

PSI (Ψ):
a Private data Sharing Interface*
(WORKING PAPER)

Marco Gaboardi[†] James Honaker[‡] Gary King[§] Kobbi Nissim[¶]
Jonathan Ullman^{||} Salil Vadhan^{**}

with contributions from

Andreea Antuca, Connor Bain, Victor Balcer, Jessica Bu, Mark Bun, Stephen Chong, Vito D’Orazio,
Anna Gavrilman, Caper Gooden, Paul Handorff, Raquel Hill, Allyson Kaminsky, Murat Kuntarcioglu,
Vishesh Karwa, George Kellaris, Hyun Woo Lim, Nathan Manohar, Dan Muise, Jack Murtagh,
Sofya Raskhodnikova, Ryan Rogers, Or Sheffet, Adam D. Smith, Thomas Steinke, David Xiao,
Haoqing Wang, and Joy Zheng

June 23, 2016

Abstract

We provide an overview of PSI (“a Private data Sharing Interface”), a system we are developing to enable researchers in the social sciences and other fields to share and explore privacy-sensitive datasets with the strong privacy protections of differential privacy.

*This work is part of the “Privacy Tools for Sharing Research Data” project at Harvard, supported by NSF grant CNS-1237235 and a grant from the Sloan Foundation. This is a working paper describing a vision for work that is still in progress, and is therefore authored by the leadership of the efforts. Future and accompanying publications that emphasize specific technical contributions will be authored by team members responsible for those contributions.

[†]Department of Computer Science and Engineering, University at Buffalo, SUNY. Work done in part while at the University of Dundee, UK and visiting the Center for Research on Computation & Society, John A. Paulson School of Engineering & Applied Sciences, Harvard University. gaboardi@buffalo.edu.

[‡]james@hona.kr; <http://hona.kr>

[§]Albert J. Weatherhead III University Professor, Harvard University, Institute for Quantitative Social Science. king@harvard.edu; <http://GaryKing.org>

[¶]Department of Computer Science, Ben-Gurion University, *and* Center for Research on Computation & Society, John A. Paulson School of Engineering & Applied Sciences, Harvard University. kobbi@seas.harvard.edu.

^{||}College of Computer and Information Sciences, Northeastern University. Work done in part while affiliated with the Center for Research on Computation & Society, John A. Paulson School of Engineering & Applied Sciences, Harvard University. jullman@ccs.neu.edu

^{**}Center for Research on Computation & Society, John A. Paulson School of Engineering & Applied Sciences, Harvard University. Work done in part while visiting the Shing-Tung Yau Center and the Department of Applied Mathematics at National Chiao-Tung University in Taiwan. Also supported by a Simons Investigator Award. salil@seas.harvard.edu.

1 The Problem

Researchers in all experimental and empirical fields are increasingly expected to widely share the data behind their published research, to enable other researchers to verify, replicate, and extend their work. Indeed, data-sharing is now often mandated by funding agencies [40, 39, 20] and journals [34, 21, 48]. To meet this need, a variety of open data repositories have been developed to make data-sharing easier and more permanent. Examples include NIH’s database of Genotypes and Phenotypes (dbGaP) [46], and repositories that use the open-source **Dataverse** platform [9, 29] (including Harvard’s Dataverse repository, which has, under some measures, the largest repository of social science datasets in the world).

However, many of the datasets in the social and health sciences contain sensitive personal information about human subjects, and it is increasingly recognized that traditional approaches such as stripping “personally identifying information” are ineffective at protecting privacy, especially if done by a lay researcher with no expertise in deidentification. This leads to two problems, one for privacy and one for utility:

1. There are numerous data sets, such as surveys, that have been "deidentified" via traditional means and increasingly being deposited in publicly accessible data repositories. As the literature has repeatedly shown, it is likely that many subjects in these surveys can be reidentified by attackers with a moderate amount of background information, and thus their privacy may not be sufficiently well-protected.
2. There are numerous other data sets that researchers do not make available at all, or only with highly restrictive and time-consuming provisions, which can include a review by the original data depositor—who may no longer be accessible—and/or an Institutional Review Board (IRB), and a lengthy negotiation between institutions on the terms of use.

Thus, an important problem is to develop and deploy methods that can be used to offer greater privacy protections for datasets of the first type, ideally at little or no cost in utility¹, and enable the safe sharing of datasets of the second type.

Differential privacy [15] offers an attractive approach to addressing this problem, enabling us to provide wide access to statistical information about a dataset without worries of individual-level information being leaked inadvertently or due to an adversarial attack. There is now both a rich theoretical literature on differential privacy and numerous efforts to bring differential privacy closer to practice. However, none of the past work simultaneously meets all of our desiderata for such a system:

- **Accessibility by non-experts:** researchers in the social sciences should be able to use the system to share and explore data with no involvement from experts in data privacy, computer science, or statistics.
- **Generality:** the system should be applicable and effective on a wide variety of heterogeneous datasets hosted in a repository such as the Harvard Dataverse.
- **Workflow-compatibility:** the system should fit naturally in the workflow of its users (e.g. researchers in the social sciences), and be positioned to offer clear benefits (e.g. more access to sensitive data or less risk of an embarrassing privacy violation) rather than being an impediment.

¹Even traditional de-identification techniques have been found to have a significant negative impact on utility [10].

2 Differential Privacy

Differential privacy is a formal mathematical framework for measuring the privacy guarantees provided by statistical computations. Consider an algorithm M that takes a dataset x as input and performs a randomized computation to produce an output y . Informally, differential privacy requires that if we change any one individual's data in x , then the distribution of y does not change much. Intuitively, this means that each individual's data is hidden from an adversary that views the output y .

To make this intuition precise, we need to define what we mean by “one individual's data,” and provide a measure of how much the distribution of y is allowed to change. For the former, a typical choice is to consider datasets x that consist of n records, where we think of each record as consisting of one individual's data, and the sample size n is public (not sensitive information). We call two datasets x and x' *neighbors* if they agree in all but one record (i.e. x' is obtained from x by changing one individual's data). Then the formal definition of differential privacy is as follows:

Definition 2.1 ([15, 14]) For parameters $\epsilon \geq 0$ and $\delta \in [0, 1]$, a randomized algorithm M is (ϵ, δ) -differentially private if for every two neighboring datasets x, x' and every set S of outputs,

$$\Pr[M(x) \in S] \leq e^\epsilon \cdot \Pr[M(x') \in S] + \delta,$$

where the probabilities are taken over the randomization of the algorithm M .

The level of privacy protection is governed by the two *privacy parameters* ϵ and δ ; the smaller they are, the closer the distributions of $M(x)$ and $M(x')$ are, and hence the greater the level of privacy. Typically, ϵ is taken to be a small constant such as .1, whereas δ is taken to be very small, like 2^{-30} .

The way that differentially private algorithms for statistical analysis are often designed are by carefully introducing a small amount of random noise into non-private algorithms for the same analyses. The more noise that is introduced, the greater the level of privacy protection (i.e. a smaller ϵ and/or δ). However, less noise produces a more accurate and useful analysis. Thus differentially private algorithms offer a privacy-utility tradeoff.

By now, there is a large literature giving differentially private algorithms for a wide variety of data analysis tasks. Often, these algorithms are accompanied by a theoretical analysis showing that their performance converges to that of the non-private algorithm as the sample size n tends to infinity. However, such asymptotic performance guarantees do not necessarily translate to good performance at a specific finite sample size, and thus a great deal of work remains to be done to engineer differentially private algorithms to be useful in practice.

In addition, one typically does not want to run just one analysis on a dataset, but rather a large collection of analyses. Fortunately, differential privacy satisfies a variety of *composition theorems* showing that the privacy protection degrades gracefully when we run many differentially private algorithms. For example:

Lemma 2.2 (basic composition [15, 14]) Let M_1, \dots, M_k be randomized algorithms where M_i is (ϵ_i, δ_i) differentially private for $i = 1, \dots, k$. Then the algorithm $M(x) = (M_1(x), \dots, M_k(x))$ that runs each of the M_i 's using independent coin tosses is $(\sum_i \epsilon_i, \sum_i \delta_i)$ differentially private.

Since the δ parameter in differential privacy is generally taken to be negligibly small, the effect of this summing is generally insignificant, so we will focus on the ϵ parameter. If we want to achieve a global, overall level of privacy protection given by ϵ_g , we can think of ϵ_g as a “privacy budget” to be spent on different analyses M_i we want to run. We can spend more of this budget on a specific analysis M_i (i.e. take ϵ_i smaller), but this will consume more of our budget, leaving less for the other analysis if we want to ensure that $\sum_i \epsilon_i \leq \epsilon_g$.

There are better bounds on the composition of differentially private algorithms than the simple summing bound given above [16, 28, 37], but they still have the same budget-like effect—a larger ϵ_i (i.e. higher accuracy, lower privacy) for one computation requires reducing the ϵ_j ’s for other computations in order to maintain the same overall level of privacy.

3 Our Contribution: PSI

In this paper, We provide an overview of PSI (“a Private data Sharing Interface”), a system we are developing to enable researchers in the social sciences and other fields to share and explore privacy-sensitive datasets with the strong privacy protections of differential privacy. It is designed to achieve all of the desiderata mentioned in Section 1 (Accessibility for Non-Experts, Workflow-compatibility, and Generality). Indeed, unique features of PSI include:

- None of its users, including the *data depositors* who have privacy-sensitive data sets they wish to share and the *data analysts* who seek to analyze those datasets, are expected to have expertise in privacy, computer science, or statistics. Nevertheless, PSI enables them to make informed decisions about the appropriate use of differential privacy, the setting of privacy parameters, the partitioning of a privacy budget across different statistics, and the interpretation of errors introduced for privacy.
- It is designed to be integrated with existing and widely used data repository infrastructures, such as the *Dataverse* project [9, 29], as part of a broader collection of mechanisms for the handling of privacy-sensitive data, including an approval process for accessing raw data (e.g. through IRB review), access control, and secure storage. Consequently, PSI can initially be used to *increase* the accessibility of privacy-sensitive data, augmenting rather than replacing current means for accessing such data, thereby lowering the adoption barrier for differential privacy.
- Its initial set of differentially private algorithms were chosen to include statistics that have wide use in the social sciences, and are integrated with existing statistical software designed for lay social science researchers, namely the Zelig [7] package in R and the TwoRavens [26] graphical data exploration interface.

A preliminary prototype of PSI is available at <http://privacytools.seas.harvard.edu/psi>. It does not yet incorporate all of the planned features described in this paper, as a number of them are still under development. The purpose of this paper is to describe the *design* of PSI, and initiate a discussion about the choices made and possible alternatives.

4 Previous work

Most of the previous work to bring differential privacy to practice can be partitioned into the following categories:

- *Programming languages and systems*: here the goal is to make it easier for users to write programs that are guaranteed to be differentially private, either by composition of differentially private building blocks [33, 42, 23], using general frameworks such as “partition-and-aggregate” or “subsample-and-aggregate” [38] to convert non-private programs into differentially private ones [44, 36], or by formal verification from scratch [5]. On one hand, these methods provide much more generality than we seek—our target users are not programmers, and it will already be very useful to provide them with a small, fixed collection of differentially private versions of statistical computations that are common in the social sciences. On the other hand, most of these tools do not provide much guidance for a lay user in deciding how to partition a limited privacy budget among many statistics or analyses he or she may want to run, or how to interpret the noisy results given by a differentially private algorithm.

In contrast to the other tools mentioned above, GUPT [36] does enable a user to specify fine-grained accuracy goals and automatically converts these into privacy budget allocations, in a similar spirit to our privacy budgeting tool (described later). However, GUPT is limited to differentially private programs obtained via the subsample-and-aggregate framework, whereas our tool has no such restriction, and can be extended to include arbitrary differentially private algorithms. Moreover, our tool allows the privacy budget allocation to be interactively adjusted by users, and supports optimal composition theorems for differential privacy [37].

- *Optimization for specific data releases*: there have been several successful applications of differential privacy to very specific and structured sources of data like commuter patterns [31], mobility data [35], client-side software data [19], and genome-wide association studies [8]. Here differential privacy experts carefully optimize the choice of differentially private algorithms and the partitioning of the privacy budget to maximize utility for the particular data source. In the context of a broad data repository in the social or health sciences, the collection of data sources and the structure of the datasets is too heterogenous to allow for such optimization. And it is not scalable to have a differential privacy expert manually involved in each instance of data sharing.

- *Optimization and evaluation of specific algorithms*: there is a vast literature on the design of differentially private algorithms for specific data analysis tasks, including substantial experimental work on comparing and optimizing such algorithms across a wide range of datasets. As an example, the recent work on DPBench [24] provides a thorough comparison of different algorithms and different ways of optimizing them. Such work is complementary to ours. Algorithms that perform well in such evaluation are natural candidates to add to our library of differentially private routines, but such evaluation does not address how to budget the privacy allocated to this one algorithm against many other analyses one might want to run on the same dataset or more generally how to enable lay users to make appropriate use of differential privacy. Moreover, our use case of a general-purpose social science data repository guides the choices of which algorithms to implement, the measures of accuracy, and the methods for evaluation, as discussed in the later sections.

There are also a number of deployed systems that provide query access to sensitive data, using heuristic approaches to protect privacy. These include systems for querying clinical health data [30, 32], education data [3], genomic data [46], and Census data [1]. However, the lack of rigorous privacy guarantees raises a genuine risk, as illustrated by attacks on the Israeli Census query system [49], on genomic data [25, 18] and more generally on releases of aggregate statistics [11, 17]. (Some of the aforementioned systems address this concern by limiting access to a more trusted set of users.)

5 Incentives For Use

Differential privacy has sometimes been critiqued for its cost in utility (coming from the noise introduced in statistics), thus one might wonder what would motivate researchers to use it in place of the current data-sharing ecosystem. We see at least three different scenarios in which differential privacy can provide a clear benefit over current approaches.

- (“DP works great”) In some circumstances, the results of differentially private analyses are virtually indistinguishable from non-private analyses. Currently, this tends to be the case when the number n of samples is large, the data is low-dimensional, and the analyses to be performed are relatively simple and few in number. In such cases, the greater privacy protections of differential privacy come essentially for free. As both theoretical and applied work on differential privacy advances and data gets “bigger” (n gets larger), we can expect an increasingly large set of data-sharing circumstances to fall in this scenario.
- (“Access is wide”) When we wish to make sensitive data available to an extremely wide community (for example, when allowing public access), we should be increasingly concerned about attacks from individuals with malicious intent. Such adversaries can include ones who have extensive knowledge about a particular data subject that can be exploited as background information. Thus, the strong protections of differential privacy, which remain meaningful regardless of an adversary’s background information, are attractive.
- (“Data is currently unavailable”) For data that is currently unavailable except possibly through restrictive and time-consuming provisions, *any* useful statistical information that differential privacy can offer is a benefit to utility, even if it does not fall in the “DP works great” category. In particular, differential privacy can offer the possibility of rough, exploratory analysis to determine whether a dataset is of sufficient interest to go through the process of applying for access to the raw data.

The architecture of PSI is designed to support all three of these scenarios.

In the near term, we expect the third scenario, namely enabling exploratory analysis of data that is currently unavailable, to be the one where PSI is most frequently used. In this scenario, PSI can provide a clear utility benefit, can be applied with the modest sample sizes that are common in social science, and does not require an extensive library of highly optimized and sophisticated differentially private algorithms. However, PSI is extensible to incorporate such a library in the future, and we hope that eventually it will be used more often in the other two scenarios as well, providing high-utility and privacy-protective access to data that is currently shared in a less safe manner.

In the future, another potential incentive for the use of a differentially private data analysis system like PSI is the automatic protection that differential privacy provides against false discovery, allowing analysts to perform adaptive data exploration (without “preregistration”) and still have confidence that the conclusions they draw are statistically valid [13, 6, 43].

We note that sometimes researchers do not wish to share their data, and are only using privacy as an excuse. A system like PSI can help eliminate the excuse. Still, other external incentives may be needed (such as from the research community, funding agencies, or journals) to encourage sharing of data.

On exploratory analysis. Since it is our initial goal for the use of PSI, we elaborate on what we mean by supporting “exploratory data analysis.” This term generally refers to a wide-ranging set of techniques to empirically learn features of data by inspection, and familiarize oneself with the nature of the data, or discover apparent structure in the data [47]. It is inspection and discovery not driven by theory or modeling. In our setting of a social science data repository, we envision at least two uses for exploratory analysis. For lay-users (e.g. members of the general public), exploratory analysis can be a way to satisfy curiosity and discover interesting facts for situations where a statistically rigorous analysis may not be necessary (e.g. for a high-school project). For a social science researcher, the goal of exploratory analysis can be to determine which of the many datasets in the repository are of most interest, so that the researchers only invest their time and effort in applying for raw access to those datasets. Any final analyses they wish to perform and publish could then be done on the raw data, not through the differentially private interface. This more modest near-term goal for PSI compensates for the fact that we cannot perform the kinds of optimizations that might be done if we had a differential privacy expert involved in each instance of data sharing.

6 Actors and Workflow

We have three different kinds of actors in PSI:

- *Data depositors:* These are users that come to deposit their privacy-sensitive dataset in a data repository, and may wish to make differentially private access to their dataset available. Based on the interaction with the data depositor, the system will gather basic information about the dataset (e.g. the types and ranges of the variables), set the overall privacy parameters, select an initial set of differentially private statistics to calculate and release, and determine how the remaining privacy budget will be partitioned among future data analysts.

Data depositors are the ones with the initial ethical and/or legal responsibility for protecting the privacy of their data subjects, and they (or their institutions) may be liable if they willfully violate their obligations. Thus, they can be trusted to follow instructions (if not onerous or confusing) and answer questions truthfully to the best of their knowledge. On the other hand, they cannot be assumed to have expertise in differential privacy, computer science, or statistics, so any questions that involve these areas need to be explained carefully.

- *The data curators:* These are the data-repository managers that maintain the hardware and software on which PSI runs and the accompanying data repository infrastructure (e.g. **Dataverse**) and associated statistical tools (e.g. **Zelig** and **TwoRavens**). They are trusted, and indeed may also have legal obligations to ensure that the repository does not violate the privacy protections it claims to offer through tools such as PSI. Data curators can be assumed to have expertise in IT systems administration and data stewardship [22] and archiving [4], and can be trained to have at least a modest background in statistics and differential privacy. But they are few in number, and cannot be actively involved in most instances of data sharing or data exploration. Thus PSI needs to be sufficiently automated to enable data depositors and data analysts to safely use it on their own.

Data curators would also be responsible for deciding whether to accept new differentially private routines into the library used by PSI and correcting bugs or security flaws found in existing routines. These can be difficult tasks even for experts in differential privacy. Thus, in a future version of the

system, it would be of interest to minimize the amount of trusted code, and have tools to formally verify the remaining components (both original components and later contributions), along the lines of the programming languages tools described in Section 4.

- *Data analysts:* These are users that come to access sensitive datasets in the repository, often with the goal of data exploration as discussed in Section 5. They will have access to all of the differentially private statistics selected by the data depositor, as well as ability to make their own differentially private queries (subject to staying within the overall privacy budget, as discussed more below).

We envision at least two tiers of trust for data analysts. PSI can make access available to a very wide community of analysts (e.g. the general public), in which case the analysts are considered completely *untrusted*. Alternatively (or in addition), we can restrict to a set of analysts that are identifiable (e.g. as registered users of the data repository), with some accountability (e.g. through their verified affiliation with a home institution). Such analysts may be considered as *semi-trusted*, as we can assume that most of them will follow basic terms of use to not abuse the system in certain ways. Specifically, we will assume that semi-trusted users will not collude to compromise privacy, and will not create phony accounts. (This will enable us to provide greater utility for such users, as discussed in Section 8.)

7 Pedagogical Materials

In order to enable PSI to be used by empirical researchers without expertise in privacy, computer science, or statistics, we have prepared pedagogical materials explaining differential privacy in an intuitive but accurate manner, with a minimum of technical terminology and notation. These materials are meant to be sufficient for data depositors and data analysts to understand and make appropriate choices in using PSI, such as those described in the forthcoming sections. Data depositors require more background material than data analysts, as the former are concerned with the privacy protections afforded to their data subjects, whereas the latter only need to understand the impact of the system on their analyses (namely, that results will be less accurate or statistically significant than would be obtained on the raw data, and that there is a limited “budget” of queries that they can perform).

Relevant extracts of the pedagogical materials will be offered to users of PSI at each decision point, and can also be included when describing data-sharing plans to Institutional Review Boards (IRBs). In addition, members of our team have started to develop rigorous arguments showing that differential privacy should be deemed to satisfy certain legal obligations of privacy protection, which can also be used to reassure data depositors, data curators, and IRBs that differential privacy is a sufficiently strong form of protection.

As discussed in Section 6, we assume that data curators have expertise in IT systems administration and data stewardship, and at least a modest background in statistics and differential privacy. Thus, they do not need any specialized pedagogical materials other than a thorough documentation of the system.

8 Privacy Budget Management

One of the challenges in enabling non-experts to use differential privacy is that it can be difficult to understand the implications of different selections of the privacy parameters (namely ϵ and δ), both in terms of privacy and utility, especially when these need to be distributed over many different statistics to be computed. To address this issue, PSI is designed to expose these implications to the user, in easy-to-understand terms, and is accompanied by a variety of simple explanations of differential privacy and its parameters that are shown to the user at relevant times.

Global privacy parameters.

The data depositor, who carries the initial responsibility for protecting the privacy of her data subjects, is charged with setting the overall (“global”) privacy parameters ϵ_g, δ_g for her dataset. To enable this choice, we provide intuitive (but accurate!) explanations of the meaning of each of these privacy parameters, and give recommended settings based on the level of sensitivity of a dataset (e.g. corresponding to Harvard’s data security levels [2] or the similar categories in the DataTags system that integrates with PSI [45]). δ_g is easily explained as the probability of arbitrary leakage of information, like the probability of an adversary breaking an encryption scheme, and thus should be set to be extremely small, like 2^{-30} . For the main privacy parameter, ϵ_g , we explain it with a table comparing an adversary’s posterior belief that a data subject has a sensitive trait to the posterior belief had the subject opted out of the study. PSI also confirms with the data depositor that each individual subject’s data corresponds to one row of the uploaded dataset (so that the per-row protections of differential privacy translate to per-subject protections). The data depositor is also asked whether the dataset is a random sample from a larger population, and whether the choice of this sample has been kept confidential. If so, the “secrecy of the sample” lemma in differential privacy allows for an effective savings in the privacy parameters corresponding to the ratio of sizes between the dataset and the larger population.² This means that correspondingly greater utility can be provided for the same level of privacy protection. (To account for the fact that, in practice, population samples are typically not perfectly random, the depositor is instructed to conservatively estimate the overall population size.)

Budgeting among different statistics.

Even after the global privacy parameters are determined, there is still the challenge of how they should be distributed among the different statistics to be computed. That is, for each statistic to be computed, we need to select privacy parameters (i.e. set ϵ_i and δ_i for statistic i) and then apply composition theorems to ensure that globally, we achieve (ϵ_g, δ_g) differential privacy. To ensure that we get the most utility out of the global privacy budget, we use the “approximate optimal composition theorem” of [37], which in fact was developed for the purpose of our privacy budget tool.

This leaves the question of how a user should select individual privacy parameters ϵ_i (and δ_i). The larger the value of ϵ_i is taken, the more utility we obtain from the i ’th statistic, but this leaves less of the global privacy budget remaining for the other statistics. Since some statistics a user is computing may be more important than others, and different differentially private algorithms

²<https://adamsmith.wordpress.com/2009/09/02/sample-secrecy/>.

have different privacy-utility tradeoffs, the “best” use of the privacy budget is likely to involve a non-uniform distribution of the ϵ_i ’s.

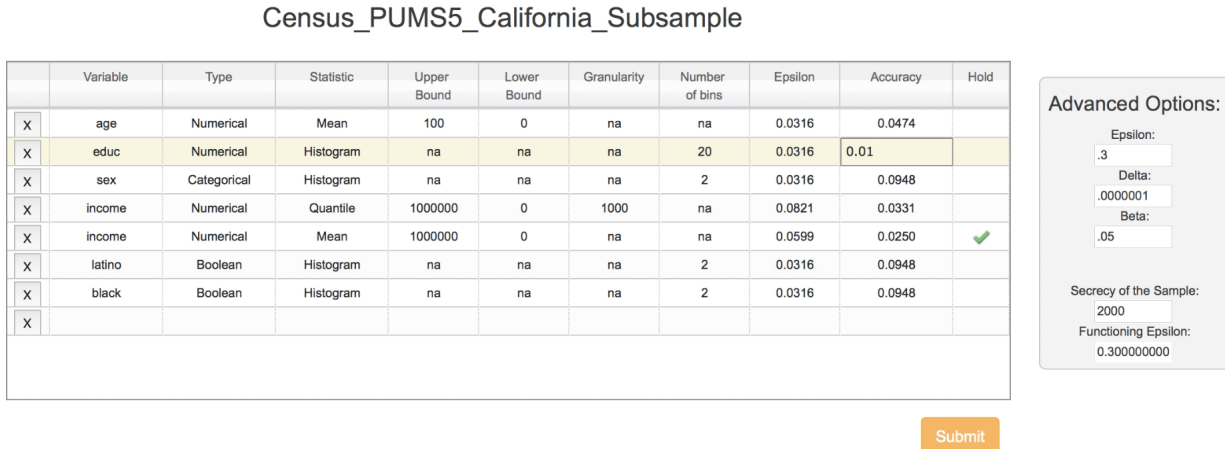


Figure 1: PSI privacy budgeting interface

To enable a user to determine this partition, we have developed a privacy budgeting tool that exposes the privacy-accuracy tradeoff to the user. (See Figure 1.) That is, rather than selecting the individual privacy parameters ϵ_i , the user can select and modify the “accuracy” that will be obtained for different selected statistics (presented as, for example, the size of 95% confidence intervals; see further discussion in the next section). For each differentially private algorithm in PSI, there are accompanying functions that translate between the privacy parameters and a measure of accuracy (also depending on other metadata, such as the range of variables involved and the dataset size n). These functions are used by the privacy budgeting tool to translate the accuracy bounds into individual privacy parameters and ensure that the global privacy parameters are not exceeded.

Budgeting among different actors.

Recall that the selection of differentially private statistics to be computed is done both by the data depositor, who selects an initial set of statistics that will be shared among all analysts that access the dataset, and by individual data analysts, who may be carrying out novel explorations of their own conception. The privacy budgeting tool described above is designed to support both types of actors (with slightly different settings for each to reflect their different roles and level of trustworthiness). The data depositor is tasked with deciding how much of the global privacy budget ϵ_g to reserve for future data analysts. For example, if the data depositor uses up ϵ_d units of privacy for the statistics she chooses to release, then at least $\epsilon_a = \epsilon_g - \epsilon_d$ units of privacy will be left for the future analysts. (ϵ_a might actually be larger, since composition theorems for differential privacy can in some cases give better bounds than simply summing the privacy parameters.)

In the case of semi-trusted data analysts (who we assume will not collude, as discussed in Section 6), PSI provides *each* analyst a *per-user* privacy budget of ϵ_a . In the case of completely untrusted analysts, we share ϵ_a among all future analysts. This model is more conservative with respect to privacy protection, and thus may be appropriate when analysts do not have the sufficient accountability or the data is highly sensitive (e.g. with life-or-death or criminal implications). The

downside of the more conservative model is that it is vulnerable to a denial-of-service attack, where the first few data analysts, intentionally or inadvertently, deplete the entire privacy budget, leaving future analysts unable to make any queries. This can be partly mitigated by rate-limiting the use of the privacy budget and by sharing all statistics computed publicly. It is also possible to reserve part of the privacy budget for untrusted analysts and part for trusted analysts, with each part being treated as described above.

9 Differentially Private Algorithms

Choice of Statistical Procedures

While it is designed to be easily extensible so as to incorporate new algorithms from the rapidly expanding literature, the initial set of differentially private algorithms in PSI were chosen to support the most necessary statistics that are needed to provide immediate utility for social science research and data exploration. Specifically, we include:

- Univariate descriptive statistics, such as means, quantiles, histograms, and approximate cumulative distribution functions. From some of these, post-processing can also provide additional descriptive statistics at no additional privacy cost.
- Basic statistical estimators, for inference about the population from which a dataset was sampled. We have selected some of the most widely used statistical inference procedures in social science, such as difference-of-means testing for causal inference, hypothesis tests for the independence of categorical variables, and low-dimensional covariance matrices from which can be extracted correlations, least-squares regressors, and principal components.
- Transformations for creating new features (variables) out of combinations of already existing ones. These allow the previously described procedures to be leveraged to do more sophisticated computations on a broader range of questions.³

These choices are also motivated in part by the data exploration tools that PSI integrates with, and which we expect our data analysts to use. In particular, the `TwoRavens` graphical data exploration tool (<http://2ra.vn>) provides descriptive statistics for each variable in a dataset, as well as graphical illustrations of its empirical distribution (e.g. a histogram or a probability density function) [12]. PSI replaces these with the differentially private descriptive statistics it computes. `TwoRavens`, together with the R package `Zelig` that it integrates with (<http://zeligproject.org>), also provides a unified, user-friendly interface for running a wide variety of statistical models. We have chosen to initially implement differentially private versions of statistical inference procedures that are widely used in social science and where the differentially private algorithms are sufficiently simple and well-understand to give good performance at finite sample sizes.

³For example, the (empirical) covariance between two attributes can be estimated by estimating the mean of a new attribute that is the product of the two original attributes (as well as the means of the original attributes), or the mean of a variable in a subpopulation can be computed from the mean of the product of that variable with an binary indicator for the subpopulation of interest, and the mean of the indicator.

Measuring Accuracy

The choice of accuracy measure, and how to represent it to the users, is important both in the privacy budgeting tool as well as for data exploration by data analysts, who need to know how to interpret the noisy statistics provided by differential privacy.

For descriptive statistics, we have determined that 95% confidence intervals are the simplest and most intuitive way to represent the noise introduced by differential privacy. For many of the basic differentially private algorithms for descriptive statistics (such as the Laplace mechanism), a theoretical worst-case analysis is also indicative of typical performance, so we can use this to calculate the a priori privacy-accuracy translation needed in the privacy budgeting tool.

For statistical inference procedures, the accuracy (e.g. size of a confidence interval obtained) is necessarily data-dependent, even without privacy. (For example, using a t -test for mean estimation gives a confidence interval of size that depends on the empirical variance of the data.) When incorporating such methods, PSI uses conservative confidence intervals (or p -values), meaning that it ensures that the differentially private confidence interval includes the true value with probability *at least* .95. Intuitively, we account for the noise introduced by differential privacy by making the confidence intervals larger — this ensures that analysts do not draw incorrect conclusions from the differentially private statistics (but more analyses may come out inconclusive, as we explain to users of the system). And to provide the *a priori* accuracy bounds needed by the privacy budgeting tool, we use “rules of thumb” based on experimental evaluation: given n , ϵ , the number of variables, etc.,

10 Architecture

The differentially private algorithms, the accuracy estimates, and the budgeting coordinated by our composition theorem, each discussed in the previous sections, are all implemented in the R programming language, which is widely used in the statistics and quantitative social science communities [41]. We expect to distribute all our routines as an R package for easy reuse within the R environment (independently of `Dataverse` and `TwoRavens`). We run the R code on a remote Apache server, as an `rApache` application [27].

At the time of budgeting by the data depositor, the thin HTML interface uses a simple Javascript interactive table, as for example in Figure 1, which has no direct access to either the data or computations on the data; whenever the table requires a new computation,⁴ it copies the contents of the current table⁵ to a remote application that can calculate composition of the privacy budget and the associated accuracies the division of privacy budget would provide. This remote process then recomputes and returns an updated version of the entire table. The frontend interface then rewrites the table with these newly provided values, and waits for more additions from the user until another round of computation is required. The backend composition process is memoryless, in the sense that no past version of the table persists or is stored, but every request of the backend begins an entirely new set of budgeting and accuracy computations. For this reason, the connection between the frontend and backend does not have to be persistent.

When the depositor has finalized the list of statistics she wishes to make available, together with their respective privacy parameters, the table is then submitted to another separate remote

⁴As when a row is completed, or deleted, or when an accuracy value, or any parameter listed in the advanced options is edited.

⁵As a JSON file, by means of an HTTP POST to the composition application running remotely as `rApache` process.

release application that then computes all the differentially private statistics requested.⁶ This is the only process that has access to the raw data which sits in secure storage in the Dataverse repository. The differentially private statistics that are generated are released in a file of meta-data associated with the securely archived data. Everything in this meta-data file can be made available for public browsing. The application that draws the differentially private releases does not reply to the depositor interface. An architecture diagram is shown in figure 2, which shows the directions of communication between every piece of the system and one can trace out from this that any path from the raw data to any data analyst (or even the data depositor), has to pass through the differentially private channel from this application to the writing of the public meta-data file.

The architecture for handling interactive queries by data analysts is more involved, as now the differentially private release tool is no longer being used only by trusted data depositors. We will describe this in more detail in a future version of the paper.

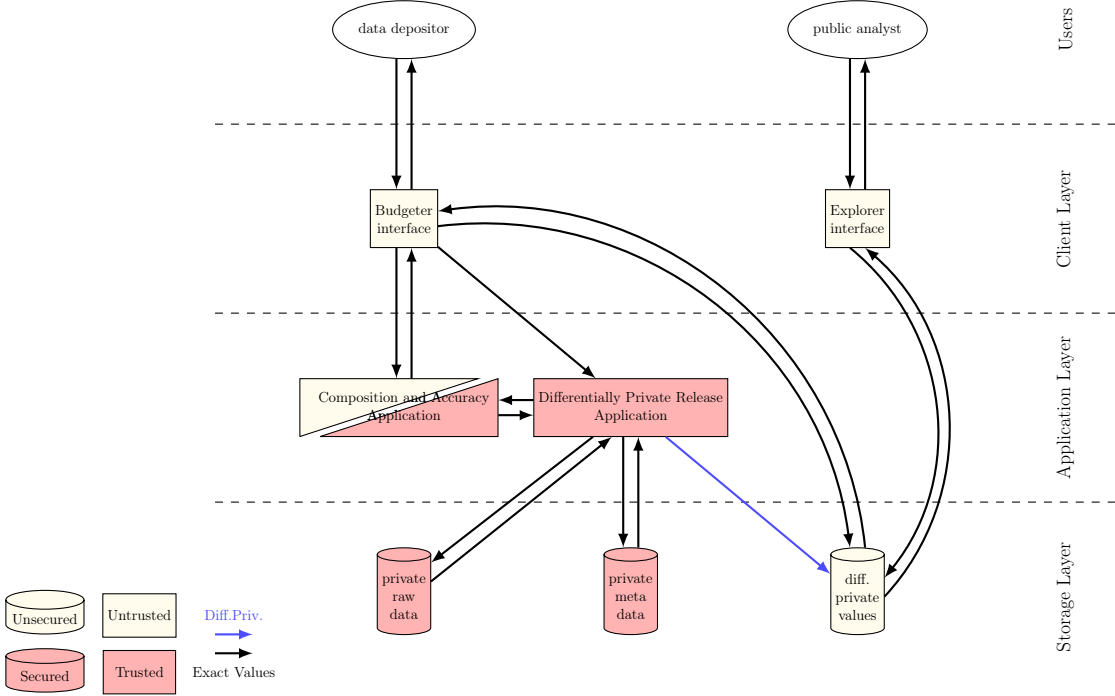


Figure 2: *Architecture diagram (for statistics selected by data depositor).*

11 Security

The initial prototypes of the PSI do not address many of the security and side-channel issues that have been raised in the literature about implementations of differential privacy [23, 36]. We feel that a higher priority is evaluating whether the design of PSI is useful for its potential user community, and if the answer is positive, security issues can be addressed in a future version, before it is used to handle highly sensitive data.

⁶This release tools checks the composition of the request with a trusted version of the composition application, which means that code to this point does not have to be trusted, so long as the global ϵ can be verified.

12 Empirical Evaluation

All of the differentially private algorithms implemented in PSI have been experimentally evaluated using a combination of real and synthetic data. For the future, we plan several forms of evaluation that are more closely tied to our use case of a social science repository.

Replications

The existing (non-sensitive) datasets in the Harvard **Dataverse** repository are also linked to numerous published research papers that include analysis of the dataset. This provides us an opportunity to evaluate the usefulness of PSI for social science research by attempting replication of these published analyses using only our differentially private interface. We can see if the differentially private results provide us with similar conclusions as the published results. This sets a higher bar for PSI than our actual initial goal, which is to support data exploration (for determining whether one should apply for access to raw data). On the other hand, open-ended data exploration may involve many more queries, stretching the privacy budget more thinly and resulting in less accuracy.

Utility for social science users

A different way we intend to evaluate PSI is by means of user experiments with the kinds of social scientists who would actually be using the system. Such experiments would test whether the users are able to appropriately set privacy parameters and interpret the results given by differential privacy, whether the system is sufficiently easy and efficient to use for practicing social scientists, and whether the system satisfies our goals for data exploration. On the latter point, we want to know whether users make the same decisions about which datasets are of most interest for investigating a particular research question as they would if they could view the raw data.

References

- [1] American factfinder. <http://factfinder.census.gov/>.
- [2] Harvard university data classification table. <http://policy.security.harvard.edu/>.
- [3] Naep tools on the web. Technical Report NCES 2011-460, U.S. Department of Education, 2011. <https://nces.ed.gov/nationsreportcard/about/naeptools.aspx>.
- [4] Micah Altman, Margaret O. Adams, Jonathan Crabtree, Darrell Donakowski, Marc Maynard, Amy Pienta, and Copeland H. Young. Digital preservation through archival collaboration: The data preservation alliance for the social sciences. *The American Archivist*, 72(1):170–184, 2009.
- [5] Gilles Barthe, Boris Köpf, Federico Olmedo, and Santiago Zanella Béguelin. Probabilistic relational reasoning for differential privacy. In John Field and Michael Hicks, editors, *Proceedings of the 39th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2012, Philadelphia, Pennsylvania, USA, January 22-28, 2012*, pages 97–110. ACM, 2012.
- [6] Raef Bassily, Kobbi Nissim, Adam Smith, Thomas Steinke, Uri Stemmer, and Jonathan Ullman. Algorithmic stability for adaptive data analysis. In *48th Annual Symposium on the*

- Theory of Computing (STOC'16)*, June 2016. To appear. Preliminary version available at <http://arxiv.org/abs/1511.02513>.
- [7] Christine Choirat, James Honaker, Kosuke Imai, Gary King, and Olivia Lau. Zelig: Everyone's statistical software (version 5), 2015.
 - [8] A community assessment of privacy preserving techniques for human genomes. Xiaoqian jiang and yongan zhao and xiaofeng wang and bradley malin and shuang wang and lucila ohno-machado and haixu tang. *BMC Medical Informatics and Decision Making*, 14(Suppl 1)(S1), 2014.
 - [9] Mercè Crosas. The dataverse network: An open-source application for sharing, discovering and preserving data. *D-Lib Magazine*, 17:1–2, 2011. doi:1045/january2011-crosas.
 - [10] Jon P. Daries, Justin Reich, Jim Waldo, Elise M. Young, Jonathan Whittinghill, Andrew Dean Ho, Daniel Thomas Seaton, and Isaac Chuang. Privacy, anonymity, and big data in the social sciences. *Communications of the ACM*, 57(9):56–63, September 2014.
 - [11] Irit Dinur and Kobbi Nissim. Revealing information while preserving privacy. In *PODS*, pages 202–210, 2003.
 - [12] Vito D'Orazio and James Honaker. *A User Guide to TwoRavens: An overview of features and capabilities*, 2016.
 - [13] Cynthia Dwork, Vitaly Feldman, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Aaron Roth. The reusable holdout: Preserving validity in adaptive data analysis. *Science*, 349(6248):636–638, 2015.
 - [14] Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. Our data, ourselves: Privacy via distributed noise generation. In *EUROCRYPT*, pages 486–503, 2006.
 - [15] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *3rd Theory of Crypt. Conf.*, pages 265–284, 2006.
 - [16] Cynthia Dwork, Guy N. Rothblum, and Salil P. Vadhan. Boosting and differential privacy. In *FOCS*, pages 51–60, 2010.
 - [17] Cynthia Dwork, Adam D. Smith, Thomas Steinke, and Jonathan Ullman. Hiding in plain sight: A survey of attacks on private data. Manuscript, April 2016.
 - [18] Yaniv Erlich and Arvind Narayanan. Routes for breaching and protecting genetic privacy. *Nature Reviews Genetics*, 15(6):409–421, 2014.
 - [19] Úlfar Erlingsson, Vasyl Pihur, and Aleksandra Korolova. RAPPOR: randomized aggregatable privacy-preserving ordinal response. In Gail-Joon Ahn, Moti Yung, and Ninghui Li, editors, *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security, Scottsdale, AZ, USA, November 3-7, 2014*, pages 1054–1067. ACM, 2014.
 - [20] National Science Foundation. Award & administration guide (AAG) chapter vi.d.4, 2014.

- [21] Sergiu Gherghina and Alexia Katsanidou. Data availability in political science journals. *European Political Science*, 12(3):333–349, Sep 2013.
- [22] Alyssa Goodman, Alberto Pepe, Alexander W. Blocker, Christine L. Borgman, Kyle Cranmer, Merce Crosas, Rosanne Di Stefano, Yolanda Gil, Paul Groth, Margaret Hedstrom, David W. Hogg, Vinay Kashyap, Ashish Mahabal, Aneta Siemiginowska, and Aleksandra Slavkovic. Ten simple rules for the care and feeding of scientific data. *PLoS Comput Biol*, 10(4):1–5, 04 2014.
- [23] Andreas Haeberlen, Benjamin C. Pierce, and Arjun Narayan. Differential privacy under fire. In *Proceedings of the 20th USENIX Security Symposium*, August 2011.
- [24] Michael Hay, Ashwin Machanavajjhala, Gerome Miklau, Yan Chen, and Dan Zhang. Principled evaluation of differentially private algorithms using dpbench. In *SIGMOD*, 2016.
- [25] Nils Homer, Szabolcs Szelinger, Margot Redman, David Duggan, Waibhav Tembe, Jill Muehling, John V Pearson, Dietrich A Stephan, Stanley F Nelson, and David W Craig. Resolving individuals contributing trace amounts of dna to highly complex mixtures using high-density snp genotyping microarrays. *PLoS genetics*, 4(8):e1000167, 2008.
- [26] James Honaker and Vito D’Orazio. Statistical modeling by gesture: A graphical, browser-based statistical interface for data repositories. In *Extended Proceedings of ACM Hypertext 2014*. ACM, 2014.
- [27] Jeffrey Horner. *rApache: Web application development with R and Apache.*, 2013.
- [28] Peter Kairouz, Sewoong Oh, and Pramod Viswanath. The composition theorem for differential privacy. In *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6-11 July 2015*, pages 1376–1385, 2015.
- [29] Gary King. An introduction to the dataverse network as an infrastructure for data sharing. *Sociological Methods and Research*, 36:173–199, 2007.
- [30] H. J. Lowe, T. A. Ferris, P. M. Hernandez, and S. C. Weber. STRIDE ?- an integrated standards-based translational research informatics platform. *AMIA Annual Symposium Proceedings*, pages 391–395, 2009.
- [31] Ashwin Machanavajjhala, Daniel Kifer, John M. Abowd, Johannes Gehrke, and Lars Vilhuber. Privacy: Theory meets practice on the map. In *Proceedings of the 24th International Conference on Data Engineering, ICDE 2008, April 7-12, 2008, Cancún, México*, pages 277–286, 2008.
- [32] Andrew J. McMurry, Shawn N. Murphy, Douglas MacFadden, Griffin Weber, William W. Simons, John Orechia, Jonathan Bickel, Nich Wattanasin, Clint Gilbert, Philip Trevvett, Susanne Churchill, and Isaac S. Kohane. SHRINE: enabling nationally scalable multi-site disease studies. *PLoS ONE*, 8(3), 2013.
- [33] Frank McSherry. Privacy integrated queries: an extensible platform for privacy-preserving data analysis. In *Proceedings of the ACM SIGMOD International Conference on Management of Data, SIGMOD 2009, Providence, Rhode Island, USA, June 29 - July 2, 2009*, pages 19–30, 2009.

- [34] R. Michael Alvarez and Jonathan N. Katz. Editors' note. *Political Analysis*, 24(2):131, 2016.
- [35] Darakhshan J. Mir, Sibren Isaacman, Ramón Cáceres, Margaret Martonosi, and Rebecca N. Wright. DP-WHERE: differentially private modeling of human mobility. In Xiaohua Hu, Tsau Young Lin, Vijay Raghavan, Benjamin W. Wah, Ricardo A. Baeza-Yates, Geoffrey Fox, Cyrus Shahabi, Matthew Smith, Qiang Yang, Rayid Ghani, Wei Fan, Ronny Lempel, and Raghunath Nambiar, editors, *Proceedings of the 2013 IEEE International Conference on Big Data, 6-9 October 2013, Santa Clara, CA, USA*, pages 580–588. IEEE, 2013.
- [36] Prashanth Mohan, Abhradeep Thakurta, Elaine Shi, Dawn Song, and David E. Culler. GUPT: privacy preserving data analysis made easy. In K. Selçuk Candan, Yi Chen, Richard T. Snodgrass, Luis Gravano, and Ariel Fuxman, editors, *Proceedings of the ACM SIGMOD International Conference on Management of Data, SIGMOD 2012, Scottsdale, AZ, USA, May 20-24, 2012*, pages 349–360. ACM, 2012.
- [37] Jack Murtagh and Salil P. Vadhan. The complexity of computing the optimal composition of differential privacy. In *Theory of Cryptography - 13th International Conference, TCC 2016-A, Tel Aviv, Israel, January 10-13, 2016, Proceedings, Part I*, pages 157–175, 2016.
- [38] Kobbi Nissim, Sofya Raskhodnikova, and Adam Smith. Smooth sensitivity and sampling in private data analysis. In *Proc. 39th STOC'07 ACM*, 2007.
- [39] National Institute of Health. NIH notice NOT-OD-03-032: Final NIH statement on sharing research data., 2003.
- [40] Office of Management and Budget. OMB memorandum M-13-13: Open data policy – managing data as an asset., 2013.
- [41] R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2016.
- [42] Jason Reed and Benjamin C. Pierce. Distance makes the types grow stronger: A calculus for differential privacy. In *ICFP*, 2010.
- [43] Ryan M. Rogers, Aaron Roth, Adam Smith, and Om Thakkar. Max-information, differential privacy, and post-selection hypothesis testing. *CoRR*, abs/1604.03924, 2016.
- [44] Indrajit Roy, Srinath T. V. Setty, Ann Kilzer, Vitaly Shmatikov, and Emmett Witchel. Airavat: Security and privacy for mapreduce. In *Proceedings of the 7th USENIX Symposium on Networked Systems Design and Implementation, NSDI 2010, April 28-30, 2010, San Jose, CA, USA*, pages 297–312. USENIX Association, 2010.
- [45] Latanya Sweeney, Mercè Crosas, and Michael Bar-Sinai. Sharing sensitive data with confidence: The datatags system. *Technology Science*, 2015.
- [46] Kimberly A. Tryka, Luning Hao, Anne Sturcke, Yumi Jin, Zhen Y. Wang, Lora Ziyabari, Moira Lee, Natalia Popova, Nataliya Sharopova, Masato Kimura, and Michael Feolo. Ncbi's database of genotypes and phenotypes: dbgap. *Nucleic Acids Research*, 42(D1):D975–D979, 2014.

- [47] John Wilder Tukey. *Exploratory Data Analysis*. Addison-Wesley, 1977.
- [48] Sven Vlaeminck and Lisa-Kristin Herrmann. Data policies and data archives: A new paradigm for academic publishing in economic sciences? In B. Schmidt and M. Dobрева, editors, *New Avenues for Electronic Publishing in the Age of Infinite Collections and Citizen Science: Scale, Openness and Trust*, pages 145–155.
- [49] Amitai Ziv. Israel’s ‘anonymous’ statistics surveys aren’t so anonymous. *Haaretz*, 7 January 2013.