Chapter 10
Valuating Privacy with Option Pricing Theory

Stefan Berthold and Rainer Böhme

Abstract One of the key challenges in the information society is responsible handling of personal data. An often-cited reason why people fail to make rational decisions regarding their own informational privacy is the high uncertainty about future consequences of information disclosures today. This chapter builds an analogy to financial options and draws on principles of option pricing to account for this uncertainty in the valuation of privacy. For this purpose, the development of a data subject's personal attributes over time and the development of the attribute distribution in the population are modeled as two stochastic processes, which fit into the Binomial Option Pricing Model (BOPM). Possible applications of such valuation methods to guide decision support in future privacy-enhancing technologies (PETs) are sketched.

10.1 Introduction

In certain jurisdictions, the right of informational self-determination implies active control of one's personal data. To exercise such control, it is crucial for people to understand the implications of data disclosure. While visions for privacy-enhanced identity management [22] seek to provide technical means for *securing the disclosure* of personal data under different threat models, it is still a challenging question how individuals can be supported in *assessing the value* of their personal data. However, the latter is a prerequisite for the former: making informed disclosure decisions depends on the ability to compare between the alternatives in the first

Stefan Berthold
Karlstads Universitet, Fakulteten för Ekonomi, Kommunikation och IT, Universitetsgatan 2, 651 88 Karlstad, Sweden, e-mail: stefan.berthold@kau.se
Rainer Böhme
International Computer Science Institute, 1947 Center Street, Ste 600, Berkeley, CA 94704, USA, e-mail: rainer.boehme@icsi.berkeley.edu

T. Moore et al. (eds.), *Economics of Information Security and Privacy,*
DOI 10.1007/978-1-4419-6967-5_10,

place. Irrespective of the concrete supporting technology, a major obstacle that prevents people from making rational decision regarding their privacy is the uncertainty about possible future consequences of data disclosure at present [1]. Similarly, known information-theoretic privacy metrics at best reflect the present value of personal data. These metrics ignore that the value of personal data, for instance for re-identification, may change over time. However, the time between disclosure and exploitation of personal data is very relevant for the inter-temporal value of personal data: the more time passes between both events, the more uncertainty arises about the value. This is so because attribute values which apply to a data subject at the time of disclosure may not be applicable to the same data subject anymore when the data is exploited. Also the distribution of attribute values in the entire population changes over time. Attribute values which uniquely describe a single data subject at present may become common in the population in the future. Accordingly, their value for the purpose of re-identification would decline over time.

In this chapter, we present a framework to model this kind of uncertainty and account it in measures of the future value of attributes that are to be disclosed at present. Although novel to the field of privacy research, modeling uncertainty about future states has a long tradition in other disciplines, such as finance and accounting. So we will draw on concepts from option pricing theory and show how this theory translates to the problem of personal data disclosure. The core idea is to interpret data disclosure as writing a call option that allows the counterpart to use the data for identification later on.

To start with a simple case and focus on the core idea, we confine ourselves in this chapter to binomial stochastic processes, similar to the Binomial Option Pricing Model (BOPM) [11]. In principle, the theory generalizes so that any stochastic process with better fit to reality can be plugged into our framework. The choice of the most appropriate process for specific attributes in a certain context is an empirical question. It thus falls beyond the scope of this work. Again for the sake of simplicity, we limit our view to a single attribute with finite and discrete attribute values. Extensions to multiple attributes are possible, but increase the dimensionality of the problem substantially. We further rule out any ambiguity or measurement error and assume that exactly one attribute value can be assigned to each data subject.

Under the above-stated assumptions, the value of an attribute to re-identify a data subject after some time is determined by a combination of two factors:

1. by the chance that the attribute value still applies to the particular data subject. This factor is governed by the individual behavior of the data subject. So we will refer to it as the *micro* level.

   And, if this condition holds,

2. by the uniqueness of the attribute value, i. e., how many other data subjects in the population do meanwhile share the same attribute value and thus form an equivalence class? This factor is driven by the aggregate behavior of all, possibly heterogeneous, data subjects in the population. So we will refer to it as the *macro* level.

In our framework, each factor is a source of uncertainty and can be modeled by a stochastic process from the point of view of a transaction counterpart, who

1. learns the attribute value of a data subject at the time of disclosure, and
2. can observe the distribution of attribute values in the population at any time (e. g., through representative anonymous surveys or observation).

Hence, changes of individual attributes remain private information of each data subject. We deem this a reasonable and practical abstraction.

The remainder of this chapter is organized as follows. Section 10.2 recalls existing approaches to quantify anonymity and privacy in databases and communication systems as well as generalizations. Since none of these metrics is designed to consider value over time, inspiration is sought from financial mathematics. We briefly review existing adaptations of quantitative financial methods to information security before we present our notion of *privacy options* in Sect. 10.3, the 'heart' of this chapter. Section 10.4 implements the ideas in a concrete proposal to model the two relevant quantities as independent stochastic processes: a state-space model is suggested for individual attribute value transitions (Sect. 10.4.1), and a binomial random walk serves as proxy for the distribution of attribute values in the population (Sect. 10.4.2). We combine both components to a valuation method in Sect. 10.5 and interpret the results in Sect. 13.5. The concluding Sect. 10.7 sketches future directions.

10.2 Related Work

We have identified two areas of relevant prior art. First, measurement of privacy with information theory and probability calculus has some tradition as a sub-field of computer science [25]. Section 10.2.1 briefly reviews this string of research. Second, another set of relevant publications are prior attempts to adopt quantitative methods from finance to information security and privacy. These are summarized in Sect. 10.2.2.

10.2.1 Measurement of Anonymity and Unlinkability

Measuring *anonymity* with information theory was—to the best of our knowledge—first motivated in the 1980s after a public debate about the census in Germany.[1] Fischer-Hübner [16, 17] uses the entropy of attributes (columns) in a database, for instance demographic data in a census survey, to measure their average information. This way, it is possible to compute the average number of records in the database that would match a given set of attributes. The degree of anonymity (or the "risk of re-identification" in [17]) is the reciprocal of this number of records. Attempts to

[1] *Confidentiality* in statistical databases has a much longer research track, e. g., [13, 35].

measure anonymity in statistical databases [42] have led to a number of combinatorial metrics, most prominently *k*-anonymity [40].

Aside from statistical databases, benchmarking anonymous communication systems has stirred a need for research on privacy metrics. Díaz et al. [14] as well as Serjantov and Danezis [36] propose Shannon entropy [37] to measure the uncertainty of an outside observer about the assignment of users to roles (sender, recipient, uninvolved) in a communication system. Shannon entropy quantifies the amount of additional information an observer would need in order to unanimously identify the role of the user. From this metric, it is possible to calculate the average size of the *anonymity set* [33] an anonymous communication system can provide. The larger the entropy the more information is effectively concealed from the observer, and hence the more anonymous the users of a system are. By contrast, Tóth et al. [41] point out that even if a communication system provides a reasonable degree of anonymity *on average*, the probability for *a single user* of being identifiable can still be unacceptably high. Therefore Tóth et al. define an upper bound for the probability of identification as *degree of anonymity*, which no user must exceed [41].

Another modification is to relax the strict focus on communication systems and model *unlinkability* between two arbitrary items [33, 39]. This view has been taken up for example by Clauß [10], who approximates unlinkability measures in a model world where each data subject's *identity* is defined by a set of finite discrete attributes. Only part of their values may be known to an outside observer. So a data disclosure decision effectively deals with the problem of whether or not an additional attribute value (previously unknown to the observer) should be disclosed. Our model assumptions later in Sect. 10.4 are compatible with this stylized view of the world. Though not carried out in this chapter, our approach is extendable to joint unlinkability measures between more than two items. Obviously, there exist infinitely many projections that map the resulting probability space over the exponentially growing number of set partitions to a scalar. Specific instances of such projections with more [18] or less [15] clear information-theoretic interpretation have been proposed in the literature as concrete metrics of unlinkability.

Most of existing privacy metrics were conceived with the aim to compare between alternative technical systems. All methods have in common that the value of personal data is measured at a single point in time[2] and not account for its value in possible future states. When the area of application shifts from comparing systems to supporting individual disclosure decisions, this limitation prevails: existing metrics neglect the fairly accepted principle that so-called *adversaries* against one's informational privacy will never forget any information disclosed to them (see for instance [33]). As already outlined in the introduction, the inability of individuals to anticipate future states in disclosure decisions is named as the main reason to explain partly puzzling results from laboratory experiments that try to measure people's valuation of personal data empirically [1, 5, 24].

[2] We are aware about only one commendable exception: a metric targeted to location-privacy [44], which accounts for changing locations over time.

10.2.2 Financial Methods in Information Security

Option pricing has its roots in financial mathematics and deals with finding the 'fair' price for contracts that allow their holders to choose between a security and a fixed amount of money at a future point in time. The field has grown rapidly since the seminal work by the meanwhile Nobel laureates Black, Scholes and Merton [6, 28] was published in the 1970s. Financial options became a popular tool for risk managers because they allow portfolio managers to 'hedge' idiosyncratic risks on financial markets, that is to shape the distribution of possible outcomes in sophisticated ways and thereby adjust it to the investor's risk appetite. But the idea soon spread to other domains than marketable securities. So-called *real options* have been proposed to gauge investment decision, in particular in project management [3]. They are tools to model project risk and opportunities with sound financial valuation methods to compare between alternatives. One advantage of real options in project management is the possibility to anticipate midcourse strategy corrections to react to uncertain future states.

Several authors have proposed to apply real options to information security investment [12, 20, 23, 26] to complement other accounting metrics, such as return on information security investment (ROSI) and annual loss expectancy (ALE) [19, 34, 38]. Interestingly, Gordon et al. [20] use real options to criticize security overinvestment, whereas Daneva [12] makes the case for higher spending.

Other applications of financial methods on specific information security problems include Matsuura's [27] option pricing approach to model the value of what he calls *digital security token*. These tokens can be thought of as media objects with attached protection, as suggested in the context of digital rights management (DRM). In [8], we have adapted the idea of prediction markets [43] to fix incentives in software vulnerability disclosure with so-called *exploit derivatives*. Ozment [30] has tackled vulnerability disclosure with auction theory.

To the best of our knowledge, this work is the first to apply option pricing theory to informational privacy. Neither are we aware of any work in other domains that suggests financial derivatives written on information measures (in Shannon's sense [37]) as underlying.

10.3 From Financial to Privacy Options

The key idea of this work is that disclosing a single attribute value can be interpreted as writing an option for exploiting the attribute in the future. Here, 'to exploit' refers to the act of using the attribute to draw inference on the data subject's identity or preference, and to base decisions on this information that may affect the data subject. One prominent example brought forward by Odlyzko [29] and Acquisti and Varian [2] is price discrimination in buyer–seller relationships. Thus, the data subject who discloses an attribute value thereby writes an option, whereas the transaction counterpart buys an option to use the information for decision-making. We

follow the convention in the information security literature and further refer to the transaction counterpart as *adversary*. This term reflects a convention and should not be interpreted as an adoption of the normative view that collecting personal data is necessarily hostile or evil.

Most elements of financial option pricing theory have direct correspondences in our notion of *privacy options*.

The *currency* in which privacy options are denominated is *information* in Shannon's [37] sense. Knowing an attribute value (i. e., holding the option), if valid, helps to reduce the uncertainty of the adversary about the identity of the data subject. The means to express uncertainty in information theory is entropy and the contribution of the attribute value has information value. The *unit* of information is *bits*.

The *underlying asset* of privacy options is the disclosed attribute value and the *market price* corresponds to the information (in Shannon's sense) which the adversary gains from the attribute value by exploiting it.

The privacy option is a *call option*, in which the data subject takes a *short position*. The asset, that is the attribute value, is handed over to the adversary (*long position*) at the time of the option purchase. The action that may be performed by the adversary is *exploiting* the underlying attribute value rather than *buying* the underlying security.

The correspondence to the *premium* is the compensation the adversary has to pay in return for the attribute value. However, this compensation is not necessarily denominated in the currency 'information'. For example, a merchant could offer a small rebate to the sales price to incentivize the use of loyalty cards from which personal data can be collected. This way, empirical measurements of this monetary premium, such as in [21, 24], could be linked to information-theoretic quantities by calibrating information-utility functions.

The increasing uncertainty about the linkability of the attribute value to the data subject can be interpreted as *interest rate* of an alternative investment: the probability of a valid link between the disclosed attribute value and the data subject decreases with the time elapsed since the disclosure of an attribute value. The value of the option decreases proportionately to the probability of a valid link because this linkability determines whether the adversary can benefit from the option at all.

Analogies also exists for the distinction of the two vanilla option styles, i. e., the *American* option and the *European* option. The difference between both styles is the time period in which the option may be exercised. An American option may be exercised at any time starting from the purchase of the option until it expires. This applies to the situation where a service provider does not depend on the assistance of the data subject for exploiting the data after the data subject has once disclosed its attribute value. An European option may only be exercised at the date of expiry. This applies to situations where the benefit for the adversary depends on some action of the data subject. For example, a personalized purchase history is only valuable to a seller if (and when) the data subject decides to revisit his store [2, 9].

Other elements of financial options do not have direct correspondences in our notion of privacy options developed in this chapter. *Put options* are impractical since 'negative information' does not exist. They could, however, make sense in special

(and largely hypothetical) cases where deletion of previously disclosed personal data can be enforced [7]. Due to the non-rivalrous nature of information goods, we were also unable to conceive a correspondence to *dividends* of the attributes underlying our privacy options. Finally, the *strike price* (or *exercise price*) is the amount of money to be paid when the option is actually exercised. If exploiting the attribute value does not depend on other attributes, then there is no way to enforce a transfer of money or information, hence the strike price is always zero. One can conceive to change this by introducing a trusted third party who acts as an information broker, or by allowing for partial disclosure of multiple dependent attributes. Another interpretation for the strike price is the effort of the adversary to retrieve the personal data at the time of exploitation. It may vary with organizational and technical factors, but it is largely determined by the adversary and not—like for financial options—by the contract itself. All this highlights that there is room for further extension of the analogy, though they are clearly beyond the scope of this chapter.

10.4 Sources of Uncertainty

In this section, we specify models for each source of uncertainty. In order to keep the calculations tractable, we model the two sources of uncertainty as *independent* stochastic processes; more precisely, the *timed linkability process* for attribute value changes of a single data subject (microscopic view, Sect. 10.4.1), and another stochastic process that drives the *distribution of attribute values in the population* (macroscopic view, Sect. 10.4.2). The latter model has many similarities with simple models of asset value fluctuations in financial option pricing.

10.4.1 Micro Model: Timed Linkability Process

Attribute values that have just been disclosed by a data subject are linkable to the data subject by the adversary. Here, we do not consider misinformation and thus assume links to be valid as long as the data subject does not change—intentionally or unintentionally—to another attribute value. We further assume that it is generally possible to change attribute values, however, the actual change, particularly its time and the new value, is not observable by the adversary.

This suggests modeling the attribute values over time as a stochastic process. The process can be expressed in a (discrete time-invariant) state-space model without inputs nor outputs. The state vector $\mathbf{x}(t)$ contains the probability of a valid link in the first element and the probability of an invalid link in the second element. The next state $\mathbf{x}(t+1)$ of this state-space model is defined in a recursive manner depending on the current state $\mathbf{x}(t)$ and a state transition matrix $\mathbf{A}$,

$$\mathbf{x}(t+1) = \mathbf{A}\mathbf{x}(t)\,. \tag{10.1}$$

Elements $a_{i,j}$ of $\mathbf{A}$ hold the probability of a state change from state j to state i. The absence of inputs allows us to simplify the model and use matrix multiplication instead of recursion to calculate a particular $\mathbf{x}(t+1)$,

$$\mathbf{x}(t+1) = \mathbf{A}^{t+1}\mathbf{x}(0)\ . \tag{10.2}$$

In the simplest case, the state matrix $\mathbf{A}$ has dimension 2×2 and is defined by only two probabilities, p and $\bar{p}$. Let p be the probability that the data subject keeps its linkable attribute value and $\bar{p}$ be the probability that a data subject, who once changed the attribute value to something unlinkable, does not revert to the linkable attribute value. The state vector $\mathbf{x}(0)$ at the time of disclosure is

$$\mathbf{x}(0) = \begin{pmatrix} 1 \\ 0 \end{pmatrix}\ . \tag{10.3}$$

The first element of vector $\mathbf{x}(0)$ holds the initial probability of linkability, which equals 1 by definition: the attribute value is definitely linkable when it has just been disclosed. Accordingly, we define the state matrix $\mathbf{A}$ as

$$\mathbf{A} = \begin{pmatrix} p & 1-\bar{p} \\ 1-p & \bar{p} \end{pmatrix}\ . \tag{10.4}$$

This allows us to model time aspects of attribute value changes. If, for instance, the attribute describes the attribute haircut and its value is *ponytail*, then the attribute might change instantly to any other value that describes a shorter haircut, but, naturally, hair cannot grow as fast as it can be cut off. And thus the probability of reverting back to *ponytail* is limited by a natural upper bound. Assume that the probability of keeping that haircut would be fairly high. Then Fig. 10.1 illustrates a hypothetical development of the probability of linkability over time. The functional form of Eq. (10.2) imposes an exponential decay.

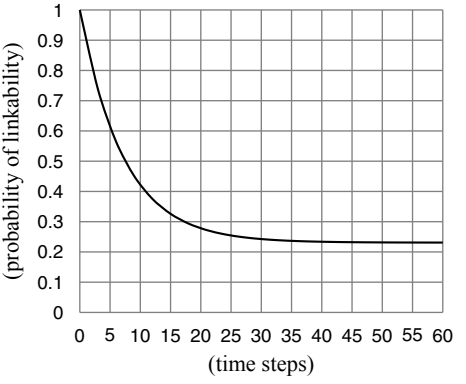

Fig. 10.1 Development of the probability of linkability if both the probability p of keeping the disclosed attribute value and the probability $\bar{p}$ of staying with anotnher attribute value are high. The diagram shows 60 time steps for $p = 0.9$ and $\bar{p} = 0.97$.

Other attributes might follow different processes, say, two attribute values and whenever the attribute has taken one value, the data subject tends to choose the other one with high probability. One can think of this as a model of fashions that alternate every couple of years. Thus, after the disclosure of the attribute value, it is possible to predict the values in the future, but with exponentially decreasing certainty. Fig. 10.2 depicts such a setting.

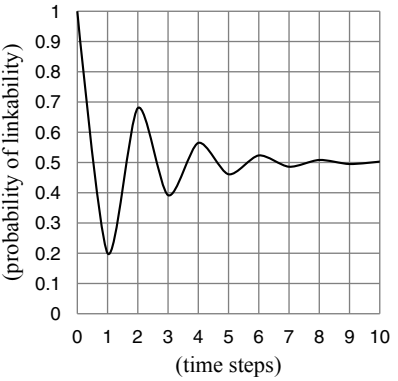

Fig. 10.2 Development of the probability of linkability if both the probability p of keeping the disclosed attribute value and the probability $\overline{p}$ of staying with another attribute value is small. The diagram shows 10 time steps for $p = 0.2$ and $\overline{p} = 0.2$.

Yet another situation emerges for attributes such as passport numbers: there is a vast number of different attribute values. The probability of requesting a new passport and therefore changing the attribute value might be small, depending on the travel habits of the data subject and on constraints imposed by the issuing country. But the probability of reverting back to exactly the same passport number is negligibly small. If we assume that this probability is in fact zero, then is it easy to see that the probability of linkability in the state-space model reduces to an exponential function of p, since for $\overline{p} = 1$, it holds that (after t time steps)

$$\mathbf{x}(t) = \mathbf{A}^t \mathbf{x}(0) = \begin{pmatrix} p & 0 \\ 1-p & 1 \end{pmatrix}^t \begin{pmatrix} 1 \\ 0 \end{pmatrix} = \begin{pmatrix} p^t & 0 \\ 1-p^t & 1 \end{pmatrix} \begin{pmatrix} 1 \\ 0 \end{pmatrix} = \begin{pmatrix} p^t \\ 1-p^t \end{pmatrix} . \qquad (10.5)$$

In Fig. 10.3, we show an example for the development of linkability, if $\overline{p} = 1$.

Note that generalizations to higher-order state-space models are possible and can be useful to represent other than binary attributes. We defer discussion of and examples for this case to future work.

10.4.2 Macro Model: Population Development

In the population, individual data subjects can be distinguished by their attribute values. A metric for the average discernibility is the self-information of an attribute value in the population. In terms of Shannon's information theory, the attribute can

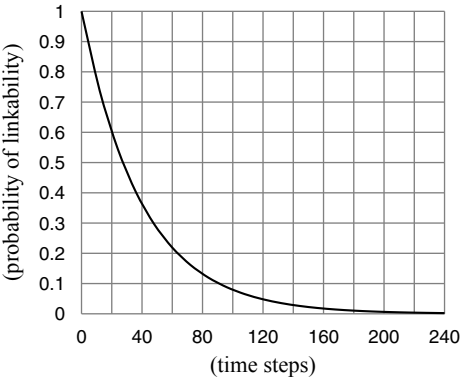

Fig. 10.3 Development of the probability of linkability if the probability $1 - \bar{p}$ of returning to the same attribute value after a change is zero. The diagram shows 240 time steps for $p = 0.975$ and $\bar{p} = 1$.

be understood as source of information, the attribute values as alphabet, and the (relative) frequency of each attribute value as the probability of the symbol. Let v be an attribute value and r_v be the relative frequency of this value in the population, then

$$H_v = -\log_2 r_v \tag{10.6}$$

is the self-information of the attribute value. It expresses the amount of information conveyed by the attribute value to the adversary, who may exploit it to re-identify the data subject.

The relative frequencies vary over time depending on how the individual data subjects change their attribute values. Thus, also the self-information of each attribute value fluctuates. The straight approach of modeling the behavior of *all* data subjects, their attribute values, and the changes over time by an aggregation of many micro-level models would be analytically intractable and computationally demanding. Moreover, generalizing one fixed state-space model of Sect. 10.4.1 to all data subjects neglects possible heterogeneity between them and is therefore debatable with theoretical arguments. Instead, we model the macroscopic changes of the distribution of attribute values in the population as a separate stochastic process.

A similar approach is taken in financial option pricing, where the market price of the underlying asset can be modeled in a similar way [11]. Both the market price and the relative frequency of the attribute value can *move up* or *down* in each single time step. This is in line with our notion that the attribute value corresponds to the underlying asset in option pricing, and the self-information, as a function of the relative frequency, can be understood as a price denoted in self-information as currency. Figure 10.4 shows a single time step of that process. The uncertainty about

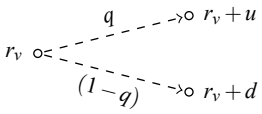

Fig. 10.4 Single time step in the development of self-information, analogous to the market price development in financial option pricing.

an increase or decrease of the relative frequency r_v is modeled by step size u (*upward move*, increase of the relative frequency) and by the probability q of an increase. Correspondingly, a *downward move* can be modeled by adding d. In line with [11], we assume u and d are chosen such that

$$d = -u \,. \tag{10.7}$$

Thus, all possible developments of the frequency for a fixed number of time steps form a lattice similar to the pricing lattice in Binomial Option Pricing. An example lattice is displayed in Fig. 10.5.

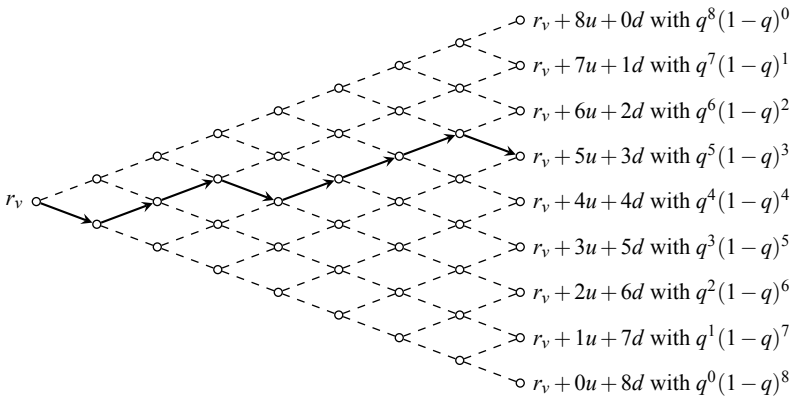

Fig. 10.5 Population model: discrete binomial stochastic process (random walk). The diagram shows all possible result states for the development of the relative frequency of an attribute value, starting with r_v, and their probabilities after eight time steps.

Even though this model is pretty simple, we can capture almost all developments of the frequency as long as we are able to choose the time steps small enough. A stagnation, for instance, can be represented by alternating up and down movements. Linear upward or downward trends of arbitrary strength can be modeled intuitively by combining upward or downward movements, respectively, with stagnation.

However, a direct analogy between market price and frequency development is not fully adequate, since market prices could increase without upper bound, but the relative frequency is defined only between zero and one. One way of dealing with the bounds would be forcing the next step of the random walk in a fixed direction, if the other direction led beyond a bound. This would simulate a stagnation at the margins of the domain. However, the approach has several drawbacks. For instance, once the upper bound is reached, any number of further upward movements would have exactly the same total effect as no further upward movements at all. In order to avoid that, we propose to transform the bounded domain of the relative frequency to an unbounded domain for the random walk. We have chosen the logit function for this transformation,

$$\operatorname{logit}(x) = \log \frac{x}{1-x} \,. \tag{10.8}$$

After running the random walk in the logit-transformed domain, we transform the value back to the frequency domain by means of the inverse logit function $\operatorname{logit}^{-1}$,

$$\operatorname{logit}^{-1}(x) = \frac{\mathrm{e}^x}{1+\mathrm{e}^x} \,. \tag{10.9}$$

The transformation to the unbounded domain allows us to rely on the same lattice process as known from Binomial Option Pricing. After the logit transformation, any number of movements in one direction is possible and exactly the same amount of movements in the other direction is necessary for compensation. Independent of the number of upward or downward moves, the outcome will remain within the bounds after the inverse transformation. Another nice property of the logit transformation is that the absolute changes in the relative frequency are the smaller the closer the level approaches the domain bounds. This captures a kind of base effect of very persistent individuals, who can be found in most heterogeneous populations.

The *information value* of an attribute value that will be exploited after $T > 0$ time steps can be computed by averaging the self-information over all possible relative frequencies, weighted with their respected probability of occurrence (right-hand side in Fig. 10.5). With $Q(n)$ being the probability of n upward moves,

$$Q(n) = \binom{T}{n} \cdot q^n (1-q)^{T-n} \,, \tag{10.10}$$

and $r_v^{(n)}$ being the relative frequency, taken from the result of the random walk with n upward moves,

$$\begin{aligned} r_v^{(n)} &= \operatorname{logit}^{-1}\Big[\operatorname{logit}(r_v) + nu + (T-n)d\Big] \\ &= \operatorname{logit}^{-1}\Big[\operatorname{logit}(r_v) + (2n-T)u\Big] \,, \end{aligned} \tag{10.11}$$

the expected self-information of the entire stochastic process after T steps is $\mathscr{H}_v(T)$:

$$\mathscr{H}_v(T) = -\sum_{n=0}^{T} Q(n) \log_2 r_v^{(n)} \,. \tag{10.12}$$

This measure of expected self-information accounts for fluctuations over time that are caused more generally by the society or the population, respectively, rather than by the individual data subject. Knowledge about an *attribute value* is the more valuable the higher the *self-information* of the attribute value becomes in the future and thus the smaller its relative frequency becomes in the population. Generalizing one step, knowledge about an *attribute* is less valuable the higher the *entropy* of the attribute is expected to grow (or remain) in the future.

Fig. 10.8 shows the development of a downward trend in a lattice diagram. Imagine an adversary who exploits technical attributes, such as browser or operating

system, of data subjects for re-identification. The parameters to be plugged into the process could be estimated from the dynamics of the market share of web browsers or operating systems in the population. The hypothetical development shows a clear downward trend in the market share of one particular browser, which had a dominant share before ($r_v = 0.7$). The downward trend might be due to data subjects switching to a competing alternative browser. The fewer data subjects use the formerly dominant browser, the higher is the value of the information that a specific data subject to be identified uses this particular browser. Thus, the expected self-information increases over time. Assuming that sufficiently accurate parameters can be estimated from historical observations, scaled down to a single time step, and predicted to remain valid for the next 100 steps, then we can continue the lattice shown in Fig. 10.8 in order to calculate the expected self-information after 100 steps. The development of the self-information for that time period is shown in Fig. 10.6.

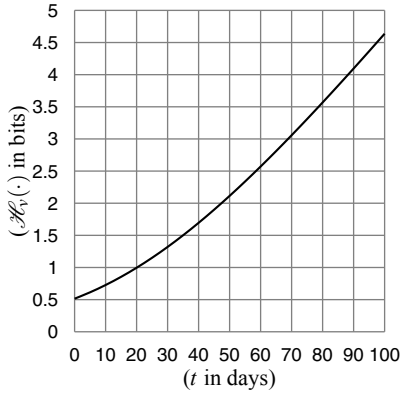

Fig. 10.6 Trend development of the expected self-information $\mathscr{H}_v(\cdot)$ for an attribute value with an initial relative frequency of the attribute value in the population $r_v = 0.7$, the probability of an increase $q = 0.3$, and the step size parameter $u = 0.1$. States that can be reached by a random walk in the first five steps are illustrated in Fig. 10.8.

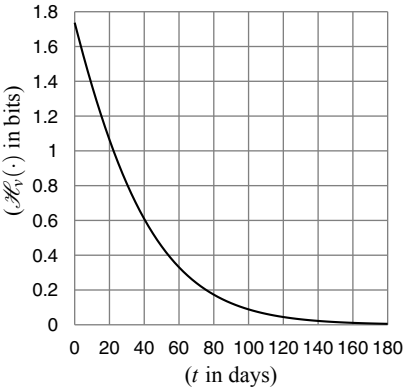

Fig. 10.7 Trend development of the expected self-information $\mathscr{H}_v(\cdot)$ for an attribute value. The parameters $r_v = 0.3$, $q = 0.7$, and $u = 0.1$ are chosen such that a positive trend can be observed for the attribute value. States that can be reached by a random walk in the first five steps are illustrated in Fig. 10.9.

Similarly, a browser or an operating system which has previously been used by a minority in the population ($r_v = 0.3$) may quickly become popular ($q = 0.7$). And therefore, the attribute value soon applies to a majority in the population. Thus, the expected self-information of that attribute value will decrease over time. We have outlined the first five steps in a lattice again, see Fig. 10.9, and continued the next 180 steps of the expected self-information in a diagram, see Fig. 10.7.

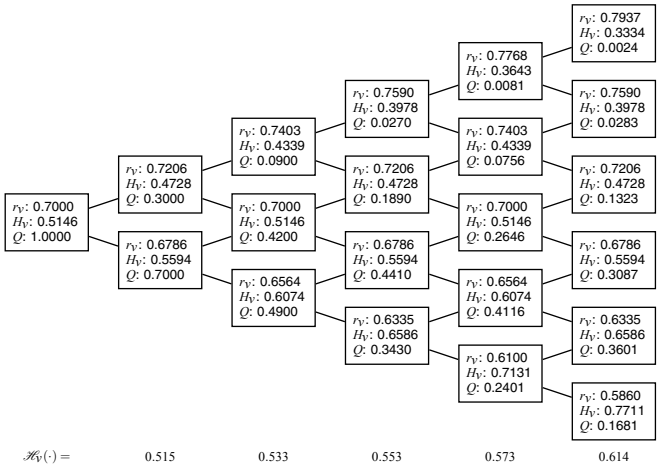

Fig. 10.8 First five steps of the development of the self-information with the parameters as described in Fig. 10.6. Each box represents a possible intermediate step of the random walk and for each step r_v denotes the relative frequency, H_v denotes the self-information, and Q denotes the probability. The expected self-information $\mathscr{H}_v(\cdot)$ after each time step is printed below each column of the lattice.

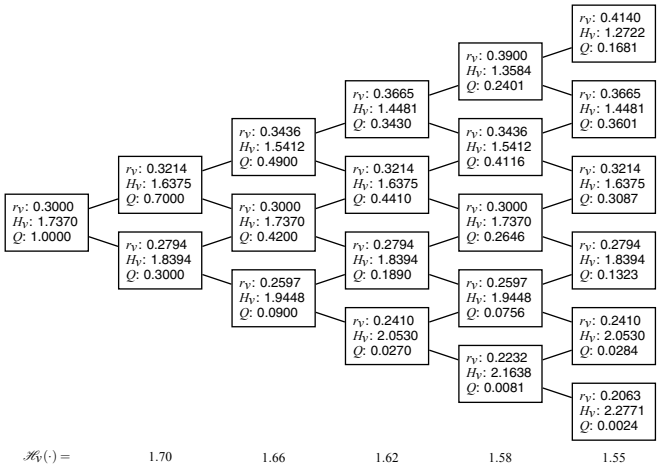

Fig. 10.9 First five steps of the development of the self-information for a positive trend development as described in Fig. 10.7. The notation is the same as in Fig. 10.8.

10.5 Valuation of Privacy Options

The main observation underlying our notion of privacy options is that disclosure of personal data and its exploitation does often not take place at the same point in time. In the previous section, we have argued that two sources of uncertainty drive the valuation of privacy options, and both can be modeled as independent stochastic processes. Now we will show how to combine the processes on the micro and macro level to obtain an inter-temporal measure of the value of personal data disclosure.

It is intuitively clear that the value of personal data (for re-identification of the corresponding data subject) depends on the self-information at the time of exploitation. The value is the lower, the lower the probability of a link between the data (i. e., attribute value) and the data subject is at that time. Thus, the probability of the link, modeled by Eq. (10.2) of the micro model, discounts the value of the (European) privacy option $\mathscr{V}_{\mathrm{Eu}}(T)$ at time T,

$$\mathscr{V}_{\mathrm{Eu}}(T) = x_1(T) \cdot \mathscr{H}_v(T) \,. \tag{10.13}$$

Recall from Eq. (10.3) that $x_1(T)$ denotes the first element of vector $\mathbf{x}(T)$, which holds the probability of a valid link.

The value of a privacy option depends on the parameters for the linkability model, namely the probabilities p and $\bar{p}$, and the parameters of the population development, namely the current relative frequency r_v of the disclosed attribute value, the probability of an upward movement in the random walk q, the step size u of an upward movement in the random walk, and the exercising time T of the option. For example, consider a privacy option with the parameters

$$\begin{aligned} p &= 0.95 \,, \quad & \bar{p} &= 1 \,, \\ r_v &= 0.5 \,, \quad & q &= 0.5 \,, \\ u &= 1.2 \,, \quad & T &= 100 \,. \end{aligned} \tag{10.14}$$

Observe in Fig. 14.4 that there is a substantial difference between the current value of personal data, i. e., the attribute value, and its information value for re-identification after several time steps. Further, it would be best to exercise the privacy option after seven time steps. Before reaching the seventh time step, the value of the option (solid line) is dominated by the increasing self-information of the attribute value (dotted line). Afterwards, the value diminishes due to the decreasing probability of a link between attribute value and data subject (dashed line).

$\mathscr{V}_{\mathrm{Eu}}(T)$ is the value of the privacy option, if it is exploited in the 'European' style, that is, the option can only be exercised when it expires. By contrast, American options can be exercised at any time between purchase and expiry. For privacy options, this means that personal data can be exploited at any time between disclosure and, for instance, the date of an obligation to erase the data. One can even think of the data being exploited more than once in the period of time. However, to allow for a better comparison, we normalize the valuation to exactly one exploitation. Thus, the value of an American privacy option $\mathscr{V}_{\mathrm{Am}}(T)$ is the average of the expected value at

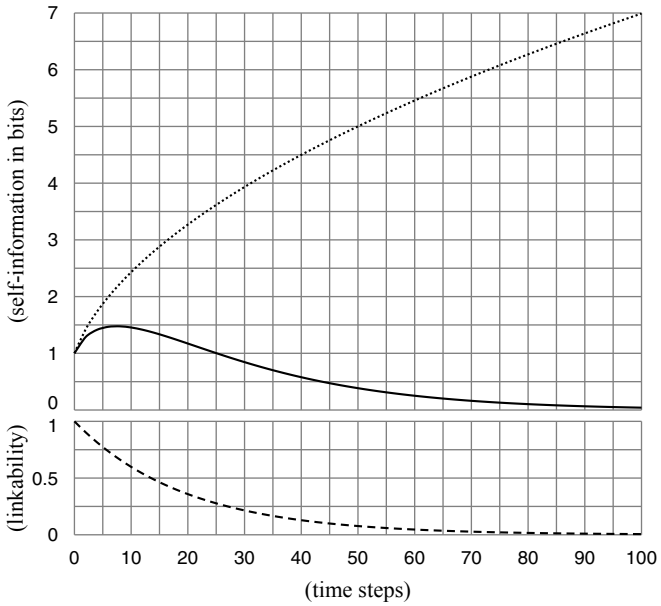

Fig. 10.10 Development of the privacy option value (solid line) with unit *bits*. The dashed line shows the corresponding probability of linkability with low probability of a change of the attribute value ($p = 0.95$), but once the attribute value has been changed, the previous linkable attribute value will never be recovered ($\bar{p} = 1$). The dotted line shows the development of the expected self-information of the attribute value (in *bits*). The distribution of the attribute value is assumed to remain the same ($r_v = 0.5$, $q = 0.5$), but the dispersion is high ($u = 1.2$).

each point in time between data disclosure and expiry of the option,

$$\mathscr{V}_{\mathrm{Am}}(T) = \frac{1}{T} \sum_{t=0}^{T} x_1(t) \cdot \mathscr{H}_v(t) \; . \tag{10.15}$$

If the attribute value is exploited several times between the attribute disclosure and the expiry of the privacy option, say, k denotes the number of exploits, then the value of the option is $\mathscr{V}_{\mathrm{Am}}(T)$, multiplied by k. One can also consider a weighted average to reflect a given prior on the possible time of exploitation.

10.6 Discussion of Results

A common assumption of privacy measures in the literature is that personal data, once disclosed, reduces the informational privacy of the data subject by its present self-information. Implicitly, this implies that the present self-information of disclosed data remain constant over time, at least until the data is exploited.

Table 10.1 Valuation of privacy options. Comparison of inter-temporal valuation with the self-information of the attribute value at present.

assumptions			valuation (in bits)			over/ under-valuation indicator[a]
stochastic process		option	point in time		time range	
linkability	population dev.	expiry date	present	future[b]	present–future[c]	
Fig. 10.1 (page 194)	Fig. 10.6 (page 199)	0	0.515			
		10		0.307	0.430	↘/↘
		25		0.293	0.342	↘/↘
		50		0.489	0.363	↘/↘
		100		1.070	0.569	↗/↗
Fig. 10.3 (page 196)	Fig. 10.6 (page 199)	0	0.515			
		10		0.565	0.595	↗/↗
		25		0.611	0.594	↗/↗
		50		0.596	0.603	↗/↗
		100		0.369	0.546	↘/↗
Fig. 10.1 (page 194)	Fig. 10.7 (page 199)	0	1.737			
		10		0.579	1.136	↘/↘
		25		0.237	0.664	↘/↘
		50		0.104	0.410	↘/↘
		100		0.020	0.231	↘/↘
Fig. 10.3 (page 196)	Fig. 10.7 (page 199)	0	1.737			
		10		1.066	1.517	↘/↘
		25		0.494	1.043	↘/↘
		50		0.127	0.654	↘/↘
		100		0.007	0.347	↘/↘

[a] Comparison between the present value, i. e., the present self-information of the attribute value, and the privacy option value, i. e., the expected self-information at a point in time, discounted by the probability of a link between attribute value and data subject. "↗" denotes that the present self-information underestimates the actual value, whereas "↘" denotes overestimation. The first arrow in this column refers to the "future" value and the other to the "present-to-future" value.

[b] This corresponds to a European option, which can be exercised at the expiry date.

[c] This is the value of the privacy option, if it is exercised at exactly *one* point in time between the date of disclosure and the expiry date. We assume that the point in time is randomly drawn from a uniform distribution.

Our examples show that the present self-information of personal data is only an appropriate measure for the information an adversary obtains when exploiting the data, if the disclosure and the exploit take place instantaneously. Otherwise, i. e., if time elapses between disclosure and exploit of personal data, the self-information at present can lead to both over- and underestimation of the 'true' value of the information passed over to the adversary.

Table 10.1 summarizes our findings by selected examples. It shows four privacy options derived from examples of the previous sections and their valuation with regard to expiry dates between 0 and 100. This corresponds to the situation where an attribute value is disclosed now and exploited either at the expiry date (column "future") or sometimes between now and the expiry date (column "present–future").

The resulting values of the privacy options, and thus the expected self-information of the underlying attribute value, is compared to the present self-information. Under- and over-valuations are indicated by arrows that point up or down, respectively.

In general, when the probability of a link between attribute value and data subject is uncertain, the value of personal data will be over-valuated by the present self-information, if the expected self-information is constant or decreasing over time. Undervaluations only occur, if an increasing expected self-information compensates the discount induced by the declining probability of linkability. In Table 10.1, this is the case for the first two privacy options, depending on the expiry date.

10.7 Conclusions and Outlook

In this chapter, we have motivated why and explained how option pricing theory can be useful for the valuation of informational privacy. In a first step towards this direction, we have proposed a very simple model that highlights the main features of our approach, namely the description of changes in each individual data subject's attribute values and the evolution of the distribution of attribute values in the population as two independent stochastic processes.

Once the realm of option pricing theory has been touched, possible extension and refinements are abundant. Most notably, it would be interesting to allow more than two attribute values in the state-space model, or to consider more than one attribute. This would not only allow to use the valuation results as guidance on which of a set of alternative attributes should be disclosed (if the data subject has a choice), but also to extract the self-information of combinations of attributes over time. Another obvious next step is to replace the binomial process with more appropriate processes. Ideally these processes should be validated with and calibrated to empirical data, e. g., from longitudinal population surveys. Replacing the discrete-time process with a continuous-time process could bring our model closer to (variants of) the Black–Scholes [6] formula, which promise closed-form solutions. This avoids computational effort when the number of time steps grows large (though at the price of additional assumptions). While the analysis in this chapter was strictly confined to expected values, one could also calculate and interpret other summary measures of the distribution functions over time. In particular small quantiles could be interesting to study (un)linkability with a security or risk management mindset by regarding the ϵ-worst case.

But there is more than just tweaks in the proposed framework: implementing true and conscious control of personal data in everyday social interactions is generally difficult. The fact that more and more social interactions happen in the digital sphere aggravates this problem substantially. Following in Baran's [4] footsteps, ideas of comprehensive privacy-enhancing technologies (PETs) have been conceived. Their vision is to cure the problems created by technology with more technology. So-called privacy-enhanced identity management is envisaged to assist people on deciding if, when, which, and at what price personal data should be disclosed. As

with every decision support system, this implies that several alternatives have to be evaluated and compared more or less automatically. And since most interactions do have consequences for the future, this evaluation would be incomplete if it does not consider time [22]. So privacy-enhancing technologies are an obvious field of application for our framework.

Existing blueprints for such PETs use so-called *privacy policies* to define how personal data should be handled (although enforcement of such policies against realistic adversaries is largely unsolved). Ideally, privacy policies are formulated in formal languages, which should support complex enough semantics to capture all relevant aspects of personal data disclosure—including time. Interestingly, a similar problem exists for modern financial contracts: nested derivatives quickly create a complexity in semantics that is manually intractable. The solution, again, lies in the intersection between finance and computer science. For example, Peyton Jones [31, 32] has proposed domain-specific languages to model complex financial constructs and enable their valuation over time. This can be seen as a generalization of classical option pricing theory. An interesting direction for future research is to adapt this to privacy policies and develop a formal language that can express aspects of time, and thereby generalize the valuation framework presented here.

Beyond direct applications in informational privacy protection through data avoidance, measuring the inter-temporal value of attribute values for linkability could also be useful in other contexts, even with opposite sign. It is conceivable that the data subject seeks to disclose as much information as possible to ensure clear identification in the future. This perspective will most likely matter when communicating bandwidth for attribute values is a scarce resource (e. g., through a hidden channel) and one must select those attributes which will be most informative later on. Moreover, although the exposition in this chapter was framed from the data subjects' perspective and targeted to protecting their personal data, the very same underlying ideas and valuation methods can also be useful for businesses to estimate the value of their customer databases. This is generally considered a hard task due to the intangible nature of personal data, so a new perspective might stimulate further advances in this area, too.

To conclude, although the idea of valuating privacy with option pricing theory sounds intriguing on paper, we have to recall that this framework is in no way a panacea. Many obstacles ignored in this exposition are likely to remain as serious limitations in practice: complexity, measurement problems, heterogeneous preferences, model mismatch, and bounded rationality, among others. So the confidence bands of our privacy metrics will most likely be loose in practice, but having a theoretically founded measurement method which can deliver some point estimates is certainly better than nothing at all.

Acknowledgments

This chapter incorporates some valuable comments by the anonymous reviewers for WEIS 2009. The first author was partially funded by the Research Council of Norway through the PETweb II project. The second author was supported by a post-doctoral fellowship of the German Academic Exchange Service (DAAD). Research leading to these results has also received funding from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement No. 216483.

References

1. Acquisti, A., Grossklags, J.: Privacy and rationality in individual decision making. IEEE Security and Privacy **3**(1), 26–33 (2005)
2. Acquisti, A., Varian, H.R.: Conditioning prices on purchase history. Marketing Science **24**(3), 1–15 (2005)
3. Amram, M., Kulatilaka, N.: Real Options: Managing Strategic Investment in an Uncertain World. Harvard Business School Press (1999)
4. Baran, P.: Communications, computers and people. Tech. rep., RAND Corporation, Santa Monica, CA (1965)
5. Berendt, B., Günther, O., Spiekermann, S.: Privacy in e-commerce: Stated preferences vs. actual behavior. Communications of the ACM **48**(4), 101–106 (2005)
6. Black, F., Scholes, M.: The pricing of options and corporate liabilities. Journal of Political Economy **81**, 637–654 (1973)
7. Blanchette, J.F., Johnson, D.G.: Data retention and the panoptic society: The social benefits of forgetfulness. Information Society **18**(1), 33–45 (2002)
8. Böhme, R.: A comparison of market approaches to software vulnerability disclosure. In: G. Müller (ed.) Emerging Trends in Information and Communication Security (Proc. of ETRICS), *LNCS*, vol. 3995, pp. 298–311. Springer, Berlin Heidelberg (2006)
9. Böhme, R., Koble, S.: Pricing strategies in electronic marketplaces with privacy-enhancing technologies. Wirtschaftsinformatik **49**(1), 16–25 (2007)
10. Clauß, S.: A framework for quantification of linkability within a privacy-enhancing identity management system. In: G. Müller (ed.) Emerging Trends in Information and Communication Security (ETRICS), *LNCS*, vol. 3995, pp. 191–205. Springer, Berlin Heidelberg (2006)
11. Cox, J., Ross, S., Rubinstein, M.: Option pricing: A simplified approach. Journal of Financial Economics (1979)
12. Daneva, M.: Applying real options thinking to information security in networked organizations. Tech. Rep. TR-CTIT-06-11, Centre for Telematics and Information Technology, University of Twente, Enschede, NL (2006)
13. Denning, D.E., Denning, P.J., Schwart, M.D.: The tracker: A threat to statistical database security. ACM Trans. on Database Systems **4**(1), 76–96 (1979)
14. Díaz, C., Seys, S., Claessens, J., Preneel, B.: Towards measuring anonymity. In: P. Syverson, R. Dingledine (eds.) Workshop on Privacy Enhancing Technologies, *LNCS*, vol. 2482. Springer, Berlin Heidelberg (2002)
15. Fischer, L., Katzenbeisser, S., Eckert, C.: Measuring unlinkability revisited. In: Proc. of Workshop on Privacy in the Electronic Society (WPES), pp. 105–109. ACM Press, New York (2008)
16. Fischer-Hübner, S.: Zur reidentifikationssicheren statistischen Auswertung personenbezogener Daten in staatlichen Datenbanken *[Towards reidentification-secure statistical data analysis of personal data in governmental databases]*. Diploma thesis, Universität Hamburg (1987). In German
17. Fischer-Hübner, S.: IT-security and privacy: Design and use of privacy-enhancing security mechanisms, *LNCS*, vol. 1958. Springer, Berlin Heidelberg (2001)

18. Franz, M., Meyer, B., Pashalidis, A.: Attacking unlinkability: The importance of context. In: N. Borisov, P. Golle (eds.) Privacy Enhancing Technologies, *LNCS*, vol. 4776, pp. 1–16. Springer, Berlin Heidelberg (2007)
19. Gordon, L.A., Loeb, M.P.: The economics of information security investment. ACM Trans. on Information and System Security **5**(4), 438–457 (2002)
20. Gordon, L.A., Loeb, M.P., Lucyshyn, W.: Information security expenditures and real options: A wait-and-see approach. Computer Security Journal **14**(2), 1–7 (2003)
21. Grossklags, J., Acquisti, A.: When 25 cents is too much: An experiment on willingness-to-sell and willingness-to-protect personal information. In: Workshop of Economics and Information Security (WEIS). Carnegie Mellon University, Pittsburgh, PA (2007). `http://weis2007.econinfosec.org/papers/66.pdf`
22. Hansen, M., Pfitzmann, A., Steinbrecher, S.: Identity management throughout one's whole life. Information Security Technical Report **13**(2), 83–94 (2008)
23. Herath, H.S.B., Herath, T.C.: Investments in information security: A real options perspective with Bayesian postaudit. Journal of Management Information Systems **25**(3), 337–375 (2008)
24. Huberman, B.A., Adar, E., Fine, L.R.: Valuating privacy. IEEE Security and Privacy **3**(1), 22–25 (2005)
25. Kelly, D.J., Raines, R.A., Grimaila, M.R., Baldwin, R.O., Mullins, B.E.: A survey of state-of-the-art in anonymity metrics. In: Proc. of ACM Workshop on Network Data Anonymization (NDA), pp. 31–40. ACM Press, New York (2008)
26. Li, J., Su, X.: Making cost effective security decision with real option thinking. In: Proc. of International Conference on Software Engineering Advances (ICSEA 2007), pp. 14–22. IEEE Computer Society, Washington, DC, USA (2007)
27. Matsuura, K.: Security tokens and their derivatives. Tech. rep., Centre for Communications Systems Research (CCSR), University of Cambridge, UK (2001)
28. Merton, R.C.: Theory of rational option pricing. Bell Journal of Economics and Management Science **4**(1), 141–183 (1973)
29. Odlyzko, A.: Privacy, economics, and price discrimination on the Internet. In: N. Sadeh (ed.) ICEC2003: Fifth International Conference on Electronic Commerce, pp. 355–366 (2003)
30. Ozment, A.: Bug auctions: Vulnerability markets reconsidered. In: Workshop of Economics and Information Security (WEIS). University of Minnesota, Minneapolis, MN (2004). `http://www.dtc.umn.edu/weis2004/ozment.pdf`
31. Peyton Jones, S.: Composing contracts: An adventure in financial engineering. In: J.N. Oliveira, P. Zave (eds.) FME 2001: Formal Methods for Increasing Software Productivity, *LNCS*, vol. 2021. Springer, Berlin Heidelberg (2001)
32. Peyton Jones, S., Eber, J.M.: How to write a financial contract. In: J. Gibbons, O. de Moor (eds.) The Fun of Programming. Palgrave Macmillan (2003)
33. Pfitzmann, A., Hansen, M.: Anonymity, unlinkability, undetectability, unobservability, pseudonymity, and identity management – A consolidated proposal for terminology. `http://dud.inf.tu-dresden.de/Anon_Terminology.shtml` (2008). (Version 0.31)
34. Purser, S.A.: Improving the ROI of the security management process. Computers & Security **23**, 542–546 (2004)
35. Schlörer, J.: Zum Problem der Anonymität der Befragten bei statistischen Datenbanken mit Dialogauswertung [On the problem of respondents' anonymity in statistical databases with dialogue analysis]. In: D. Siefkes (ed.) 4. GI-Jahrestagung, *LNCS*, vol. 26, pp. 502–511. Springer, Berlin Heidelberg (1975)
36. Serjantov, A., Danezis, G.: Towards an information theoretic metric for anonymity. In: P. Syverson, R. Dingledine (eds.) Workshop on Privacy Enhancing Technologies, *LNCS*, vol. 2482. Springer, Berlin Heidelberg (2002)
37. Shannon, C.E.: A mathematical theory of communications. Bell System Technical Journal **27**, 379–423, 623–656 (1948)
38. Soo Hoo, K.J.: How much is enough? A risk-management approach to computer security. In: Workshop on Economics and Information Security (WEIS). Berkeley, CA (2002). `http://www.sims.berkeley.edu/resources/affiliates/workshops/econsecurity/`
39. Steinbrecher, S., Köpsell, S.: Modelling unlinkability. In: R. Dingledine (ed.) Workshop on Privacy Enhancing Technologies, *LNCS*, vol. 2760, pp. 32–47. Springer, Berlin Heidelberg (2003)

40. Sweeney, L.: *k*-anonymity: A model for protecting privacy. International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems **10**(5), 571–588 (2002)
41. Tóth, G., Hornák, Z., Vajda, F.: Measuring anonymity revisited. In: S. Liimatainen, T. Virtanen (eds.) Proc. of the Ninth Nordic Workshop on Secure IT Systems, pp. 85–90. Espoo, Finland (2004)
42. Willenborg, L., De Waal, T.: Statistical Disclosure Control in Practice. Springer, New York (1996)
43. Wolfers, J., Zitzewitz, E.: Prediction markets. Journal of Economic Perspectives **18**(2), 107–126 (2004)
44. Xiaoxin, W., Bertino, E.: Achieving *k*-anonymity in mobile and ad hoc networks. In: Proc. of IEEE ICNP Workshop on Secure Network Protocols, pp. 37–42. IEEE Press, New York (2005)

List of Symbols

General Symbols

v	attribute value
t	discrete point in time
T	total number of time steps
$\mathscr{V}_{\mathrm{Eu}}(T)$	value of a European privacy option at time T
$\mathscr{V}_{\mathrm{Am}}(T)$	value of a American privacy option at time T

Timed Linkability Process (Micro Model)

p	probability of an attribute which is currently linkable to a data subject to remain linkable in the next time step
$\bar{p}$	probability of an attribute which is currently not linkable to a data subject not to become linkable in the next time step
$\mathbf{x}(T)$	state vector in the state-space model at time T
$x_1(T)$	probability of a valid link between disclosed attribute value and data subject after T time steps
$\mathbf{A}$	state transition matrix of the state-space model
$a_{i,j}$	elements of $\mathbf{A}$

Population Development (Macro Model)

r_v	relative frequency of attribute value v in the population
H_v	self-information of attribute value v
q	probability of an upward move in the random walk
u	step size of an upward move in the random walk
d	step size of an downward move with $d = -u$
$Q(n)$	probability of n upward moves in T moves in total
$r_v^{(n)}$	relative frequency of attribute value v after randomly walking through T time steps of which n are upward moves
$\mathscr{H}_v(T)$	expected self-information of attribute value v after T time steps