Quantifying Location Privacy Leakage
from Transaction Prices

Arthur Gervais[1(✉)], Hubert Ritzdorf[1], Mario Lucic[1], Vincent Lenders[2],
and Srdjan Capkun[1]

[1] ETH Zurich, Zurich, Switzerland
`arthur.gervais@inf.ethz.ch`
[2] Armasuisse, Thun, Switzerland

Abstract. Large-scale datasets of consumer behavior might revolution-
ize the way we gain competitive advantages and increase our knowledge in
the respective domains. At the same time, valuable datasets pose poten-
tial privacy risks that are difficult to foresee. In this paper we study the
impact that the prices from consumers' purchase histories have on the
consumers' location privacy. We show that using a small set of low-priced
product prices from the consumers' purchase histories, an adversary can
determine the country, city, and local retail store where the transaction
occurred with high confidence. Our paper demonstrates that even when
the product category, precise time of purchase, and currency are removed
from the consumers' purchase history (e.g., for privacy reasons), infor-
mation about the consumers' location is leaked. The results are based
on three independent datasets containing thousands of low-priced and
frequently-bought consumer products. The results show the existence of
location privacy risks when releasing consumer purchase histories. As
such, the results highlight the need for systems that hide transaction
details in consumer purchase histories.

1 Introduction

Making data publicly available creates unexpected privacy risks. Recent exam-
ples include AOL's release of users' search keywords [30], which has led to the
identification of users and their profiles [1]. Data released by Netflix was de-
anonymized by leveraging IMDB and dates of user ratings [28], showing that
the release of data cannot be analyzed in isolation. The privacy risks of com-
bining different public records have led to several [36] de-anonymization attacks.
Recent studies of anonymized mobility data showed that mobility traces can be
de-anonymized by leveraging a few observations [19]. One source of consumer
information involves their spending patterns. To date however, it was unclear to
what extent consumer prices leak information about the respective purchase.

Consumer purchase histories are typically recorded by store chains with loy-
alty programs and are used to compute consumer spending profiles [6]. Banks,
payment card issuers, and point-of-sale system providers collect this data at dif-
ferent levels of granularity. In a number of scenarios, it might be desirable to
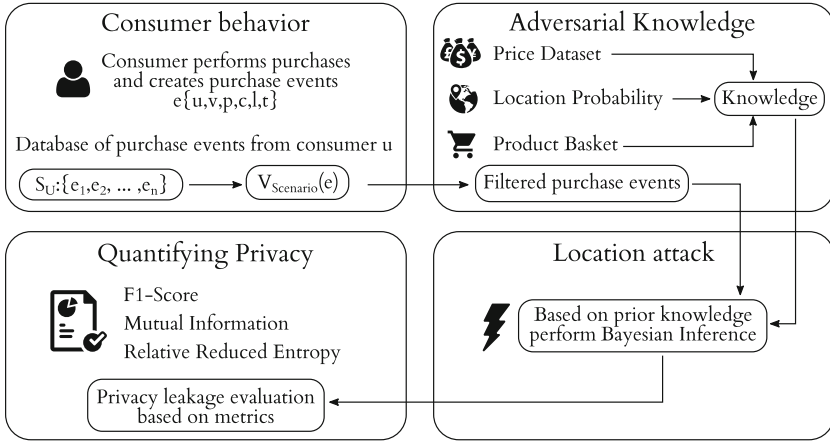
Fig. 1. Framework overview for quantifying location privacy leakage from consumer price datasets.

share this data within different departments of a company, across companies, or with the public [7]. Before disclosure, the data is sanitized so that it does not leak sensitive data, such as personally identifiable information and that it (partially or fully) hides location information. In new digital currency systems such as Bitcoin [33] and Ripple [10], transaction values are stored on a public ledger. Irrespective of whether transaction values are made available so that a system can fulfill its functions or are being disclosed for research purposes, it is important to understand the privacy implications of such disclosures.

In this paper we focus on quantifying location disclosure resulting from the release of prices from consumer's purchase histories. Intuitively, the price distribution for a product differs from country to country, which allows us to identify possible purchase locations. We focus on consumer products which are generally inexpensive ($\leq$ 25 USD) and frequently-bought. More precisely, based on global prices (leveraging the Numbeo dataset [9]), we show that given access to a few consumer prices (and even without the product categories, precise times of purchase or currency), an adversary can determine the country in which the purchase occurred. Similarly, given the country, the city can be determined and within a city (leveraging the Chicago dataset [11]), the *local store* can be identified. We further demonstrate that it is possible to distinguish purchases among store chains (leveraging the Kaggle dataset [7]).

We present a generic framework (cf. Fig. 1) that allows the modeling and quantitative evaluation of location leakage from consumer price datasets. In our framework we model the adversarial knowledge, composed of a public dataset of consumer prices and location-specific information. We assume that the adversary has access to the individual product prices of a purchase (similar to the Kaggle dataset) and a coarse-grained value of the purchase time. In order to make the framework more flexible, our model supports different prior knowledge

scenarios, e.g., the adversary additionally has access to the merchant category (e.g., knowledge that the product was bought in a market or a restaurant) or the product category (e.g., apples). Furthermore, we model the adversarial attack by detailing the corresponding probability functions. In particular, we point out how the adversary leverages multiple product prices in order to increase the probability of identifying the correct location.

Within our framework, we quantify the location privacy of consumer purchases in relation to different dimensions. For example, we measure how well the adversary estimates the location probability of the purchases with the F_1-score [35], capturing the test's accuracy. Furthermore, we use mutual information [18] to quantify the absolute location privacy loss of consumers, based on the considered price dataset. In addition, we capture the relative privacy loss by measuring the reduction in entropy. The proposed metrics are independent of the choice of adversarial strategy and therefore allow us to quantitatively measure the privacy loss induced from any price dataset known to the adversary.

We apply our framework to three real-world datasets: (i) the *Numbeo dataset* [9] contains, after outlier filtering, crowd-sourced real-world consumer prices from 112 countries and 23 US cities for 23 distinct product categories; (ii) the *Chicago dataset* [11] contains 24 million prices for 28 product categories capturing on average of 6304 products sold in Dominick's stores within the Chicago metropolitan area; finally, (iii) the *Kaggle dataset* [7] contains 350 million purchases from 311,541 consumer across 134 store chains.

Our evaluation shows that in order to infer the country based on a vector of purchases, an adversary often needs to observe less than 30 prices. Similarly, after having identified the country of the purchases and given roughly 30 prices, we show that we can reliably predict among 23 major cities within the United States. Finally, when the adversary narrowed down the coarse location, such as the Chicago metropolitan area, we show that based on a regional price dataset, and given a vector of purchases, an adversary can distinguish with high confidence *among local stores* using 100 purchases. For comparison, a weaker adversary with access only to coarse-grained time, i.e., the day of the purchase and price information, requires 50 purchases to identify the country. Furthermore, to establish practical utility of our methodology, we evaluate it on a dataset of purchase records (Kaggle [7]) and show that an adversary requires approximately 250 purchases to distinguish with high confidence among 134 store chains.

The main contributions of this paper are as follows:

– We propose a generic quantitative framework for evaluating attacks against the location privacy of consumer purchases. We validate our framework on three independent price datasets of real-world consumer prices and show that location information can be extracted reliably.
– We introduce three privacy metrics to capture the performance of the adversary in the attack as well as the extent to which location privacy of consumers is reduced when the adversary has access to a specific dataset of purchases.

To the best of our knowledge, this is the first work to infer the location of a purchase based on the price value in consumer purchases. The remainder of

this paper is organized as follows. In Sect. 2, we model purchase history and describe the adversarial model. In Sect. 3, we present the datasets selected for our evaluation in Sect. 4. We survey the related work in Sect. 5 and conclude the paper in Sect. 6.

2   Model

In this section we introduce our system and adversarial model. We present the privacy metrics that quantify the probability of location disclosure based on the assumption that the adversary has access to a part of a consumer's purchase history.

2.1   System Model

A consumer interacts with merchants and performs purchases of one or more products. This interaction leaves a trace of purchase activity as a sequence of *purchase events*. We model each of the consumer's purchase events together with their contextual information as e: {consumer u, value v, product p, product category c, location l, time t}, where v is the price value spent on product p of product category c at location l and time t. In our model, one purchase event is limited to one product, similar to the data contained in the Kaggle dataset. In addition, the price value is given in a global currency, which usually is different from the local currency of the purchase (e.g., the original price is SEK, but recorded in USD). The trace of purchases performed by the target consumer U, given as a series of purchase events, is denoted by S_U:{$e_1, e_2, \ldots, e_n$}. We define the following functions to represent the adversarial knowledge:

LOCATION PROBABILITY: It describes the prior probability of a purchase event taking place in a specific location, e.g., $P(\text{USA})$ is the prior probability with which a random purchase event e has $e.l = \text{USA}$. We define $\mathbb{L}$ as the set of all considered locations.

CATEGORY PROBABILITY: Given location l, $P(c \mid l)$ describes the conditional probability of a purchase event to belong to a certain product category, e.g., $P(\text{Milk} \mid \text{USA})$ is the conditional probability with which a random event e from the USA has $e.c = \text{milk}$. This conditional probability models the product category preferences in a location. We define $\mathbb{C}$ as the set of all considered product categories.

VALUE PROBABILITY: Given location l and product category c, $P(v \mid l, c)$ describes the conditional probability of a purchase event at a given price value. It models the price distributions for different product categories in different locations, e.g., $P(1.5 \mid \text{USA}, \text{Milk})$ is the conditional probability with which milk can be bought in the USA for 1.5 worth of a global currency.

The adversary can now model the spending behavior and identify likely candidate locations. Specifically, the adversary computes the posterior probability

that a single price value v for a product category c originated from a location l. The computation involves the prior and the conditional probabilities described above and the application of Bayes' theorem:

$$P(l \mid c, v) = \frac{P(l) \cdot P(c, v \mid l)}{P(c, v)} \tag{1}$$

In order to infer the location without knowing the product category, the adversary computes the probability that a price value v originates from location l:

$$P(l \mid v) = \frac{P(l) \cdot P(v \mid l)}{P(v)} \tag{2}$$

2.2   Adversarial Model

The adversary's goal is to identify the location of the events in S_U. In this section we present two different adversaries: (1) an adversary with complete knowledge and (2) an adversary with only public knowledge.

Adversary with Complete Knowledge. The ideal adversary represents a strong adversary with complete access to global purchase events. In particular, the adversary has access to the following prior knowledge:

GLOBAL PURCHASE HISTORY: The complete series of purchase events in the history of global purchases[1], denoted by $\mathcal{H}_G$. The adversary computes the posterior probability of a location based on $\mathcal{H}_G$.
HISTORY FOR TARGET CONSUMER: The adversary might have access to prior information about the target consumer's purchase history, denoted by $\mathcal{H}_U$. This could help the adversary to optimize the model for the target consumer[2].

Based on this knowledge, the ideal adversary computes the probabilities in Eqs. 1 and 2.[3]

Adversary with Public Knowledge. Our second adversarial model is a more realistic one, where the adversary only makes use of public information.

POPULATION: Given the population at each location, the adversary estimates the location probability $P(l)$.
PRODUCT BASKET: A product basket indicates which products an average consumer purchases during a year, both in terms of quantity and monetary amount. We leverage the product basket in order to estimate the probability of a product category given the location $(P(c \mid l))$[4].

---

[1] The area of the attacker's interest can be restricted, e.g., when the adversary knows that its victim is somewhere in that restricted area.
[2] For example, by only considering the locations of previous purchases.
[3] The intermediate steps are given in the Appendix A.
[4] We currently use a single product basket for all locations.

PRICE DATASET: For each location and product category combination, a price value distribution D is available, e.g., the Numbeo or the Chicago dataset. The adversary can use the distribution to estimate $P(v \mid l, c)$. We define $D(l, c, v)$ as the number of occurrences of price value v for product category c in location l and $D(l, c)$ as the number of price values for product category c and location l.

Since D might be imperfect, the adversary can have incomplete or incorrect knowledge about the price value probabilities (i.e. unknown or rounded product prices). In this case the adversary should perform additive smoothing, which assigns a small probability α to each event [26]. On the contrary, if the adversary has or assumes complete knowledge of the price value probabilities, additive smoothing is not required.

The adversary with public knowledge computes the following probabilities:

$$P(l) = \frac{\text{Population}(l)}{\sum\limits_{l' \in \mathbb{L}} \text{Population}(l')} \tag{3}$$

$$P(c \mid l) = \frac{\text{Basket}(l, c)}{\sum\limits_{c' \in \mathbb{C}} \text{Basket}(l, c')} \tag{4}$$

$$P(v \mid l, c) = \frac{D(l, c, v) + \alpha}{D(l, c) + \alpha \cdot |S_U|} \tag{5}$$

In order to compute the probabilities defined earlier in Eqs. 1 and 2, the adversary requires access to either $P(l \mid c, v)$ or $P(l \mid v)$. Next, we describe how the adversary computes these probabilities and we define the adversary's knowledge.

2.3   Knowledge Scenarios

As mentioned, the adversary's objective is to identify the location of the events in S_U. The adversary is given a finite set of events S_U on which the attack is executed—the adversary is not allowed to choose or request new purchase events e. We consider an adversary with public knowledge and distinguish among three distinct adversarial knowledge scenarios, each consisting of a subset of the public knowledge. Depending on the knowledge scenario, the adversary might not have access to all information from a purchase event e. Therefore, we define a family of functions $V_{\text{scenario}}(e) = V(e)$ that filter, depending on the given scenario, the public knowledge accessible to the adversary.

PRICE: This scenario corresponds to an adversary that has access to multiple purchase events e, *only* the corresponding *price value* and a notion of the purchase time $e.t$. The adversary is not aware of the product $e.p$ or the product category $e.c$. The precision of the purchase time depends on further specifications of the scenario. More formally, $V_{\text{price}}(e) = \{e.v, e.t\}$. Given the public knowledge modeled by Eqs. 3, 4 and 5, the adversary computes the posterior probability

$P(l \mid v)$ of a price value v from location l. The intermediate steps for computing $P(v \mid l)$ and $P(v)$ are detailed in the Appendix A in Eqs. 10 and 12.

PRICE_MERCHANT: Similar to the former knowledge scenario, the adversary here has access to S_U, a series of multiple purchase events. In this scenario, however, the adversary knows the price value $e.v$ of the event as well as which merchant category m sold the product. Formally, for each purchase event e, $V_{\text{price}}\text{\_merchant}(e) = \{e.v, e.t, m\}$, where $V_{\text{price}}\text{\_merchant}$ requires a function $M(e) = m$. We consider three merchant categories: restaurant, market and local transportation. The $V_{\text{price}}\text{\_merchant}(e)$ function estimates the merchant category m from the product category $e.c$ of the respective event[5]. Analogously, using Eq. 1, the adversary computes the probability of a location, based on the merchant and the price value:

$$P(l \mid m, v) = \frac{P(l) \cdot P(m, v \mid l)}{P(m, v)} \qquad (6)$$

where $P(m, v \mid l)$ is computed as follows:

$$P(m, v \mid l) = \sum_{c \in M^{-1}(m)} P(c, v \mid l) \qquad (7)$$

PRICE_PRODUCT-CATEGORY: This scenario corresponds to the most knowledgeable adversary with public knowledge. Similarly to the former scenarios, the adversary receives multiple purchase events S_U. In addition, the adversary has access to the product category $e.c$ as well as the price value $e.v$. Note that $e.c$ implicitly assumes knowledge of the merchant, resulting in more formally $V_{\text{price}}\text{\_product-category}(e) = \{e.v, e.t, e.c\}$.

Given the public knowledge described in Sect. 2.2, the adversary computes the probability $P(l \mid c, v)$ of a purchase event with product category c and price value v originating in location l. The intermediate steps for computing $P(c, v \mid l)$ and $P(c, v)$ are detailed in the appendix in Eqs. 11 and 13.

In the following section we provide an intuitive perspective on the probabilities $P(l \mid v)$ and $P(l \mid c, v)$.

2.4 Conditional Probability Intuition

$P(l \mid v)$ is the probability of a location, given a price value in a purchase event. An example plot based on our evaluation can be found in Fig. 2. We have chosen the purchase event e with a price value of $e.v = 1$ Euro and estimated the location of the price. The figure shows that the most likely location for 1 Euro is France, closely followed by Germany, Italy and Spain. The plot also shows $P(l \mid c, v)$ for a purchase event with $e.v = 1$ Euro and the product category is milk. The most likely country is again France, followed by Germany and Italy. Surprisingly, China ranks as 5^{th}. This can be explained by the fact that (i) some prices from

---

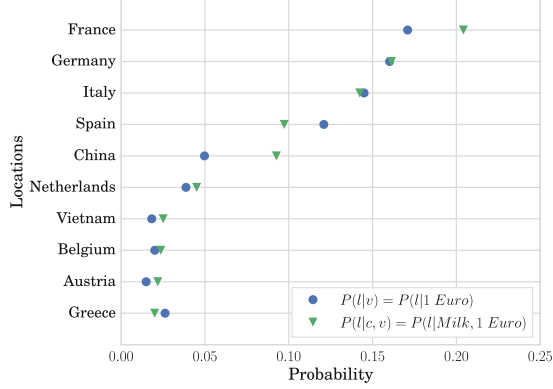[5] In the following we refer to the merchant category as merchant.

Fig. 2. Probability distribution of $P(l \mid v)$ and $P(l \mid c, v)$, given 1 Euro and milk.

China in the dataset were erroneously reported in Euros and (ii) that the location probability $P(l)$ influences the overall outcome, and, since China's population is considerable, there is an increased probability of purchases occurring there. Overall we observed that the probability distribution changes when the product category is known, i.e., France is more likely to have a 1 Euro price for milk, than a 1 Euro price in general.

2.5 Multiple Purchase Events

Up to this point, the analysis has been based on a single purchase event. To naturally combine multiple purchase events, we assume that the purchase events are conditionally independent, given the location l. Therefore, the probability of a location l, given a set of purchase events S_U, is calculated as follows:

$$P(l \mid S_U) = P(l \mid V(e_1), V(e_2), \ldots, V(e_n))$$

$$= \frac{P(l) \cdot \prod_{e \in S_U} P(V(e) \mid l)}{P(V(e_1), \ldots, V(e_n))} \tag{8}$$

The intermediate steps for computing $P(l \mid S_U)$ can be found in the appendix in Eq. 18. We experimentally verified the conditional independence of $V(e)$ given l for the three knowledge scenarios and therefore Eq. 8 applies equally to the different adversarial knowledge scenarios. Note that we effectively weaken the adversary by considering the products of different purchases independent from each other.

2.6 Privacy Metrics

We introduce three privacy metrics in order to capture the privacy of consumers revealing their purchase histories across different dimensions: We (i) measure the

performance of the adversary in identifying the true location with the F_1-score. Then, (ii) using the notion of mutual information [18], we quantify the absolute privacy loss of the consumer due to the adversary's knowledge of a price dataset. Finally, (iii) we use the relative reduced entropy as a relative privacy metric[6].

F_1-**SCORE:** The objective of the adversary is to assign the purchase events to the correct location. In the worst case, the adversary is forced to randomly guess among all possible locations. If the adversary, however, can estimate location probabilities more accurately, location privacy is reduced. Our problem corresponds to a multi-class classification problem and we therefore quantify the adversarial performance by averaging the F_1-score [35] of each individual class. The F_1-score corresponds to the harmonic mean of recall and precision, measuring the test's accuracy.

MUTUAL INFORMATION: A purchase event dataset enables the adversary to infer the distribution of prices among locations. Therefore, we want to measure how much privacy consumers lose when their purchase events are revealed and when the adversary has access to a dataset of purchase events. We quantify this privacy objective by measuring the *absolute* reduced location entropy given the purchase events. To this extent, we use the Mutual Information [18], denoted by $I(l, V(e))$, which measures how much the entropy of the locations is reduced given the purchase events (cf. Eq. 9).

$$I(l, V(e)) = \sum_{l \in \mathbb{L}, e \in S_U} P(l, V(e)) \cdot \log_2 \frac{P(l, V(e))}{P(l)P(V(e))} \tag{9}$$

RELATIVE REDUCED ENTROPY: Recall that the mutual information quantifies what we call the absolute privacy loss. In fact, there is an inherent randomness in the price distribution among locations. It is important to capture to what extent the original uncertainty about the locations can be reduced when a dataset of purchase events is given. The relative reduced entropy therefore captures the *relative* privacy, as the complement of the fraction of the conditional entropy over the location entropy. Given $H(l) = I(l, V(e)) + H(l \mid V(e))$, we compute the relative reduced entropy as $1 - \frac{H(l|V(e))}{H(l)}$ over all purchase events.

The proposed evaluation metrics are independent of a particular adversarial strategy. In return, the output of the privacy leakage quantification only depends upon the employed dataset of purchase events. In the next section we present the datasets utilized for our experimental evaluation.

3   Datasets

There are only a couple of datasets accurately accumulating the worldwide product price information. For individual products (e.g., a Big Mac [5] or Starbucks

---

[6] Defined as the complement of the fraction of conditional entropy over the location entropy.

coffee [8]), the average price values per country are available. Because a product often appears multiple times with different price values in the same country or city, the average is not a good estimator for elaborate studies. In the following, we describe the three independent price datasets considered in our work.

The first dataset, Numbeo [9], is a crowd-sourced dataset containing world-wide price values per product category, city and country. It is the most complete dataset of worldwide harvested prices available to our knowledge. We restricted our analysis to 23 frequently bought product categories, and split the Numbeo dataset into two separate datasets: (i) two years of data as the Numbeo dataset and (ii) five months of data as the Numbeo test dataset (cf. Table 3). Numbeo performs sanity checks on the crowdsourced inputs, and we additionally filtered extreme outlier [3][7] from the data to account for possible mistakes from crowd-sourced data. We identified 112 countries, with a total of 328,720 price values. Note that the provided data mostly contains prices from the US (18 %) and India (14 %).

The second dataset, referred to as the Chicago dataset [11], covers 84 stores in the Chicago metropolitan area over a period of five years. The data is sourced on a weekly basis from Dominick's supermarket stores. We sample 85 weeks with the most data, each containing on average 283,181 prices, spanning 28 product categories for an average of 6304 different products.

The third dataset originates from Kaggle [7], a Machine Learning competition platform. The dataset contains 350 million purchase events from 311,539 consumers across 134 store chains. The data is anonymized, but contains the individual product price, product category, date of purchase and purchase amount. Most purchase events cost less than 25 USD. The country of the dataset is not disclosed, but purchase prices are given in USD and purchase amounts are described in the imperial system.

In order to estimate the location probability, an adversary requires the knowledge of the population in each location. On the country granularity, we use the data available from the World Bank [12] for the year 2013, while for the US city granularity we used the data from the US Census Bureau [37].

As described in Sect. 2.2, we increase the knowledge of the adversary with the product basket. A product basket details which and how many products an average person purchases, both in terms of quantity and monetary amount. We leverage a national product basket [4] from 2010 containing over 300 product categories in order to infer the ratio in which different products are bought over the year.

4   Experimental Evaluation

In this section we evaluate the adversarial models designed in Sect. 2.2. We start by presenting the assumptions and choices made for the evaluation.

---

[7] $price < 25^{th}$ percentile $- 3 \cdot$ interquartile range, and
$price > 75^{th}$ percentile $+ 3 \cdot$ interquartile range.

4.1   Experimental Considerations

With respect to the value probability $P(v \mid l, c)$, we assume that the frequency of price values in the Numbeo dataset reflects the frequency of real-world purchase events with the corresponding price values. This is a natural assumption and is further motivated by the fact that e.g., Numbeo contributors likely entered the most popular price values for the considered product categories. Because our datasets contain a limited amount of products and product categories, our analysis is naturally confined to the available products. Note that, if the adversary knows the product categories of the purchases, e.g. milk, other categories such as apples can be ignored, which allows precise predictions with knowledge about few products. In order to compute the product category probability, $P(c \mid l)$, we only consider one national product basket and apply it to every country. Note that we do not use the product basket as an indicator of how much money is spent on average by a person, but rather as an indicator in which ratio products are bought.

SAMPLING PRICE VALUES: Given a location l, we generate *synthetic* consumer purchase events by sampling price values from the respective dataset. For the three datasets we consider adversaries with complete knowledge of the price values. In addition we instantiate an adversary with incomplete knowledge with the Numbeo test dataset. Given the product basket of the location l we compute the probability of a product category being sampled (cf. Eq. 4). Thus, we sample each product category with the product category probability $P(c \mid l)$. For each location we repeat the sampling of the price values $n = 1000$ times and average the result.

ADDITIVE SMOOTHING PARAMETER: In the case of an adversary with incomplete knowledge, we make use of additive smoothing to avoid zero probabilities when aggregating the probabilities of multiple purchase events for locations (see Sect. 2.2). We choose a smoothing parameter $\alpha = 0.01$ which provides us with the best results on our data (cf. appendix Fig. 6).

In the following, we evaluate up to three knowledge scenarios (cf. Sect. 2.3) for four location granularities: (i) across 112 countries worldwide; (ii) across 23 cities within the United States; (iii) across 84 stores within the Chicago metropolitan area; (iv) we distinguish among 134 store chains in a country.

4.2   Country Granularity

The adversary has to distinguish 112 candidate countries for each purchase event. We quantify the privacy given the three privacy metrics defined in Sect. 2.6. In particular, we performed our study in two settings. First, (i) we assumed that the adversary does not have complete knowledge. This means that the adversary receives purchase events from the Numbeo test dataset and estimate their location based on the Numbeo dataset. In the second case, (ii) the adversary assumes complete knowledge of price values, and therefore, the sampled prices are included in the price dataset which is adversarial knowledge.

Fig. 3. F_1-score for identifying the country given purchase events *sampled from the Numbeo test dataset*, corresponding to incomplete knowledge. We are not overfitting as we successfully classify new prices based on previously known prices.

Figure 3 shows the F_1-score for the first case based on the number of purchase events accessible to the adversary. Given one purchase event, the price, price_merchant and price_product-category knowledge scenario achieve an average of 0.38, 0.41 and 0.49 respectively. The high F_1-score after one purchase event shows, that even one event allows a decent prediction. We observe that the adversary is more likely to identify the correct location when it knows the product category of the purchase event. On the contrary, if the adversary has access to 10 purchase events, the respective F_1-scores are 0.80, 0.85 and 0.90. In other words, 10 purchase events significantly improve the ability of the adversary to identify the location of the purchase events. The reported values are averaged over $n = 1000$ iterations.
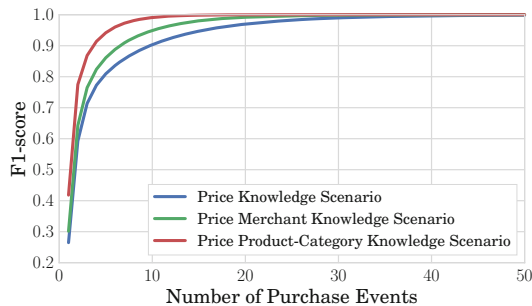

Fig. 4. F_1-score for identifying the country given purchase events *sampled from the Numbeo dataset*, corresponding to complete knowledge. Averaging does not hide poorly performing countries (cf. appendix).

Figure 4 corresponds to the second case, where the adversary assumes complete knowledge of the price values. We observe that the adversary can distinguish more accurately between the possible locations. The F_1-scores are averaged

Table 1. Mutual information and relative reduced entropy for the three knowledge scenarios when estimating the country, city, store or chain of purchase events. The respective abbreviations P., PM., PPC. stand for Price, Price Merchant and Price Product-Category knowledge scenario respectively.

Knowledge scenarios	112 countries			23 US cities			84 stores		134 chains	
	P	PM	PPC	P	PM	PPC	P	PPC	P	PPC
Mutual information	0.539	0.841	1.703	0.368	0.572	1.164	0.280	0.569	0.456	2.256
Relative reduced E.	0.114	0.178	0.360	0.101	0.157	0.319	0.044	0.089	0.068	0.337

over all considered countries. For each considered country in the price knowledge scenario, we verify that averaging does not hide poorly performing countries (cf. Fig. 7 in the appendix).

Table 1 presents the results of the mutual information and the relative reduced entropy for each knowledge scenario. We observe that the price_product-category knowledge scenario reduces the entropy more significantly than the other knowledge scenarios. Naturally, this is because the price_product-category knowledge scenario provides the adversary with more information than the price knowledge scenario, thus effectively reducing uncertainty when identifying the location.

4.3   US City Granularity

In this section we analyze an adversary that aims to distinguish among the purchase events of 23 US cities. As before, we quantify the privacy based on the three privacy metrics defined in Sect. 2.6. We sample and test purchase events on the Numbeo dataset only, since our test dataset does not contain sufficiently many purchase events per considered US city.

Figure 10 illustrates the F_1-score depending on the number of purchase events. We observe, that after 10 purchase events, the F_1-score is greater than 0.7. Therefore, our methodology also provides accurate estimations on a city granularity. Table 1 reports the mutual information and relative reduced entropy when estimating the US city. We observe that the relative reduced entropies of country and city granularity match across the knowledge scenarios. This exemplifies the usefulness of the relative reduced entropy to highlight similarities across different price datasets.

4.4   Chicago Metropolitan Granularity

In this section, we analyze an adversary that aims to distinguish among the purchase events of 84 Dominick's stores within the Chicago metropolitan area.

We sample the price values from the Chicago dataset, and assume an adversary with complete knowledge; we therefore do not apply additive smoothing. We consider the location prior probability $P(l)$ to be uniform, because we do not have reliable store popularity information for the Chicago area.

In Fig. 11 we can observe that the adversary can identify a local store given 100 purchase events with high confidence. We expected a weaker result, since all stores are operated by the same chain, implying relatively similar price structures. We ran our attack on each of the 85 weeks with most data, averaged the results and report the standard deviation as shown in the blue area of Fig. 11.

Table 1 shows that the Chicago price dataset reveals less information about the considered locations than the Numbeo dataset. This observation holds for both knowledge scenarios, and is consistent with the result that more price points are required to localize purchase events within the Chicago area.

4.5   Store Chain Granularity

The large-scale Kaggle dataset does not provide precise location information of purchase events, but allows the adversary to distinguish among 134 store chains. Knowing the store chain of purchase events effectively reduces the possible locations of the purchases. Note, that the prices of Kaggle are distributed over a year and the adversary therefore does not know the precise time of the purchase events.

We uniformly sample purchase events of different consumers and perform our attack on the Kaggle dataset. Figure 5 reveals that given approximately 250 price values we achieve an F_1-score of over 0.95 for the origin of the purchase events. Note, that the price_product-category knowledge scenario is particularly strong due to many product categories. This is reflected by the particularly high Mutual Information (cf. Table 1).
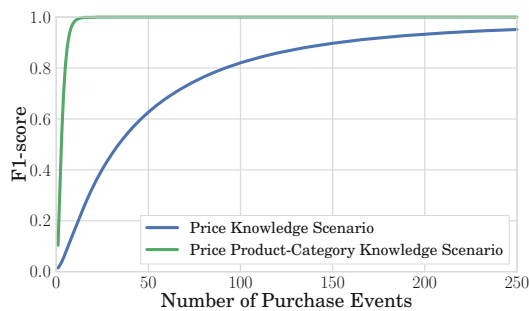

Fig. 5. F_1-score for identifying the store chain. The purchase events are sampled from the Kaggle dataset.

Given these results, we conclude, that our framework and methodology apply to a wide variety of different price datasets and allow us to quantitatively com-

pare their respective privacy leakage. In the following, we extract further insights from our data to strengthen the attack.

4.6 Most Revealing Product Category

In this section we investigate which of the 23 considered product categories from the Numbeo dataset leak more information. This is a useful insight since an adversary would pick purchase events of this product category in order to increase the probability of correctly identifying their location. Therefore, with the mutual information we measure the extent to which the location entropy is reduced, given the purchase events of a particular product category. Contrary to the previous analysis, we evaluate the mutual information *per product category* based on the price_product-category knowledge scenario defined in Sect. 2.3. More specifically, we compute the mutual information using only purchase events of a particular product category.

The results of the evaluation can be found in Fig. 13. According to this metric, the most revealing product categories are milk, a one-way ticket for local transportation, and a loaf of white bread. On the contrary, the product categories that disclose less information about a location are oranges, chicken breasts and rice.

4.7 Required Time Precision

Previously, we assumed that knowledge of the exact currency conversion rates is required to compare non-localized purchase events. Exact currency conversion rates, however, require a precise knowledge of the purchase event times. In this section, we show that our attack does not require the exact currency conversion rates, but also works if the adversary knows only the date or even week of the purchase, i.e. it has an uncertainty of 24 h or 7 days in relation to the conversion rates. We therefore relax the requirements on the time precision.

Due to the conversion rate differences, the adversarial estimation of $P(v \mid l, c)$ is inaccurate. To compensate for the conversion rate differences, the adversary can use a price tolerance. We study two options for the tolerance: a static tolerance and a dynamic tolerance. For the static tolerance, the adversary estimates $P(v \mid l, c)$ in the presence of uncertainty by considering price values in the interval $[v - tol_s, v + tol_s]$ where the static tolerance tol_s is a small amount in global currency (e.g., 0.02 USD). The dynamic tolerance value tol_d is a percentage-wise estimate of uncertainty (e.g., 2 %). To estimate $P(v \mid l, c)$ the adversary considers price values from the interval $[v \cdot (1 - tol_d), v \cdot (1 + tol_d)]$.

We evaluated the attack to infer the country of purchase events with imprecise purchase times and compensated the time error with different tolerance values. To simulate imprecise purchase times, we converted the adversarial knowledge using conversion rates of 30 different days from the year 2014 and then converted the non-localized purchase events S_U using the previous days' conversion rates. As before, we computed the F_1-score to evaluate the quality of the estimated $P(l \mid S_U)$.

For static and dynamic tolerance values, we found that the attack is still accurate, i.e. reaches an F_1-score above 95% with less than 50 purchase events. A higher tolerance value has two opposing effects: (i) it compensates for differences in currency conversion rates and increases the number of correctly considered price values; (ii) a higher tolerance, however, also increases the number of incorrectly considered price values which fall into larger intervals. Therefore, the tolerance value presents a trade off between the true-positive and true-negative rate. Our experimental results reflect this trade off both for static and dynamic tolerance values (cf. Appendix B). Based on our experimental results we propose a dynamic tolerance of 2% for a $24\,h$ time imprecision.

We also evaluated the uncertainty of one week on the currency conversion rates. We used real-world currency conversion rates that were seven days apart from each other. Figure 14 shows the result of this experiment for the different knowledge scenarios and a dynamic tolerance value of 2% on the Numbeo dataset. We conclude that our attack does not require precise purchase event times.

5   Related Work

Location Privacy. Blumberg [16] *et al.* provide a non-technical discussion of location privacy, its issues and implications. Gruteser and Grunwald [23] initiate major research in the area of the anonymization approaches to location privacy. Further, Narayanan *et al.* [29] investigate location privacy from a theoretical standpoint and present a variety of cryptographic protocols motivated by and optimized for practical constraints while focusing on proximity testing. Shokri *et al.* [34] propose a formal framework for quantifying location privacy in the case where users expose their location sporadically. They model various location-privacy-preserving mechanisms, such as location obfuscation and fake location injections. This work is orthogonal to ours, since in our setting the consumers are not willingly revealing their locations. Voulodimos *et al.* [38] address the issue of privacy protection in context-aware services through the use of entropy as a means of measuring the capability of locating a user's whereabouts and identifying personal selections. Narayanan [28] and Shmatikov propose statistical de-anonymization attacks against high-dimensional micro-data. We do not rely on their methods, since we are not aiming to de-anonymize the consumers. De Montjoye *et al.* [39] show that consumers can be uniquely identified within credit card records with only a few spatiotemporal triples containing location, time and price value. Contrary to their work, we focus on the price values and we localize instead of identify consumers.

Payment systems. The privacy implications of public transaction prices have been widely ignored. One prominent example is Bitcoin [17,33], where transactions are exchanged between peers by means of pseudonyms. The actual transaction prices are archived and publicly available. The literature features many different methods for analyzing the privacy implications of Bitcoin, e.g., by means

of appropriate heuristics [13], tainting [22], or other techniques [21,32]. Reid and Harrigan [31] analyze the flow of Bitcoin transactions in a small part of the Bitcoin log, and show that external information like publicly-announced addresses, can be used to link identities and organizations to some transactions. In [27] the authors propose Zerocoin, a cryptographic extension to Bitcoin that augments the protocol to allow for fully anonymous currency transactions using a distributed ECash scheme. To the best of our knowledge only two contributions [14,15] have aimed to hide the transaction prices in Bitcoin.

Price rigidity. Herrmann and Moeser [24] perform a quantitative analysis on price variability and conclude that prices are often rigid for several weeks. Pricing strategies for identical brands, however, vary significantly among retailers. Their observations match the studies of the Big Mac index [5] (the Economist), the Starbucks coffee index [8] (the Wall Street Journal) and the Ikea Billy Bookshelf index [2] (Bloomberg). The former studies show that prices of identical products from a single brand vary across locations. Dutta *et al.* [20] find that retail prices respond promptly to direct cost changes as well as upstream manufacturers' costs. Hosken and Reiffen [25] find that each product has a price mode—a price that the product stays at most of the time. Note that Hosken's non-public dataset contains nearly as many price observations as our Numbeo dataset.

6   Conclusion

Having a systematic methodology to reason quantitatively about the privacy leakage from datasets containing price relevant information is a necessary step to avoid privacy leakages. While further tests with more datasets will help to generally claim that price values alone can reveal the location of a purchase, our empirical results provide evidence that with relatively few purchase events it is possible to identify a consumer's location. In this paper, we have raised the following two questions: How much location information is leaked by consumer purchase datasets? How can it be quantified with the considered adversarial model and knowledge? In our proposed framework, we have modeled several adversaries and quantified the privacy leakage according to different dimensions. We make extensive use of Bayesian inference in our framework to model the different attack strategies. Our framework can be easily applied to any price dataset of consumer purchases and allows one to compare the privacy leakage of different datasets. We applied our methodology to three real-world datasets and achieve comparable results. The results presented in this paper strongly motivate the need for careful consideration when sharing price datasets and should be considered when designing public ledger cryptocurrencies.

Appendix A: Probability Calculations

In the following we clarify the individual steps for calculating the probabilities derived in Sect. 2.

$$P(v) = \sum_{l \in \mathbb{L}} P(l, v) = \sum_{l \in \mathbb{L}} \sum_{c \in \mathbb{C}} P(l, c, v)$$
$$= \sum_{l \in \mathbb{L}} \sum_{c \in \mathbb{C}} P(l) \cdot P(c \mid l) \cdot P(v \mid l, c) \tag{10}$$

$$P(c, v) = \sum_{l \in \mathbb{L}} P(l, c, v)$$
$$= \sum_{l \in \mathbb{L}} P(l) \cdot P(c \mid l) \cdot P(v \mid l, c) \tag{11}$$

$$P(v \mid l) = \frac{P(l, v)}{P(l)} = \frac{\sum_{c \in \mathbb{C}} P(l, c, v)}{P(l)}$$
$$= \sum_{c \in \mathbb{C}} P(c \mid l) \cdot P(v \mid l, c) \tag{12}$$

$$P(c, v \mid l) = \frac{P(l, c, v)}{P(l)} = \frac{P(l) \cdot P(c \mid l) \cdot P(v \mid l, c)}{P(l)}$$
$$= P(c \mid l) \cdot P(v \mid l, p) \tag{13}$$

$$P(l \mid v) = \frac{P(l) \cdot P(v \mid l)}{P(v)}$$
$$= \frac{P(l) \cdot \sum_{c \in \mathbb{C}} [P(c \mid l) \cdot P(v \mid l, c)]}{\sum_{l' \in \mathbb{L}} \sum_{c \in \mathbb{C}} [P(l') \cdot P(c \mid l') \cdot P(v \mid l', c)]}$$
$$= \frac{P(l) \cdot \sum_{c \in \mathbb{C}} [P(c \mid l) \cdot P(v \mid l, c)]}{\sum_{l' \in \mathbb{L}} P(l') \cdot \sum_{c \in \mathbb{C}} [P(c \mid l') \cdot P(v \mid l', c)]} \tag{14}$$

$$= \frac{\frac{\text{Population}(l)}{\sum_{l' \in \mathbb{L}} \text{Population}(l')} \cdot \sum_{c \in \mathbb{C}} [\frac{\text{Basket}(l,c)}{\sum_{c' \in \mathbb{C}} \text{Basket}(l,c')} \cdot \frac{D(l,c,v)}{D(l,c)}]}{\sum_{l' \in \mathbb{L}} \frac{\text{Population}(l')}{\sum_{l'' \in \mathbb{L}} \text{Population}(l'')} \cdot \sum_{c \in \mathbb{C}} [\frac{\text{Basket}(l',c)}{\sum_{c' \in \mathbb{C}} \text{Basket}(l',c')} \cdot \frac{D(l',c,v)}{D(l',c)}]} \tag{15}$$

$$P(l \mid c, v) = \frac{P(l) \cdot P(c, v \mid l)}{P(c, v)}$$
$$= \frac{P(l) \cdot [P(c \mid l) \cdot P(v \mid l, p)]}{\sum_{l' \in \mathbb{L}} [P(l') \cdot P(c \mid l') \cdot P(v \mid l', c)]} \tag{16}$$

$$= \frac{\frac{\text{Population}(l)}{\sum\limits_{l' \in \mathbb{L}} \text{Population}(l')} \cdot \frac{\text{Basket}(l,c)}{\sum\limits_{c' \in \mathbb{C}} \text{Basket}(l,c')} \cdot \frac{D(l,c,v)}{D(l,c)}}{\sum\limits_{l' \in \mathbb{L}} [\frac{\text{Population}(l')}{\sum\limits_{l'' \in \mathbb{L}} \text{Population}(l'')} \cdot \frac{\text{Basket}(l',c)}{\sum\limits_{c' \in \mathbb{C}} \text{Basket}(l',c')} \cdot \frac{D(l',c,v)}{D(l',c)}]} \tag{17}$$

$$
\begin{aligned}
P(l \mid S_U) &= P(l \mid V(e_1), V(e_2), \ldots, V(e_n)) \\
&= \frac{\prod\limits_{i=1..n} P(V(e_i))}{P(V(e_1), V(e_2), \ldots, V(e_n))} \\
&\quad \frac{\prod\limits_{i=1..n} P(l \mid V(e_i))}{P(l)^{n-1}} \\
&= \frac{P(l) \cdot \prod\limits_{e \in S_U} P(V(e) \mid l)}{P(V(e_1), \ldots, V(e_n))}
\end{aligned}
\tag{18}
$$

Appendix A.1: Probability Calculations

Based on its knowledge, the ideal adversary computes the following probabilities by computing the fractions of events.

$$P(l) = \frac{|\{e|e \in \mathcal{H}_G : e.l = l\}|}{|\mathcal{H}_G|} \tag{19}$$

$$P(v) = \frac{|\{e|e \in \mathcal{H}_G : e.v = v\}|}{|\mathcal{H}_G|} \tag{20}$$

$$P(c,v) = \frac{|\{e|e \in \mathcal{H}_G : e.c = c \wedge e.v = v\}|}{|\mathcal{H}_G|} \tag{21}$$

$$P(v \mid l) = \frac{|\{e|e \in \mathcal{H}_G : e.l = l \wedge e.v = v\}|}{|\{e|e \in \mathcal{H}_G : e.l = l\}|} \tag{22}$$

$$P(c,v \mid l) = \frac{|\{e|e \in \mathcal{H}_G : e.l = l \wedge e.c = c \wedge e.v = v\}|}{|\{e|e \in \mathcal{H}_G : e.l = l\}|} \tag{23}$$

Appendix B: Further Experimental Results

Appendix B.1: Required Time Precision

Figure 8 shows, that a larger tol_s will improve the overall F_1-score, but more purchase events are needed to filter out the false positives. Similarly, for the dynamic tolerance in Fig. 9, a higher value for tol_d provides a better prediction for many purchase events, but a worse prediction for few purchase events. The figures show the experiments for the price_product-category knowledge scenario, however, we note that the results are analogous to the other scenarios. Based on these results we propose a dynamic tolerance of 2 % in the case of a 24 h time imprecision on the conversion rate.
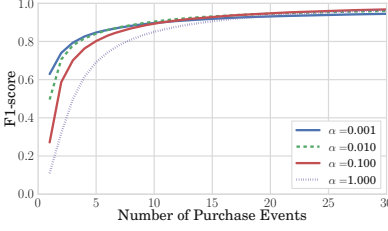
Fig. 6. Comparison of different α-parameters for additive smoothing based on the price_product-category knowledge scenario.
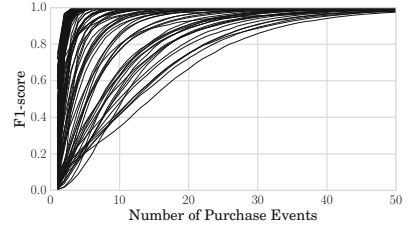

Fig. 7. F_1-score of each individual country for the price knowledge scenario. The purchase events are sampled from Numbeo. We observe that no country performs poorly.
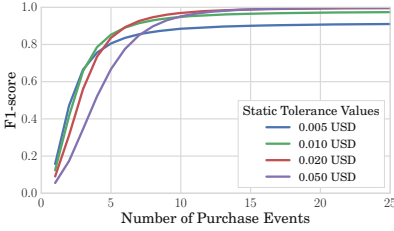

Fig. 8. Using static tolerance values to compensate for imprecise time information (one day uncertainty) in the price_product-category knowledge scenario.
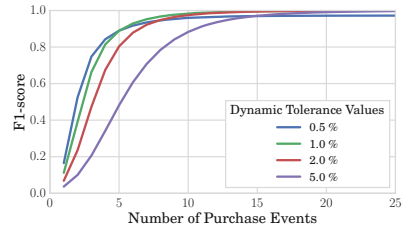

Fig. 9. Using dynamic tolerance values to compensate for imprecise time information (one day uncertainty) in the price_product-category knowledge scenario.
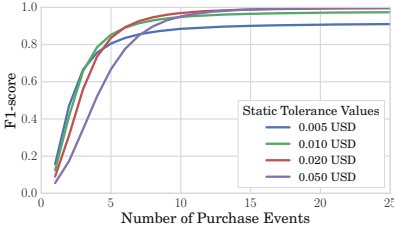

Fig. 10. F_1-score for identifying the US city given purchase events for different knowledge scenarios. The purchase events are sampled from the Numbeo dataset.
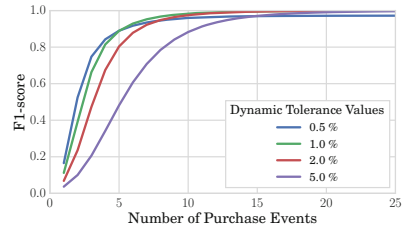

Fig. 11. F_1-score and standard deviation over 85 weeks for identifying the store in the price knowledge scenario. Data sampled from the Chicago dataset among 84 stores.

Appendix B.2: Motivating Example

Since products appear in a multitude of price values, it is at first unclear how accurately price values can identify a location. To illustrate why purchases can be localized, we focus on an example of the product category *domestic beer (0.5 L bottle)*, which can be bought in nearly every country. The price values are taken from the Numbeo dataset [9]. Figure 12 shows the distribution of price values of

beer in USD for four countries. We observe that ranges of prices clearly differ for India and the other countries, while prices in Australia are more likely to be higher than in the US and Canada, where distributions of prices are similar. Given a beer price above 3 USD, in this case, it is highly likely that the purchase has not occurred in India.
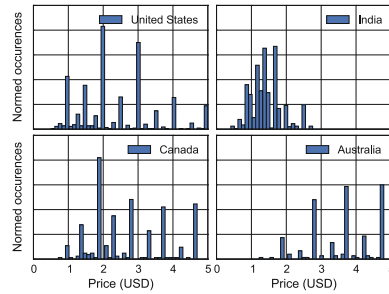

Fig. 12. Distribution of domestic beer prices (0.5 L) in 4 countries from Numbeo in USD.

Appendix C
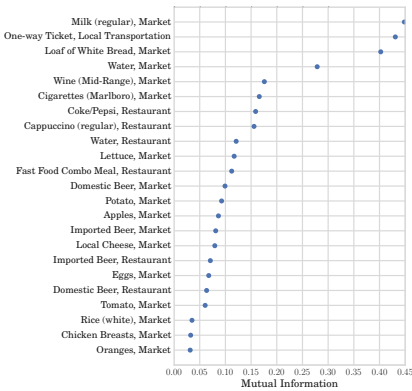
See Figures 13, 14, Tables 2 and 3.

Fig. 13. The higher the mutual information, the more revealing is the product category.

Fig. 14. Dynamic tolerance of 2 % with one week time uncertainty on the Numbeo dataset while estimating the country. Precise time allows an F_1-score of 0.95 after 10 purchase events whereas a one week time uncertainty achieves an F_1-score of 0.63.

Table 2. Product categories of the Numbeo dataset.

Category and merchant	Unit	Prices
Market product categories		
Apples	1 kg	11876
Chicken Breasts	1 kg	11893
Cigarettes (Marlboro)	1 pack	12712
Domestic Beer	One 0.5 L bottle	10243
Eggs	12 units	14617
Imported Beer	One 0.33 L bottle	9484
Lettuce	1 head	8966
Loaf of White Bread	0.5 kg	14633
Local Cheese	1 kg	10975
Milk (regular)	1 L	17197
Oranges	1 kg	10289
Potato	1 kg	10891
Rice (white)	1 kg	10924
Tomato	1 kg	10539
Water	1.5 L bottle	12762
Wine (Mid-Range)	1 bottle	11893
Restaurant product categories		
Cappuccino (regular)	1 unit	21539
Coke/Pepsi	One 0.33 L bottle	21351
Fast Food Combo Meal	1 unit	21794
Domestic Beer	One 0.5 L bottle	19128
Imported Beer	One 0.33 L bottle	18048
Water	One 0.33 L bottle	21691
Local transportation categories		
One-way Ticket	1 unit	15275

Table 3. Statistics about the three price datasets

Numbeo dataset (2 years)	
Number of countries	112
Number of prices	328,720
Number of cities in the US	23
Number of prices in the US cities	11,686
Number of distinct product categories	23
Numbeo test dataset (5 months)	
Number of countries	47
Number of prices	40,968
Number of distinct product categories	23
Chicago dataset (5 years)	
Number of stores	84
Number of total prices in top 85 weeks	24,070,437
Average number of prices per week	283,181 ± 6790
Number of distinct product categories	28
Average number of products per week	6304 ± 461
Kaggle dataset (1 year)	
Number of store chains	134
Number of purchase events	349,655,789
Number of consumers	311,539
Number of distinct product categories	836

References

1. A Face Is Exposed for AOL Searcher No. 4417749 (2006). http://www.nytimes.com/2006/08/09/technology/09aol.html
2. Ikea Billy Bookshelf Index, Bloomberg (2009). http://www.bloomberg.com/apps/news?pid=newsarchive&sid=a.K4T4ypP9ko
3. NIST/SEMATECH e-Handbook of Statistical Methods (2013). http://www.itl.nist.gov/div898/handbook/
4. Anonymized for review (2015)
5. Big Mac Index, The Economist (2015). http://www.economist.com/content/big-mac-index
6. Consumer panel data and retail scanner data across the United States (2015). http://research.chicagobooth.edu/nielsen/
7. Kaggle, Acquire Valued Shoppers Challenge (2015). https://www.kaggle.com/c/acquire-valued-shoppers-challenge
8. More (or Less) Brew for your Buck, Starbucks coffee price (2015). http://online.wsj.com/news/articles/SB10001424127887324048904578319783080709860
9. Numbeo, database of user contributed data about cities and countries worldwide (2015). http://www.numbeo.com
10. Ripple, cryptocurrency (2015). https://ripple.com/
11. Store-level scanner data collected at Dominick's Finer Foods (2015). http://research.chicagobooth.edu/kilts/marketing-databases/dominicks/dataset
12. World Population, The world bank (2015). http://data.worldbank.org/indicator/SP.POP.TOTL?order=wbapi_data_value_2009+wbapi_data_value+wbapi_data_value-first&sort=asc
13. Androulaki, E., Karame, G.O., Roeschlin, M., Scherer, T., Capkun, S.: Evaluating user privacy in bitcoin. In: Sadeghi, A.-R. (ed.) FC 2013. LNCS, vol. 7859, pp. 34–51. Springer, Heidelberg (2013). http://eprint.iacr.org/2012/596.pdf
14. Androulaki, E., Karame, G.O.: Hiding transaction amounts and balances in bitcoin. In: Holz, T., Ioannidis, S. (eds.) Trust 2014. LNCS, vol. 8564, pp. 161–178. Springer, Heidelberg (2014)
15. Ben-Sasson, E., Chiesa, A., Garman, C., Green, M., Miers, I., Tromer, E., Virza, M.: Zerocash: decentralized anonymous payments from bitcoin. In: 2014 IEEE Symposium on Security and Privacy (SP). IEEE (2014)
16. Blumberg, A.J., Eckersley, P.: On locational privacy, and how to avoid losing it forever. EEF (2009)
17. Bonneau, J., Miller, A., Clark, J., Naryanan, A., Kroll, J.A., Felten, E.W.: SoK: bitcoin and second-generation cryptocurrencies. In: IEEE Security and Privacy, May 2015
18. Cover, T.M., Thomas, J.A.: Elements of Information Theory. John Wiley and Sons, Hoboken (2012)
19. de Montjoye, Y.-A., Hidalgo, C.A., Verleysen, M., Blondel, V.D.: Unique in the crowd: the privacy bounds of human mobility. Sci. Rep. **3** (2013)
20. Dutta, S., Bergen, M., Levy, D.: Price flexibility in channels of distribution: evidence from scanner data. J. Econ. Dyn. control **26**(11), 1845–1900 (2002)
21. Meiklejohn, S., et al.: A fistful of bitcoins: characterizing payments among men with no names. In: Proceedings of the 2013 Conference on Internet Measurement Conference, IMC 2013, pp. 127–140. ACM, New York (2013)
22. Gervais, A., Karame, G., Capkun, S., Capkun, V.: Is bitcoin a decentralized currency? IEEE Secur. Priv. Mag. **12**, 54–60 (2014)

23. Gruteser, M., Grunwald, D.: Anonymous usage of location-based services through spatial and temporal cloaking. In: Proceedings of the 1st International Conference on Mobile Systems, Applications and Services, pp. 31–42. ACM (2003)
24. Herrmann, R., Möser, A.: Price variability or rigidity in the food-retailing sector? theoretical analysis and evidence from german scanner data. Technical report (2003)
25. Hosken, D., Reiffen, D.: Patterns of retail price variation. RAND J. Econ., 128–146 (2004)
26. Manning, C.D., Raghavan, P., Schütze, H.: Introduction to Information Retrieval, vol. 1. Cambridge University Press, Cambridge (2008)
27. Miers, I., Garman, C., Green, M., Rubin, A.D.: Zerocoin: anonymous distributed e-cash from bitcoin. In: 2013 IEEE Symposium on Security and Privacy (SP), pp. 397–411. IEEE (2013)
28. Narayanan, A., Shmatikov, V.: Robust de-anonymization of large sparse datasets. In: IEEE Symposium on Security and Privacy, SP 2008. IEEE (2008)
29. Narayanan, A., Thiagarajan, N., Lakhani, M., Hamburg, M., Boneh, D.: Location privacy via private proximity testing
30. Pass, G., Chowdhury, A., Torgeson, C.: A picture of search. In: Proceedings of the 1st International Conference on Scalable Information Systems, InfoScale 2006. ACM, New York (2006)
31. Reid, F., Harrigan, M.: An analysis of anonymity in the bitcoin system
32. Ron, D., Shamir, A.: Quantitative analysis of the full bitcoin transaction graph (2013). http://eprint.iacr.org/2012/584.pdf
33. Nakamoto, S.: Bitcoin: a peer-to-peer electronic cash system (2009)
34. Shokri, R., Theodorakopoulos, G., Danezis, G., Hubaux, J.-P., Le Boudec, J.-Y.: Quantifying location privacy: the case of sporadic location exposure. In: Fischer-Hübner, S., Hopper, N. (eds.) PETS 2011. LNCS, vol. 6794, pp. 57–76. Springer, Heidelberg (2011)
35. Sokolova, M., Lapalme, G.: A systematic analysis of performance measures for classification tasks. Inf. Process. Manag. **45**(4), 427–437 (2009)
36. Sweeney, L.: Simple demographics often identify people uniquely. Health (San Francisco) **671**, 1–34 (2000)
37. U.S. Census Bureau, Population Division. Annual Estimates of the Resident Population for Incorporated Places of 50,000 or More, Ranked by July 1, 2013 (2014)
38. Voulodimos, A.S., Patrikakis, C.Z.: Quantifying privacy in terms of entropy for context aware services. Identity Inf. Soc. **2**(2), 155–169 (2009)
39. Singh, V.K., Pentland, A.S., de Montjoye, Y.-A., Radaelli, L.: Unique in the shopping mall: on the reidentifiability of credit card metadata. Science **347**, 536–539 (2015)