get rid of the card (e.g., animal mobility, the most discarded card) they thought would pull the benefits ratings of their proposal down, or increase privacy concerns. This does not seem to be the case: for example, animal mobility data is seen as low-invasiveness, a purportedly positive characteristic, but is locked by far the least and discarded by far the most. Health data is seen as highly invasive, which one might assume to be unattractive to players, but it is locked second-most frequently out of all the data cards.

The variation of how cards are valued when they are locked, discarded, or transacted, and their influence in the overall assessment of the benefits and privacy concerns of the player's solutions indicates that data cards do not have absolute values, but rather transactional, changing with the flow of the game.

Another analysis of the transactional values of the data cards we tested referred to whether proposals with the highest benefits and lowest privacy concerns would receive more investments. Considering the global uneasiness with data handling by companies and governments, we expected to see a clear and positive correlation between benefits and investments. We find that investments are positively correlated with benefit ratings but have a smaller (though still statistically significant) correlation with invasiveness ratings, indicating that players may take an

idea's benefit into account more than its invasiveness when deciding which ideas to invest in (Fig. 7).

Another question we investigated is whether players attribute the same or different values to data cards depending on whether they are assessing their own proposals or other players' proposals.

We compared "other-ratings"—privacy and benefits ratings given by players to other players' proposals—to "self-ratings"—ratings given by players to their own proposals. We found, once more, that data cards do not have absolute values; rather, they vary depending on when the player is assessing their own ideas or the ideas proposed by other players. Specifically, players tend to view personal profile, electronic transaction, health, and social network data as more beneficial when rating their own proposals as opposed to when rating others'. This situational divergence between self-assessment and other-assessment is an intentional feature of the game design, allowing us to observe how personal investment or detachment influences perceived benefits and privacy concerns. During the evaluation phase, players assess proposals using a 5-point Likert scale for both benefits, ranging from low to high, and privacy concerns, ranging from low to high invasiveness. When evaluating their own proposals, players often emphasize the benefits of their chosen data combinations while minimizing privacy concerns, reflecting a personal bias. Conversely, when evaluating the proposals of others, players tend to be more critical, highlighting privacy concerns while perceiving fewer benefits. In the digital version, this phase is conducted asynchronously, with ideas displayed anonymously from a shared

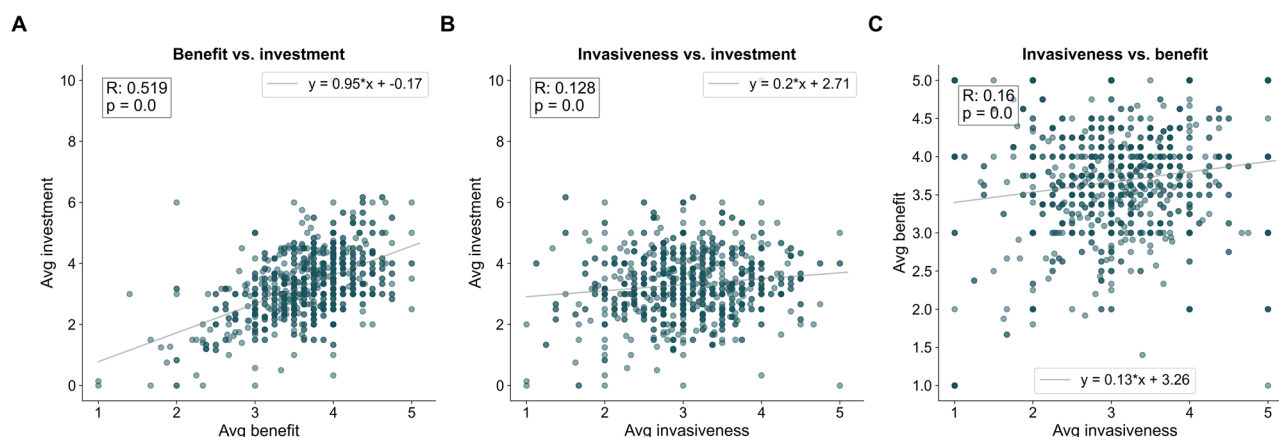


Fig. 7 Transactional values of the data cards. Correlation analysis between benefit and investment (A), invasiveness and investment (B), and invasiveness and benefit (C). Gray lines represent the best fit in an OLS regression. The total sample is $N = 991$ observations for which there are no missing invasiveness, benefit, or investment scores.

pool of proposals. This ensures consistency across evaluations and prevents biases stemming from personal familiarity or the influence of group dynamics. The absence of direct communication during this phase ensures that evaluations reflect individual perceptions rather than being influenced by persuasion or negotiation. These findings, that players rate their own proposals as having both higher benefits and higher privacy concerns compared to others', are not an artifact of unclear rules or limited interaction but instead reflect the deliberate separation of perspectives embedded in the game design. This flows naturally as part of the discussion, avoiding lists while still addressing the reviewer's concern comprehensively.

The change in perceived Invasiveness and Benefit for each data card is provided in SI.

A final question concerned each scenario, independent of the combination of cards: Do privacy concerns and benefits of data-driven solutions vary according to specific scenarios? The average benefits and invasiveness ratings do not vary much by scenario, either in average or in distribution, indicating that either data-driven solutions are not seen as generally more beneficial or privacy-invasive in any given scenario or that ratings are given relative to other ideas within the given scenario (Panels D, E, F in Fig. 8).

In summary, the results show that the values attributed to data are combinatorial (values vary depending on which cards compose the set, without a definite subset of cards driving the results), contextual (values vary according to each scenario, with some cards driving positive results in one scenario and negative in a different scenario), they are transactional (values change with the dynamic of the game), and are situational (vary depending on the player's role: when they are proposing ideas or assessing other players' ideas).

Some of our findings echo previous research, which substantiates our work: for example, Marwick and Hargittai (2019) [information type, context, institution controlling information] aligned with Acquisti et al. (2015) [uncertainty about consequences, context, commercial and governmental malleable/manipulated interests]. Our research reinforces these results by bringing another perspective, because it brings a bottom-up perspective of data handling and solutions. For instance, while Marwick and Hargittai (2019) find that who controls information matters when it comes to privacy concerns and that participants in their focus groups were distrustful of both government and corporate actors, in our work the decision of which data to use, how to combine them, for which purposes, were completely transparent during the entire game, and players themselves were coming up with solutions they proposed to their peers.

Comparing physical and digital plays. For most of the analysis we combined the in-person and digital versions of Data Slots, as they have the same items and rules and the summary statistics show they are comparable (for example, see similarities across the physical and digital game in likelihood to lock and discard given data cards in Fig. 3). Nevertheless, there are a few differences between the two versions that are worth exploring. The most salient one is the fact that, in the in-person version, each player plays the game twice. Thus, we would like to know whether values change from the first to the second round. We found that the cards more often locked and first discarded don't change from round one to round two. Likewise, the combination of cards driving the assessment of the benefits and privacy concerns doesn't change. We find that the distribution of both the benefit and the invasiveness of the proposals remained statistically the same between the first and second rounds (Kolmogorov–Smirnov test statistic = 0.014 for comparison of benefit scores between rounds, 0.021 for invasiveness scores) (see Fig. 8).

Cultural background. Although we have played Data Slots in-person in 18 countries, and people from 74 countries have played the game online (Table in SI), we are very cautious in deriving any cultural generalizations from our sample. Also, in the in-person version, players often had some form of institutional connection with the researchers (for instance, they were invited by institutions that collaborated with the authors' institution). In the digital version, only players with access to the Internet were able to participate. For those using cell phones, only smartphone users could play. In both cases, the player needed to comprehend one of the 12 languages in which the game had been translated. Thus, the results in this section are only indicative.

In spite of these caveats, some interesting results emerge. In terms of self-declared gender and age group, we didn't find any significant difference. When it comes to cultural background, we grouped the countries following the Moral Machine experiment (Awad et al., 2018), and considered the following criteria: most locked card, first card discarded, average benefits assessment, average privacy concerns assessment, the correlation between benefits and investment, and the correlation between privacy concerns and investments; as shown in Table 1.

Discussion

Ethics is at the forefront of public discussions involving how governments and companies handle data, and technologies that produce and make sense of this data. Widely used techniques

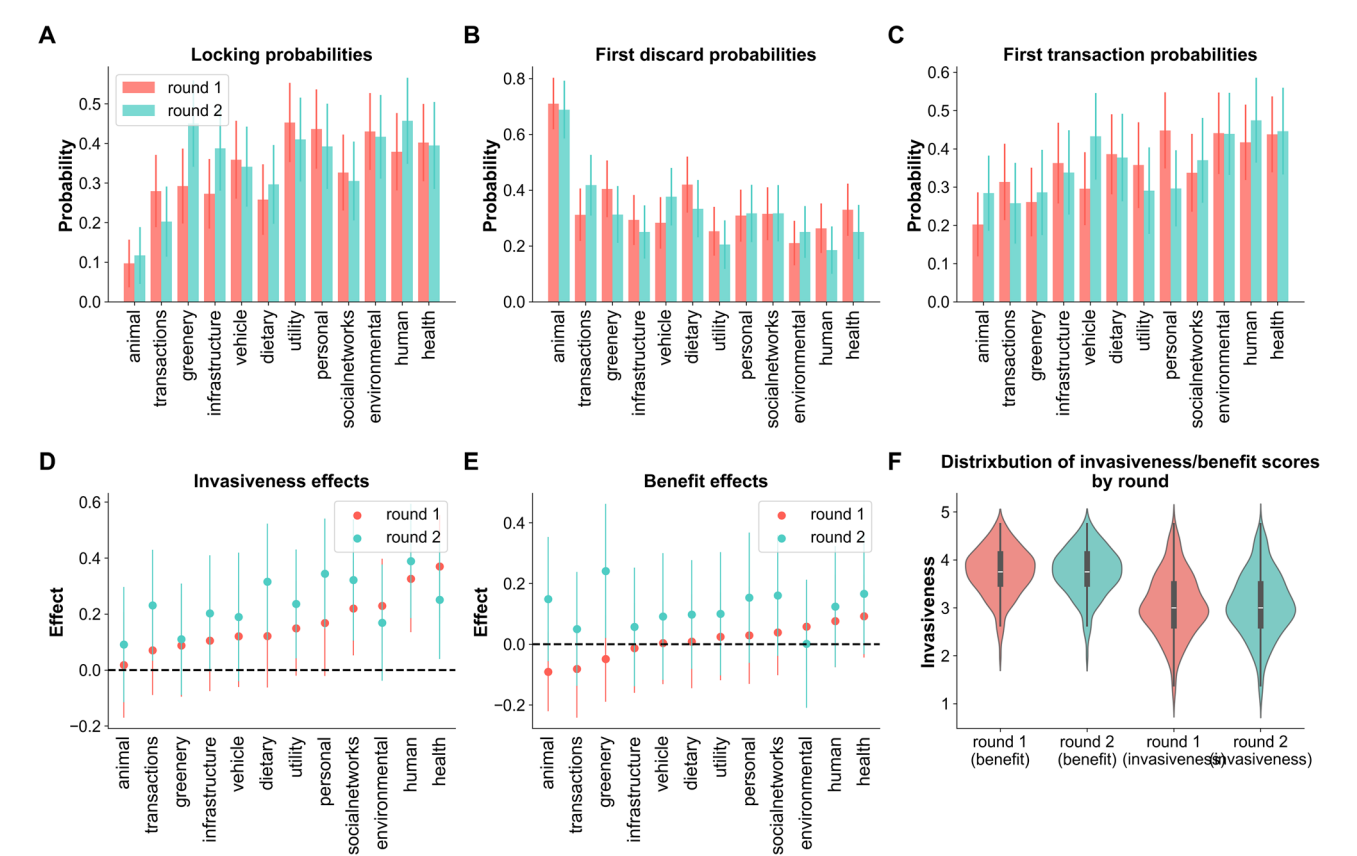


Fig. 8 Privacy concerns and solutions benefits according to different scenarios and different rounds of plays. Locking (A), discard (B), and first transaction (C) probabilities, as displayed in Fig. 4, broken down by individual card and round played. D, E show the benefit and invasiveness effects of individual cards as described in Fig. 5, broken down by round played. F shows the distribution of invasiveness and benefit scores by round, represented by a kernel density estimate as well as a box plot where the box center represents the median and edges represent the 25th and 75th percentiles. Data from in-person plays only ($N = 675$).

Table 1 Different cultural backgrounds and their relationship with participants' decisions in the game.						
Cultural group	Most locked	First discarded	Avg. benefits	Avg. privacy	Benefits vs. Investment	Privacy vs. Investment
Western	Environmental	Animal mob.	3.81	3.06	$R = 0.442, p < 0.001$	$R = 0.092, p < 0.05$
Eastern	Human mob.	Animal mob.	3.93	3.25	$R = 0.523, p < 0.001$	$p > 0.1$
Southern	Health	Animal mob.	3.71	3.02	$R = 0.542, p < 0.001$	$R = 0.154, p < 0.05$

such as machine learning or artificial intelligence often work as black boxes, with their operations not clear even to analysts. The lack of transparency of how governments and companies are using the data we generate increases the concerns people have about data privacy, and the balance between privacy and the benefits brought by data-driven solutions is often tenuous, when not contentious.

Also, values attributed to specific data vary widely across different societal groups and, within them, according to specific situations. As Andreas Theodorou and Virginia Dignum point out, algorithmic approaches to ethical decisions need to take these variations into account, identifying “different orderings of the values and which ethical theory is most suitable in a given situation” (Theodorou and Dignum, 2020).

Our analysis shows that data types do not have intrinsic values in terms of privacy invasiveness (bank transactions are not necessarily more invasive than social media posts), benefits (greenery data is not rated higher for its benefits than solutions based on self-tracking data), or investments (ideas with personal

profile data do not receive necessarily more investments than those that include animal mobility data).

We show that values attributed to data are combinatorial, situational, transactional, and contextual.

Combinatorial because the same data card receives widely diverse assessments in terms of privacy concerns and the benefits it brings, depending on the other three cards it is paired with in the final set. Among the 12 data cards, none on its own had a significantly positive relationship with benefit ratings. The largest individual effect of any given card on invasiveness and investment ratings was similarly relatively small (0.14 standard deviations and 0.12 standard deviations, respectively). These relatively small individual effects indicate that it is not single cards driving invasiveness or benefit ratings or investments, but rather combinations of cards.

Situational because cards have different values when the player is assessing their own proposals and when they are assessing other players’ proposals—people tend to both see more benefits and have more privacy concerns with their own proposals than

the proposals of others, even when the cards used are the same (see Table in SI).

Transactional because the same data card receives diverse evaluations in different phases of the game, which is based on a series of transactions. For instance, although locked cards are supposed to be the most valued by the players (they lock them so nobody can take the card from them, and therefore can use it in their proposal); the most locked cards are not the same set which forms the cards more often assessed as those with the highest benefits or with the lowest privacy invasiveness. Conversely, the first cards to be discarded, which were supposed to have the least value to the player, do not form the set of data cards deemed as the most privacy invasive or with the lowest benefits ratings.

Contextual because each card is assessed differently in terms of privacy concerns and benefits when they are used to propose ideas for each of the scenarios (home, work, and public spaces). Even the same set of cards is assessed differently depending on the scenarios. For example, utility data has a nearly double probability of being locked by a player when the home scenario is selected compared to when the public space scenario is selected. These findings emphasize the importance of understanding how respondents evaluate data types and combinations, reflecting privacy concerns and perceived benefits in varying contexts. This aligns with broader observations in privacy research, such as the context-dependence of privacy concerns (Acquisti et al., 2015) and the limited awareness individuals often have of how data can be aggregated and mined (Marwick and Hargittai, 2019). Such factors can lead to privacy resignation, where respondents may feel they have limited control over their data. The analysis highlights that data valuation is not fixed but depends on multiple factors (combinatorial and situational). If individual data types may have relatively small effects on invasiveness and benefit ratings, it gives greater meaning to combinations in changing benefits and privacy perception. This resonates with Solove's (2021) argument that privacy attitudes are shaped by specific contexts and that general attitudes about privacy concerns or its value should not be reduced to overly simplistic metrics. By adopting a bottom-up approach, this study allows respondents to define the problems and solutions most relevant to their specific contexts. This design choice reflects the dynamic interplay between benefits and privacy concerns, capturing how respondents' perceptions shift when datasets are combined or assessed in different scenarios. These findings and the analytical framework could inform public policies that have been increasingly using data directly collected from citizens or from third parties. This way, governments and stakeholders have a clear and critical understanding of how certain datasets combined in specific ways might impact the acceptance and adoption of certain policies. Future research could further explore how these choices reveal the nuanced trade-offs individuals make between privacy and the perceived utility of data-driven solutions.

Conclusion

In this article, we measured how multiple stakeholders, from diverse demographic groups in terms of gender and age, and people with different cultural backgrounds around the world, perceive the trade-off between privacy concerns and the benefits of data-driven solutions. Through digital and in-person Data Slots plays, we collected opinions from more than two thousand people.

This experiment brings empirical evidence that specific data do not have intrinsic values in terms of privacy invasiveness or benefits it can bring. Rather, they vary. This variation, considering it is combinatorial, situational, transactional, and contextual, if

taken into consideration, could inform better the inevitable trade-offs between privacy concerns and the benefits of data-driven decisions.

Data availability

The dataset generated from the in-person and digital plays is available upon request.

Code availability

The code to replicate the results in the paper can be obtained from the authors upon request.

Received: 5 July 2024; Accepted: 17 March 2025;

Published online: 13 May 2025

References

- Acquisti A, John LK, Loewenstein G (2013) What is privacy worth? *J Leg Stud* 42:249–274
- Acquisti A, Brandimarte L, Loewenstein G (2015) Privacy and human behavior in the age of information. *Science* 347:509–514
- Andrew Collinge SR, Mochizuki Y (2024) Open data is a game changer for cities worldwide—here's how to use it best. <https://www.weforum.org/agenda/2022/11/lessons-from-36-cities-on-having-impact-with-open-data/>
- Ashworth L, Free C (2006) Marketing dataveillance and digital privacy: using theories of justice to understand consumers' online privacy concerns. *J Bus Ethics* 67:107–123
- Awad E, Dsouza S, Kim R, Schulz J, Henrich J, Shariff A, Bonnefon J-F, Rahwan I (2018) The moral machine experiment. *Nature* 563:59–64
- Baldwin E (2024) Sidewalk Labs cancels Quayside smart city project in Toronto (2024). <https://www.archdaily.com/939152/sidewalk-labs-cancels-quayside-smart-city-project-in-toronto>
- Coeckelbergh M (2020) AI ethics. MIT Press
- Commission E (2020) White paper on artificial intelligence a European approach to excellence and trust. https://commission.europa.eu/publications/white-paper-artificial-intelligence-european-approach-excellence-and-trust_en
- Dinev T, Hart P (2004) Internet privacy concerns and their antecedents—measurement validity and a regression model. *Behav Inf Technol* 23:413–422
- Draper NA, Pieter Hoffmann C, Lutz C, Ranzini G, Turow J (2024) Privacy resignation, apathy, and cynicism: introduction to a special theme. Sage Publications
- Galič M (2022) Smart cities as 'big brother only to the masses': the limits of personal privacy and personal surveillance. *Surveill Soc* 20:308–311
- Gilliom J (2011) A response to Bennett's 'In defence of privacy'. *Surveill Soc* 8:500–504
- House TW (2024) Blueprint for an AI bill of rights. <https://www.whitehouse.gov/ostp/ai-bill-of-rights/>
- Il-Horn Hann S-YTL, Hui Kai-Lung, Png IP (2007) Overcoming online information privacy concerns: an information-processing theory approach. *J Manag Inf Syst* 24:13–42
- Kilian G (2020) White paper on artificial intelligence—a European approach to excellence and trust. Policy Commons
- Lutz C, Hoffmann CP, Bucher E, Fieseler C (2018) The role of privacy concerns in the sharing economy. *Inf Commun Soc* 21:1472–1492
- Marwick A (2022) Privacy without power: what privacy research can learn from surveillance studies. *Surveill Soc* 20:397–405
- Marwick A, Hargittai E (2019) Nothing to hide, nothing to lose? Incentives and disincentives to sharing information with institutions online. *Inf Commun Soc* 22:1697–1713
- Network BC (2022) America join city data alliance to chart new frontiers for data-driven government. <https://bloombergcities.jhu.edu/news/22-mayors-north-and-south-america-join-city-data-alliance-chart-new-frontiers-data-driven>
- Petrosyan A (2019) Share of internet users who are more concerned about their online privacy compared to a year ago as of February 2019, by country. <https://www.statista.com/statistics/373322/global-opinion-concern-online-privacy/>
- Sareen S, Saltelli A, Rommetveit K (2020) Ethics of quantification: illumination, obfuscation and performative legitimization. *Palgrave Commun* 6:20
- Savulescu J, Kahane G, Gyngell C (2019) From public preferences to ethical policy. *Nat Hum Behav* 3:1241–1243
- Serafimova S (2020) Whose morality? Which rationality? Challenging artificial intelligence as a remedy for the lack of moral enhancement. *Humanit Soc Sci Commun* 7:119
- Sheehan KB, Hoy MG (2000) Dimensions of privacy concern among online consumers. *J Public Policy Mark* 19:62–73

- Solove DJ (2021) The myth of the privacy paradox. *Geo Wash L Rev* 89:1
- Steinfeld C, Ellison NB, Lampe C (2008) Social capital, self-esteem, and use of online social network sites: a longitudinal analysis. *J Appl Dev Psychol* 29:434–445. Social Networking on the Internet
- Tan SB, Chiu-Shee C, Duarte F (2022) From SARS to COVID-19: digital infrastructures of surveillance and segregation in exceptional times. *Cities* 120:103486
- Thatcher J, O'Sullivan D, Mahmoudi D (2016) Data colonialism through accumulation by dispossession: new metaphors for daily data. *Environ Plan D Soc Space* 34:990–1006
- Theodorou A, Dignum V (2020) Towards ethical and socio-legal governance in AI. *Nat Mach Intell* 2:10–12
- Wong RY, Madaio MA, Merrill N (2023) Seeing like a toolkit: how toolkits envision the work of AI ethics. *Assoc Comput Mach* 7. <https://doi.org/10.1145/3579621>

Acknowledgements

The authors would like to thank the financial support received from the members of the Senseable City Lab consortium, including Toyota Woven City, the AMS Institute, Dubai Future Foundation, Arnold Ventures, Toyota, UnipolTech, Consiglio per la Ricerca in Agricoltura e l'Analisi dell'Economia Agraria, Tele2, Volkswagen Group America, FAE Technology, Fédération Internationale de l'Automobile, Amsterdam, and Rio de Janeiro.

Author contributions

Martina Mazzarello, Fábio Duarte, Simone Mora: conceptualization, methodology, investigation, resources, data curation, writing—original draft, writing—review & editing, project administration, funding acquisition. Cate Heine: methodology, formal analysis, data curation, validation, writing—review & editing. Carlo Ratti: supervision, funding acquisition.

Competing interests

One of the authors is part of the editorial board of Humanities and Social Sciences Communications.

Ethical approval

This study received exempt status by the Institutional Review Board (IRB) at the Massachusetts Institute of Technology (decision 4268). The study did not collect any personal

information from participants. The exempt status was granted on April 18th, 2024, and it is valid until August 31st, 2026. This research was conducted in accordance with Belmont Principles.

Informed consent

Informed consent was acquired verbally by all participants during in-person workshops with the exception of workshops held in Norway where written consent was acquired. The digital version of Data Slots included statements on data anonymization and data usage.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1057/s41599-025-04776-1>.

Correspondence and requests for materials should be addressed to Fábio Duarte or Simone Mora.

Reprints and permission information is available at <http://www.nature.com/reprints>

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



Open Access This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

© The Author(s) 2025