Chapter 14
How to Explore Consumers' Privacy Choices with Behavioral Economics

Sören Preibusch

14.1 Introduction: The Economic Understanding of Online Privacy Beyond Data Protection

14.1.1 The Convenience of Online Consumption

The Web has enabled consumers to access and share an unprecedented amount of information, quickly, conveniently, and cheaply. Companies have embraced new information and communication technologies and moved offline phenomena such as shopping, entertainment, or social networking into the virtual realm. With the advent of the Web, services of a new kind have emerged, such as web search or blogging. In the United Kingdom, digital value creation accounted for 7.2 % of the gross domestic product in 2010 [1]; the share of the digital economy is predicted to continue growing rapidly to reach $4.2 trillion in the G-20 nations by 2016 [2].

Retailing is among the industries that have been fundamentally disrupted by the Internet. For products such as clothing and shoes, Internet sales in the United Kingdom now account for 11.3 % of total sales, with 17.6 % growth year-on-year [3]. Average weekly spending online in May 2014 was £727.5 million [3]. In parallel, the high street is expanding into multi-channel retailing to combine the benefits of an online and offline presence. Much of the reconfiguration of the value chain has happened behind the scenes, but there is a tangible impact for consumers as well. Being greeted by name and receiving personalized product recommendations used to be a distinction of up-market boutiques. Today, it is the sign of mass-personalized online shopping, and one of the ways in which a "data culture" is implemented. Harnessing the power of ubiquitous computing enables organizations to turn data into fuel for insight [4]. Given that they are a valuable resource that

S. Preibusch (✉)
Microsoft Research, Cambridge, UK
e-mail: mail@soeren-preibusch.de

S. Zeadally and M. Badra (eds.), *Privacy in a Digital, Networked World*,
Computer Communications and Networks,
DOI 10.1007/978-3-319-08470-1_14

touches all aspects of society and shapes new forms of production and consumption, personal data are said to be "the new oil of the Internet and the new currency of the digital world." [5].

Consumers enjoy the resulting personalization [6, 7: Sect. 21.3.1]. It reduces the time and effort they have to spend on finding and judging products. For companies, personalization is a powerful tool since it allows lock-in and efficient customer value extraction in fiercely competitive markets. Retailers create and satisfy needs their customers were not yet aware of. Amazon indicates that around 30 % of all purchases result from recommendations [8]; 27 % of European consumers indicate they have bought a product in the past twelve months because it was recommended by the retailer [9]. In short, data-driven personalization works, and benefits both consumers and companies.

14.1.2 Monetization of Personal Data and Mainstream Privacy Concerns

The monetization of personal data as a commodity, through targeted advertising or otherwise, also allows many expensive services to be offered free of charge [10]. It is estimated that UK consumers enjoy an annual surplus of £5 billion from free online content, or twice what they pay to access the Internet [2]. The World Economic Forum observes that "in practical terms, a person's data would be equivalent to their 'money'" [11], and foresees that consumers could control, manage, and exchange their data as they do with cash in their bank account. The European Data Protection Supervisor similarly observes that "personal information has become a form of currency to pay for so-called 'free' online services" [12]. From this follows a close interplay between data protection and consumer protection, and that privacy cannot be achieved through technical means alone.

Consumers experience invasions of privacy as the flip-side of data-powered high-value services, including personalization [13]. Their development and ongoing provision requires far-reaching collection of data and its long-term storage. Until recently, only a few privacy-aware consumers and data protection advocates were aware of the broad consent obtained by businesses through their privacy policies. Post-Snowden, privacy issues are now making headlines in mainstream media [14].

We live in a networked world where ubiquitous Web tracking of consumers and planet-scale government surveillance of citizens are not capabilities but realities. The resulting privacy challenges are calling for privacy-enhancing technologies (PETs), policies, and practices. As outlined above, advances in technical data protection are only part of the picture. The preferences and the incentives behind the choices of companies and their customers are equally important. Economics provide the tools for their study. Privacy failures have been caused at least as often by bad incentives as by bad system design. Ignoring potential users' privacy needs leads to PETs failing in the marketplace despite good engineering [15] (e.g., fully

anonymous search engines). Economics provide tools and theories to reason about privacy failures, to suggest remedies, and to positively understand superior privacy practices as the source of competitive advantage.

14.1.3 Studying the Economics of Privacy

As a discipline, the consumer-centric economics of privacy study the value that consumers attach to items of personal information. The objects of analysis are exchanges of personal data. These happen in an environment where data protection is an unquestioned, constitutional, and human right, that provides minimum protection guarantees and remedies.

Other disciplines, beyond the scope of this chapter, challenge the assumption of guaranteed privacy. They discuss the welfare or political economy of information. Privacy as a right is questioned, by establishing its worth for society, or its impact on markets' efficiency. Recently, the blogosphere has restarted the debate under the concept of "post-privacy" [16].

The economics of privacy recognize that personal data have been commodified into a tradable asset. This empirical reality is embraced by studies of markets for personal information and of the behaviors of companies and consumers on such markets. Like many other markets, the market for personal information is far from perfect. It is a defining trait of behavioral economics to embrace these imperfections and make them the object of study: information asymmetries, barriers to entry and exit, externalities, monopolies, and oligopolies. With a focus on actual behavior observed in market players, research is descriptive rather that prescriptive. Experimental designs are inspired by theory, but the evidence need not be rationalized post hoc.

In the tradition of behavioral economics, consumers' reactions to systematically manipulated experiment conditions are observed. The experimental stimulus is an intervention and allows establishing causal relationships, for instance, between data-item sensitivity and consumers' propensity to protect those data. The influence of confounding factors (e.g., visual web site design, trust in companies and brands) can be abstracted away when held constant across treatments. These can be studied through research into human–computer interaction that complements economics experiments in deriving an overall successful user experience.

14.1.4 Supply and Demand for Privacy

On the supply side of data markets, barriers to entry are mostly immaterial: whereas up-front investments into data centers become dispensable when the cloud provides compute/storage infrastructure to new entrants without fixed costs, incumbents profit from previously collected data records. They can improve their offerings

through 'learning by doing,' leading to economies of skill. The data records accumulated by a company are an intangible asset [17]. On the demand side, data subjects are typically unable to observe how their data is used and potentially shared and misused, creating an information asymmetry to their detriment. At the same time, poor data portability between alternative services and positive network effects create switching costs; the resulting lock-in makes contractual hazards more likely. Start-ups such as Mydex position themselves as intermediaries to profit from the market frictions by offering personal data vault services. The aim of regulation and enforcement is to create rules for the market of personal information that protect the consumer and increase efficiency and social welfare [18]. An example of such an initiative is the "midata" vision put forward by the UK government: consumers should be given access and insight into their personal data, including usage logs, to migrate these to an alternative supplier if desired [19].

14.1.5 Consumers' Choices for Price–Privacy Trade-Offs

Consumer empowerment relies on their effective ability to transact with a company that suits their preferences. Potential customers have the choice between alternative suppliers that compete on price and non-price attributes. This is true for electronic markets as well as traditional markets. In the grocery store, shoppers not only consider the price tag for a bag of apples, but also the quality of the produce, whether it is grown locally, and farmed according to ecological standards. In electronic retailing, the non-price attributes of a company's offering include its privacy practices.

When consumers engage in transactions that involve exchanges of goods or services, money, and their personal data, they may choose to withhold some of their details. The resulting decrease in service quality or an increase in price is the cost they have to bear to maintain their privacy. Behavioral studies allow measuring a lower bound for the value of privacy by observing consumers' willingness to pay for avoiding data collection or other invasions of privacy.

The issue that researchers and practitioners are facing today is the lack of studies that provide reliable and valid insight into consumers' privacy concerns and behaviors. Looking back, this lack can be explained by the relative recency of the field, even within the study of human–computer interaction. However, looking forward, the ability of new studies to deliver actionable insights hinges on a methodological reboot.

14.1.6 Structure of This Chapter

As a solution, this chapter aims at equipping researchers, practitioners, and policymakers with the tools and the evidence to understand consumers' privacy

behaviors. I begin by explaining why experiments rather than surveys or hypothetical choices are needed for delivering valid insights to decision makers. After an exhaustive review of the existing empirical evidence into the value that consumers attach to their privacy, I explain the methodological requirements of valid privacy experiments and offer practical advice for conducting privacy choice experiments. The research presented in this chapter will help in developing privacy-enhancing solutions and policies that meet consumers' needs.

14.2 Surveys Versus Experiments into Privacy Behaviors

14.2.1 Divergence of Privacy Attitudes and Behavior: A Fresh Look at the Privacy Paradox

When surveyed about data protection issues, consumers repeatedly report high concerns about their information privacy [20]. In the 2011 Eurobarometer on data protection, 70 % of respondents, representatively sampled from the EU population were concerned that their personal data held by companies may be used for a purpose other than that for which they were collected [13]. At the same time, the online population increasingly engages in online activities deemed privacy-threatening, namely online social networking [21]. Concern reported in surveys is higher than what can be inferred from observed real-life behaviors.

This discrepancy between attitudes and behaviors, called the *privacy paradox*, has mainly been described with regards to the interplay between privacy and online personalization: consumers want to enjoy the benefits from profiling, but they do not want to be profiled [7]. Disclosure on online social networking sites has also been described as a privacy paradox [22], although the combined horizontal and vertical relationships amongst users, and between users and the platform operator respectively, is harder to interpret.

Establishing the privacy paradox requires observing a divergence of privacy preferences and behaviors within the same population or between two representative samples thereof. Experimental studies provide such an opportunity to observe stated privacy attitudes and actual privacy-related behavior within subjects. In laboratory experiments, participants who reported high privacy concerns exhibited behavior that diminished their information privacy [23]. Looking at information-only transactions (Sect. 14.3.1), a similar discrepancy has been observed: they actually provided more information than they had previously stated to be willing to share [24].

However, other experiments do not necessarily support the notion of a paradox: individuals with stronger privacy concerns were found to place higher values on privacy in information-only transactions [25]. In a 2013 experiment on privacy in web search, participants' stated willingness to pay for privacy-enhancing features did not explain their behavior, but it also did not contradict their actual choices. Both variables were recorded as part of the same experiment [26].

It has also been argued that disclosure seemingly diverging from attitudes may be explained by strong beliefs in the confidentiality of the disclosed data. Divergence would originate in experimenter trust, framing effects, or deception used by the experimenter [27]. It seems that the supposed paradox would be an artefact, or mode effect, originating in measuring the varying behaviors and attitudes with experimental or survey methodologies respectively. A deeper understanding of privacy concerns and behaviors, therefore, requires valid survey instruments as well as behavioral studies.

14.2.2 When to Use Privacy Surveys and When not to

Even if existing empirical studies do not necessarily support the notion of a privacy paradox, they also show how behavioral intent or self-professed behavior from a survey has little predictive power for actual behavior. A recent study commissioned by Microsoft for Data Privacy Day 2014 serves as an example. The survey specifically recruited "technology elites," characterized, for instance, by self-identifying as influencers on technology and as early adopters of new technology. Amongst the 1,075 respondents in the United States and in the European Union, more than three quarters indicated that they read privacy policies before clicking "accept"; almost a quarter even indicated they read the terms in full [28]. However, web server logs of actual privacy policy visits suggest that this proportion is lower by several orders of magnitude, even amongst advanced users. Only a small minority of Web users actually read the privacy policies of sites they interact with.

The lack of commitment is a major reason why statements about behavior do not reflect real choices. In the absence of real-world transactions, a survey creates an artificial context, influenced by mode effects [29]. The incentives for respondents on how and what to reveal are different from real transactions, typically in a way that works against truthful revelation. One of the biases in a survey is respondents' tendency to give socially desirable answers. Furthermore, surveys have a "research appeal," which makes respondents disclose more information about themselves [30]. Yet, neither privacy nor money are ultimately at stake in a survey. In consequence, predictive power and ecological or external validity are largely reduced.

Despite the inability of surveys to be a reliable and valid predictor for consumer behavior, they do have their rightful place. In the early design stages of an experiment, low-cost surveys can help identify questions of interest an experiment should focus on. In the area of privacy economics, for instance, a pilot study could incorporate a Conjoint Analysis that helps researchers making a more substantiated decision about price differences in an experiment (Sect. 14.2.3). Once an experiment is about to be deployed, screening questionnaires can help in recruiting suitable participants, for instance when sampling a population with high privacy concerns. Survey elements are also crucial in complementing an experiment session as entry- and exit-questionnaires, and comprehension tests typically before the

experiment. They deliver insights into the demographics of the sample, their attitudes and personality—in particular when well tested, validated instruments are used. A recent review of survey instruments to measure privacy concerns provides guidance on which methodology to use and how to deploy scales for privacy concern [31].

In summary, we acknowledge that privacy attitudes and privacy behaviors do not always agree. The methodological consequence is to measure both in their own right and with their dedicated procedures. Preference should be given to experimental procedures when studying privacy behavior; surveys lend themselves to assess attitudes. Both approaches must be subjected to the same scrutiny of reliability and validity [31]. In reviewing previous research into the behavioral economics of privacy, I therefore proceed by the methodology used, distinguishing between survey-like approaches (Sect. 14.2.3) and approaches relying on experiments (Sect. 14.3).

14.2.3 The Failure of Hypothetical Privacy Choices

The problems of privacy surveys also apply to survey-like methodologies when participants make hypothetical choices. Often, these works are erroneously labeled as experiments. In a typical survey-like procedure, participants are confronted with scenarios and asked whether they would be concerned about their privacy in such a scenario.

Sometimes a single scenario is used. For instance, participants are asked to imagine a university alumni association shares its members' names, contact, and other information with a car insurance company for a 30 % discount [32]. Other single-scenario work claimed to trial membership in an online bookstore, for which some personal information would be necessary, in exchange for some discount. The amounts of both varied by treatment [33]. In these two studies, the respondents had to report how happy, satisfied, or concerned they would be with the deal presented in the scenario. In a similar vein, participants have been presented with a simulated online shopping web site and asked whether they would intend to buy from that site [34]. Although this is a slight improvement, because the participants can experience the stimulus (i.e., the web site), it still remains a hypothetical choice.

Another strand of hypothetical choice studies presents participants with multiple scenarios: potential job-seeking university students were given four channels to advertise their talents and job interests, including three web sites [35]. All varied by privacy intrusion and chances of success. Participants indicated their preferred option amongst these four. There is, however, an undeniable framing bias: participants are given the impression they are supposed to have different preferences for different types of data collectors, even if the question of advertising their professional skills never occurred. The design is as flawed as asking a vegetarian whether they prefer their steak rare or well done.

Hypothetical choices are also the standard for studies using conjoint analysis. Conjoint analysis tries to decompose the joint influence of several factors on a respondent's preference for one option. There are several variables (e.g., price), each with multiple categorical levels (e.g., 1, 2, 5 euro). Several stimuli or scenarios are created by systematically combining different levels across the attributes. These can be aspects of privacy intrusion, monetary incentives, or prices. To keep the total number of stimuli manageable, an orthogonal design is often preferred over a full factorial design: one does not present all possible combinations of attribute values to the participants. Instead, multiple attributes are varied at once. Participants then rank the scenarios in decreasing appeal [36, 37]. Their rankings have no impact on payoff, but their responses were still interpreted as if they were valid.

An alternative to ranking multiple alternatives is to present scenarios in pairs; participants indicate the preferred one. When combined with an outside option ("neither"), the responses can be analyzed with choice-based conjoint analysis. Binary logistic regression can also be used. In one study, participants indicated the one preferred out of two web pages, in the absence of an outside option [38]. The stimulus was only the mock-up of a single page, not a full, interactive web site. Again, participants' payoff was independent of their choices.

14.3 Review of Privacy Choice Experiments

In contrast to surveys, choice experiments put participants in a decision-making context where their preference for one of the alternatives will have an impact on their lives. In privacy economics, decisions are made as part of a transaction between a consumer and a company (Fig. 14.1). A typical transactions involves the flow of money, personal information, and goods or services. Money may flow in either direction: customers pay a price; companies can offer vouchers. Composite transactions that include the exchange of money, personal data, and goods are common in online shopping. Money may be absent for services provided free of charge (e.g., web search) and goods are not provided when the information receiver collects data from the consumer in return for data (e.g., prize draws). Information-only transactions are observed when consumers volunteer data without compensation, such as in a poll.

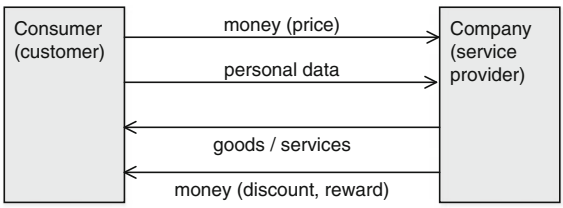

Fig. 14.1 Companies and their customers exchange personal information and goods or services when transacting online

In replicating composite transactions in a laboratory or in a field study, one can measure the monetary value that consumers attach to pieces of personal information. In this section, I review existing empirical studies, grouped by transaction type. Given the paucity of true experiments to date, this is an exhaustive review.

14.3.1 Experiments into Information–Money Exchanges

14.3.1.1 Experiment Design Varieties

In information–money exchanges, consumers receive payments in return for disclosing personal details. Three varieties of experiments have been conducted regarding *information–money exchanges*: incentivized disclosure for an unstated purpose [39], for actual or decoy research purposes [25, 40], and for a deceitfully stated and not implemented purpose [41, 42]. Only the most recent research on information-transactions, which is also the most robust in its design, did not involve deception, but told the participants up-front that the experiment was studying their privacy preferences when browsing the Web [43].

14.3.1.2 Measuring Willingness to Pay

Different mechanisms have been used to elicit willingness to pay. A reverse second-price auction is the most common [39, 41, 43]: participants put in their bids, stating how much they would want to be paid for releasing a specific item of personal information. The winner is determined by the lowest bid, and will be paid the second-lowest bid for disclosing his or her data. Participants have an incentive to bid their true valuation: winning with a bid below their true valuation will make them sell at a loss; asking for too much compensation puts them at risk of not being considered at all. It has also been noted that the auction mechanism is easy to implement and easy to explain to participants [43]—which are important practical considerations.

Amongst the auction mechanisms, the recent work into valuing the privacy of browsing behavior is most interesting and relevant by its design. Recruited though a survey on a major web portal in Spain, 168 participants installed a browser plugin, which invited them at intervals to place a bid for selling personal information relating to the web site they currently viewed [43]. In addition, bids were also solicited for various items of personal information detached from a browsing context. The median bid value across data categories was much higher for context-independent data (€25) than for context-dependent data (€7). A single piece of data was valued similarly to ten pieces of the same kind. A follow-up questionnaire further indicated that users approved of exchanging their data in return for improved service, but refused to have their data monetized by those same providers [43].

As an alternative to auctions, fixed amounts of money have been used, followed by observing whether and how many participants would accept the offer. Participants are asked whether they would disclose their data for a given compensation, such as $2 (in a field experiment, [40]) or for amounts varying between $0.25 and $1 (in a laboratory experiment by the same authors, [25]). The spreads can also be larger, varying between SG$1 and SG$9, equivalent to $0.60 and $5.40 (in a field experiment, [44]). In the latter experiment, participants were invited to a web form, disguised as a consumer research survey into mobile devices, which required items of personal information [44]. The number of data items, the compensation for completing the entire form, and the presence of privacy assurance through a statement or statement plus seal, were manipulated in the different treatments. Privacy assurances and monetary incentives both had a positive influence on disclosure [44], although disclosure was already very high without any of the two.

14.3.1.3 Volunteering of Personal Information

In one of my own studies, we explored the lower bound of what companies would need to pay their customers to stimulate data disclosure. Deployed as an online experiment, we recruited 1,500 web users to complete a form asking for ten items of personal data [30]. Items spanned identity and profile information of varying levels of sensitivity, such as first name and date of birth, as well as health and spending habits. The web form was chosen for its role as the primary mechanism to collect personal data from individuals on the Web. We manipulated the number of mandatory fields (none vs. two out of ten) and the compensation for participation ($0.25 vs. $0.50) to quantify the extent of over-disclosure, the motives behind it, the resulting costs and privacy invasion. A fully rational participant, eager to minimize her exposure and effort, would be expected to leave blank all fields but the mandatory. Quite the opposite, we observed a high prevalence of deliberate and unpaid over-disclosure of data. Participants regularly completed more form fields than required, or provided more details than requested. For instance, when asked when they had last spent $100, some not only provided the date, but also the purpose of the expenditure. We saw that more than two thirds of participants volunteered their date of birth and other personal details; disclosure rates, which were later confirmed in another study. Through careful experimental design, we verified that participants understood that additional data disclosure was voluntary, and the information provided was considered sensitive.

The experiment provides evidence that companies may be able to collect personal details without compulsion or offering incentives, but instead by leveraging consumers' psychological drivers towards completing optional web form fields. Through two manipulations, we benchmarked the efficiency of compulsion and incentives against volunteering. First, when two of the ten fields were marked as mandatory, disclosure rates for the remaining optional fields dropped. A company that forces its customers to complete certain fields reduces the amount of

volunteered details. Second, monetary incentives for completing those same fields yielded positive spillover by increasing revelation ratios for other optional fields. Both effects are statistically highly significant [30]. The effects suggest that the transaction is not perceived as a market transaction but instead as a social exchange, that can be broken (through compulsion) or reinforced (through gifting).

14.3.1.4 Challenges for External Validity

In information–money exchanges, money compensates consumers for their loss of privacy. As with other setups, it is, therefore, important that participants incur a true loss of privacy, which is typically achieved through data verification and with the transaction having an impact beyond the protected realm of the study. However, some experiments have tried to create personal information artificially in a laboratory context: "the experimental instrument separated subjects from their natural identities and allowed information and privacy values to emerge endogenously in the laboratory" [27: 8]. These studies confound personal information with the economic notion of private information [45]. By design, the information to be disclosed is no longer personally identifiable. Such studies, therefore, do not measure an invasion of privacy, but participants' avoidance of embarrassing or socially undesirable disclosure.

As a more general critique, the absence of goods or service consumption in pure information transactions creates an incentive structure which resembles paid surveys. Although it may be interesting to estimate the minimum amount of money payable to consumers to reveal some personal or demographic information, this price tag does not implement the purpose-binding of personal information. This is a systematic flaw: if a purpose is unstated, participants are tempted make up a purpose in their mind in an uncontrolled manner. If instead a purpose is stated, but not implemented, participants are deceived. Even when researchers truthfully state and implement data usage, participants trust the researchers and they are biased towards helping research, resulting in personal data disclosure for low monetary values.

Information-only transactions are not happening at large on the Web today. One should be cautious not to generalize the results of information-only transactions to composite transactions. The incentive structures in an online shopping or social networking context are quite different.

14.3.2 Experiments into Information–Service Exchanges

In the early 2000s, laboratory experiments examined consumers' willingness to disclose items of personal information in return for a personalized shopping experience; in this section, I review two early studies. Personalized shopping is one scenario where the privacy paradox could be observed (Sect. 14.2.1).

The common design for *information–service exchanges* experiments is as follows: while participants proceed through an online shop, they can unlock personalization features by disclosing additional personal information. Importantly, participants who disclose more do not get higher monetary payoffs, but may enjoy personalization benefits. Payments are made to the participants, as show-up fees and subsidies to purchases, but these do not depend on the amount or kind of personal information revealed. Their aim is to increase overall participation and purchase ratios. Payments were unconditional [46] or—which is less preferable—distributed in a lottery amongst all buyers in the experiment [23]. Participants were also informed that the experiment studies a personalization scheme; this framing has been criticized for biasing participants towards voluntary disclosure in an attempt to help research [27].

In the experiment by Spiekermann et al. [23], 171 student participants visited a web site to shop for digital cameras or winter jackets, choosing from a broad assortment of 50 and 100 models respectively. While shopping, they could interact with an "anthropomorphic 3-D shopping bot that assisted participants" through a sales dialogue involving 56 questions relating to product attributes, usage, but also personal questions (e.g., "What is your motivation when taking photographs?" or "How important are trend models to you?") [23]. Responses to these questions allegedly served to compile a ranked list of the top ten products. The authors do not report whether this ranking was truly dependent on participants' responses.

In a later experiment by Kobsa and Teltzrow [46], 52 student participants could browse an online book store. A series of 32 questions spread over nine pages would help them navigate the assortment. Each page displayed a book counter, decreasing from 1 million to 50 matching books. However, the matching was an illusion, created by decreasing the counter. The participants ignored the fact that the final selection was predetermined by the authors based on assumed general appeal, and independent of participants' responses [46]. Although all personalization questions would seem plausible in a book store context, they were far more intrusive than in the shopping bot study [23]. For instance, participants were asked for political and religious interests, their preferences for erotic literature and interest in certain medical subareas [46]. All questions featured a "no answer" option. Interestingly, the authors implemented an ID check on the buyers: this may have been the first time truthful revelation of personal data was enforced in a laboratory context.

In another strand of research, observational studies and surveys have tried to measure social capital returns from disclosing personal data online, in particular on online social networking sites [47]. It has been argued that participation in a social networking site would indeed negatively impact on privacy; however, usage would also result in so strong a gratification for the users to the extent that it warrants self-disclosure [48]. Participating in a social network despite privacy concerns would not necessarily be a privacy paradox (Sect. 14.2.1). This stream of work opens up towards non-economic, but social exchanges.

14.3.3 Experiments into Information–Money–Goods Exchanges

The body of research into the value of personal information as part of goods transactions stands out by its paucity. Compared to the ever growing number of commercial opinion polls and academic surveys there are surprisingly few experimental studies into privacy economics. A recent literature review [8] only identified the work by Tsai et al. [49] and my own 2009 DVD experiment [50] as experimental studies; it also included an information-only experiment [44], which has already been discussed. Another comprehensive literature review into the behavioral privacy economics observed that such experimental designs were rare [51: Sect. 4.2.2]. Their comprehensive enumeration only included the works by Beresford et al. [52], Tsai et al. [49], Gideon et al. [53], and Jentzsch and Giannetti [54]. At the time, the latter was still in the design phase; in its current stage, it mixes the concepts of personal and private information [54]. Concordantly with the categorization used here, another featured experiment was classified as an information-only transaction [39].

Besides my own three experiments to date [50, 51, 55], detailed below, it therefore seems that the body of experimental works to study the privacy economics of composite transactions is limited to two studies: [49] and [53]. These works have shared authorship; their designs are similar and they build on one another. Both studies invited participants to a laboratory, where they shopped online and considered privacy issues on a competitive market. The experiments feature a field component in the form of external order fulfillment. The later, more sophisticated design is described first.

14.3.3.1 Experiment "Vibrators Versus Batteries"

Tsai et al. [49] consider consumers' trade-offs as they choose between competing sellers for the same good that differ by price and privacy. The authors have republished their findings several times. The following analysis is based on their initial report [49], which also gives the most detailed account of the experiment procedures.

The goal of the experiment "was to determine whether the prominent display of privacy information in search engine results causes privacy-concerned users to take privacy into account when making online purchasing decisions" [49]. The study was further aimed at determining "whether privacy-concerned users are willing to pay a premium to make their purchases from the more privacy-friendly merchants" [49]. As part of the study, 48 participants were invited to a laboratory session, spread across three treatments, followed by an exit-questionnaire. Participants were paid a show-up fee of $45.

The distinction between privacy-friendly and privacy-unfriendly was created and made salient in the laboratory through icon-annotated result listings in a product

search engine. The privacy rating effectively took four levels, from missing to zero, two, and four out of four stars. Ratings of one or three stars were not encountered in the study.

Participants were instructed to perform searches for a series of products; the search terms were prescribed to match a single item sold by several retailers. Products included batteries and a vibrator, which can be considered as prototypical examples of office supplies and sex toys. In a preceding exploratory survey, these product categories were identified to engender low to medium concerns and high to medium purchase likelihood respectively.

The product search result was the main stimulus. The appearance of the first four results was controlled. It is reported that the order of the results was such, that lower rank was associated with a higher price and a better privacy rating [49: 12, 35]. The prices in the experiment were not controlled; the original, varying retail prices by the merchants were used.

The results were the following: guided by the visual four-star privacy rating, participants were willing to pay a premium of around $0.60 when shopping for vibrators and batteries respectively. The actual price differences between the different retailers varied. One cannot conclude that consumers paid $0.60 to shop with a well-rated merchant. In a control treatment, the rating was relabeled as "Handicap Accessibility" instead of "Privacy Report." Participants still preferred to pay higher prices to shop with a four-star merchant, although the difference in average prices was not significant in this case [49].

Although this study implemented real purchase transactions during which participants paid with their own money and released their personal credit card details to a commercial entity of their choice, they could provide a dummy shipping address instead of their own postal details. This resulted in a refund of the purchase price by the experimenters. The authors do not report the proportion of participants who placed orders with no intention to actually receive and use the purchased product [49].

The "vibrator vs. batteries" study improved upon an earlier study by Gideon et al. [53] from the same research group. Both studies used the same privacy-enhanced product search engine, and participants could choose amongst competing sellers. In the earlier study, 24 participants were recruited and paid a show-up fee of $10. Again, products varied by privacy sensitivity, with surge protectors and condoms as the extremes. Although prices differed amongst sellers, participants did not need to pay a premium for privacy, because all expenses were reimbursed by the experimenters. The main conclusion from this study would, therefore, be that consumers prefer privacy-friendly designs so long as they come for free.

14.3.3.2 Experiment "Gourmet Food"

In an earlier, unpublished experiment by Preibusch [55]—some material of which is depicted in Fig. 14.2—72 participants were invited to shop for gourmet food within

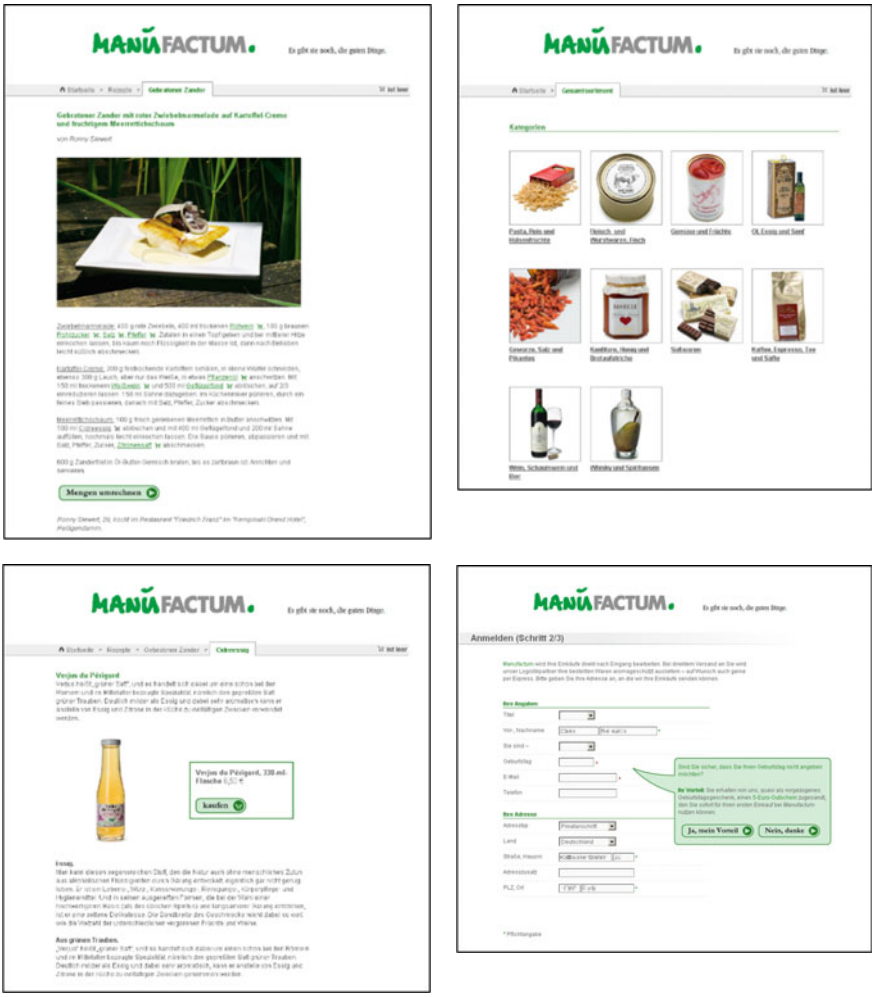

Fig. 14.2 Screenshots from a 2007 laboratory experiment, featuring an online store for gourmet food that could be browsed by product categories or through recipes by famed chefs

a single online store. Some 330 products were available along with 75 recipes by famous chefs to ease product selection. Participants were paid a €10 show-up fee and distributed evenly between two treatments. In one treatment, shoppers could twice receive an extra €5 for indicating their date of birth and their email address when making a purchase. Of all participants, 39 % placed an order, 15 of them in the incentivized treatment, 13 in the non-incentivized treatment. Through the incentives, the data disclosure ratio for date of birth could be increased from 75 % to 92 %; and from 81 % to 92 % for email. In the non-incentivized treatment, only 64 % provided their date of birth.

14.3.3.3 Experiment "DVDs"

In the 2009 DVD study, 225 participants had the choice between two DVD retailers that offered the same range of films. Thirty bestsellers were preselected and presented in a color-printed folder, but buyers had access to the entire Amazon product range through a real-time search API, offering around 100 thousand titles. In fact, almost half of the buyers (47 %, 35 in 74) made their purchase after having requested titles not in the original catalogue. We partnered with an existing bricks-and-mortar retailer of new and used CDs and DVDs [50].

Buyers had the choice between two competing branches, Cologne and Frankfurt. The order forms listed the movie titles with their prices side-by-side to the personal details required for the checkout, so that neither the prices nor the privacy aspects where given priority. Frankfurt was the privacy-invasive retailer, always asking for income and phone number when Cologne only required favorite color. In one treatment, prices were the same; in the other treatment, Frankfurt was €1 cheaper. When prices were the same, buyers seemed to pick an online store at random. They did not systematically prefer the privacy-friendly branch. When prices differed, very few were willing to pay an extra euro for not revealing their mobile phone number and income. However, they were retrospectively less satisfied with the seller's privacy practices—as found in the exit-questionnaire [52]. We also saw that the discount was overriding participants' privacy preferences: for Frankfurt buyers, there was a significant negative association between their willingness to provide the data items required by the privacy-invasive retailer, and their actual data disclosing behavior.

14.3.3.4 Experiment "Cinema Tickets"

The most thorough experiment into privacy economics to date is the 2012 cinema ticket study. It builds on the earlier DVD study, described in Sect. 14.3.3.3. The cinema ticket experiment took into consideration the lessons learnt from the DVD study. In the face of an overhaul of the EU legislation on data protection, the study on "monetizing privacy" was commissioned and funded by the European Network and Information Security Agency (ENISA) and done in collaboration with researchers at the German Institute for Economic Research. The over-arching research questions were:

- Do some customers of online services pay for privacy?
- Do some individuals value their privacy enough to pay a mark-up to an online service provider who protects their information better? [51]

To answer these questions, we created an online shopping experience where consumers faced a trade-off between privacy and price. Ultimately more than 500 laboratory participants were invited to buy up to two cinema tickets. Purchase ratios were high (43 %) and most of the buyers bought two tickets. The report published by ENISA [51] gives the results for the first 443 participants, who purchased a total

of 344 tickets—here I am reporting the results for loyal buyers, who purchased two tickets from the same firm. Upon checkout, buyers had a choice between two different retailers, shown side-by-side (Fig. 14.3). One of them asked for their mobile phone number in addition to the basic data of name, email, and date of birth. More than 80 % of buyers chose the privacy-friendly seller when prices were the same. The privacy-friendly retailer continued to attract a demand when its prices were higher. Around a third of the loyal buyers paid €1 extra for keeping their phone number private. These results are statistically highly significant. We also fielded this experiment nationwide and the results were corroborated, providing strong evidence that the results from the laboratory do generalize.

The design of the cinema ticket study closely followed the earlier DVD experiment, with some improvements. Again, a laboratory experiment was

Fig. 14.3 Public-facing web site deployed for the cinema ticket study, featuring the price and privacy points of two alternative sellers side-by-side

implemented, but complemented with a hybrid and a field deployment. In the field experiment, all interactions took place on a public web site. Participants ignored the fact that they were partaking in an experiment. In the hybrid, participants interacted with the same public web site, but were explicitly invited to participate, using university mailing lists. Consequently, the hybrid and the laboratory experiments build on a student-dominated participant pool, whereas the field experiment samples from the general online population.

The experiment was framed as a study into how consumers make purchase decisions. Online sales of cinema tickets was taken as an example; cinema going is a broad social phenomenon [56], and ticket purchases and the consumption of culture are widespread activities on the Web [57: Table 10]. The advantages of DVDs, including low price and homogeneity, also apply to cinema tickets. The main difference in the experimental setup is a more pronounced privacy gradient. For the DVD study, data collection did not differ between the retailers in the number of data items required, and participants needed to inspect the two order forms closely to spot the difference. In contrast, the cinema ticket sellers differed in the number of data items collected and, with four versus three items, the variation is relatively high. Furthermore, the side-by-side display on-screen made comparisons easy.

14.4 How to Run Experiments in Behavioral Privacy Economics

14.4.1 Measuring Willingness to Pay

Experiments into the behavioral economics of privacy aim to measure the value that consumers attribute to their privacy or to privacy-enhancing features. Examples of privacy-enhancing features can be found easily in digital goods and services: a web browser with enabled tracking protection, a webmail provider that refrains from scanning messages, or a search engine that offers its users the ability to disable or to curate their search history. There is a research and business interest in measuring how much these privacy enhancements appeal to users, in absolute monetary terms (e.g., for pricing subscriptions) and relative to other features such as search result quality (e.g., for prioritizing engineering efforts).

Examples of enhanced privacy are often found in the way companies provide goods and services to their customers. An online retailer may refrain from asking sensitive personal information, or may not use the order confirmation email address provided by the customer to send them unsolicited newsletters. Better privacy is thus operationalized along one of the privacy dimensions of data collection, use, retention, and sharing. Research and business are interested in two ways in which consumers articulate their value of privacy. First, would they pay money or give up other desirable things such as personalization to enjoy more privacy? Second,

framed inversely, would they give up privacy to receive discounts, higher payoffs, or to enjoy more functionality and convenience? Whereas the first question examines a willingness to pay for privacy, the second question looks at the willingness to accept incentives towards increased data disclosure. It has been speculated that the willingness to pay and to accept differ; however, available evidence is inconclusive [40, 58].

An economic experiment into the value of privacy places consumers into a decision-making situation where they have to trade off privacy against money/convenience/functionality. Their choices are observed in a laboratory or field study. Willingness to pay for privacy is revealed through controlled variation of the stimulus across treatments, such as the discount an online shop grants the customer for revealing his or her mobile phone number.

As a discipline, experimental economics have developed principles on how to conduct such experiments. In some aspects, these differ from other disciplines that are also looking at decision making, including psychology [59]. First, experiments in economics are scripted: participants' progress through the experiment, their possible choices and payoffs are set forth in a detailed protocol. Second, the payoffs are variable and depend on the choices participants have made and their performance. Third, deception is avoided throughout the experiment [59].

The methodological differences between economists and psychology researchers can be a practical challenge. The design of an experiment may face opposition when reviewed by an ethics committee that subscribes to the respectively other research standard. Especially performance-dependent payoffs may face resistance amongst psychology scholars.

14.4.2 Essential Stages of the Experiment

14.4.2.1 Sign-up and Participation

An experiment that measures consumers' value of privacy usually progresses through several stages, some of which happen before and some after the session (Fig. 14.4). Potential participants sign up beforehand, at which a screening questionnaire may be deployed to sample a specific population. At the time of the session, identity checks are carried out; only registered participants are admitted, without walk-in participation. Although a single session is attended by multiple participants, they progress at their own pace and must not communicate with each other, unless the procedures explicitly foresee teaming.

14.4.2.2 Instructions and Consent

The session starts by explaining the procedures to the participants, the choices they will have during the session and how their choices will impact their payoffs.

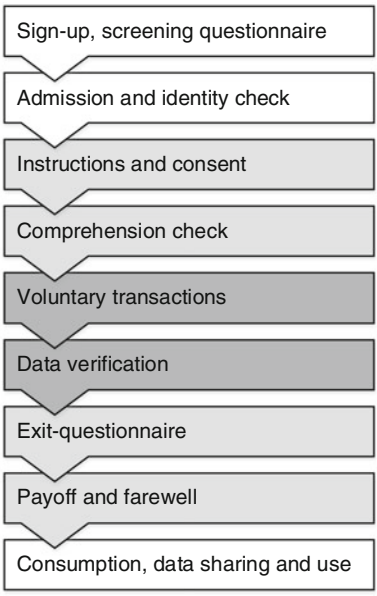

Fig. 14.4 Stages of an experiment session into the behavioral economics of privacy

The instructions, succinct yet complete, are distributed as hard-copies to the participants, so that they can refer to them easily later on. Experimenters may also read them out loud; and arising questions are answered. As an ethical obligation, participants are also informed about potential harms and about what will happen with their data. Participants may still withdraw at this stage without being sanctioned, as participation is voluntary and requires consent. It is important that participants understand the procedures, because otherwise the nexus between the stimulus and their actions is broken.

14.4.2.3 Voluntary Transactions and Data Verification

Voluntary transactions are at the core of the session. They bring a real-world scenario into the laboratory and at best replicate every aspect of the consumption scenario. The creation of a realistic environment is resource-intensive; further details are provided below. The interaction is typically fully computerized although ancillary non-digital materials may be used. Computer-mediated delivery allows high levels of instrumentation so that participants' actions are logged precisely to be analyzed later on. It also makes data verification possible, which is crucial in privacy experiments: if participants provide fake data, they can avoid privacy risks and thereby contravene the experimental protocol. Contact details including mailing addresses, email addresses, and mobile phone numbers can be verified through delivery checks, when a sent confirmation code has to be rekeyed into a web site.

Personal details including name, date of birth, and nationality can be verified using ground truth such as identity cards, leveraging the face-to-face interaction the laboratory offers. Biometrics can be checked in place, through observation (e.g., gender) or by measurement (e.g., body height and weight).

14.4.2.4 Payment Collection

If the transaction involves payment by the participant, such as for an online purchase, payment collection should happen as part of the transaction. This includes electronic payment, where PayPal can be offered as long as the extra third party does not subvert the experiment design. Preference should otherwise be given to a white-label credit card acquirer—however, this leads to substantial overhead. Offering card payments contributes to a realistic shopping environment and makes sure that those without cash at hand can engage. Cash payments may be settled as part of the final payoff.

14.4.3 Creating Real-World Shopping Scenarios

14.4.3.1 Purchase Ratios and Product Selection

True choices require voluntary transactions. This can be challenging when studying privacy economics in electronic retailing as only buyers contribute observations into price–privacy trade-offs. Researchers must, therefore, achieve high purchase ratios to make the most of the recruited sample. At the same time, participants' decisions to make a purchase should not be systematically associated with their privacy attitudes or demographics, which can be checked through the exit-questionnaire. Purchase ratios between 40 % and 60 % can be achieved. Product selection is key and there are five guiding principles. First, the product should appeal universally, regardless of age, gender, or education. Second, the product should be affordable, especially for cash-strapped student samples. Third, the product should lend itself to impulse purchases, without requiring much thought or outside information seeking. Fourth, the product must be homogenous: its quality should be unaffected by whoever sells it. Finally, regulation rules out certain products: in most countries, age restrictions apply to alcohol, tobacco, or pornography; licenses may also be required. Train tickets, although an otherwise suitable product, may only be sold by authorized persons in the United Kingdom. A fixed subsidy can be offered to stimulate purchases; but Germany, for instance, forbids discounts on books. The research design guides whether multiple products or quantities can be bought or whether unit demand is enforced.

14.4.3.2 Real and Fulfilled Transactions

The laboratory is the gold standard for human subject experiments: it allows rich data collection in a controlled environment that rules out confounds beyond the stimuli of the experiment. All the same, results established in the laboratory are sometimes criticized for low generalizability beyond this protected realm. In economics, behaviors observed in the laboratory have been found repeatedly to be a good predictor for the outside world, even when student volunteers are recruited [60]. Having a realistic decision-making environment is key in delivering valid results that generalize beyond the laboratory.

It is, however, not sufficient for transactions in a laboratory experiment to be realistic; they need to be real. This mandate follows from the proscription of deception. Deception pollutes the shared resource of a participant pool by creating distrust [61]. This ethical argument against deception is complemented by an argument of scientific validity. Deception has been found to have an impact on participants' behavior [62]: their actions become inconsistent, a sign that they stop taking the experiment seriously. It also has a negative effect on return rates and may yield self-selection biases. Deception, therefore, is a serious threat to the validity of the findings.

In an experiment into privacy behaviors, all aspects of the transaction must be real. Taking the example of online shopping, participants who make a purchase will have to pay for it with their own money, and they need to be provided with the good they bought. Researchers must be able to fulfill orders, although this will typically be easier for products that also improve purchase ratios, as discussed above. The exchange of personal information must be real and any uses or third-party sharing that were communicated to the participant have to be executed. For instance, if allergies are collected to filter products in an online grocer, this functionality must be implemented. If participants were told that their data would be shared with a company, this data transfer needs to happen, because a deflected privacy threat also counts as deception. It is intended that participants' choices during the experiment have an impact on their lives after the session [62].

14.4.3.3 Partnering with External Companies

One of the challenges in creating a real transaction is to bring a commercial service provider or retailer into the laboratory. When studying exchanges of money and personal data between companies and their data, one cannot achieve valid results when making participants interact with a trusted institution like a university. In the instructions, participants must be told truthfully that they transact with an existing company. Researchers, therefore, have to collaborate with established firms: the setup of a working relationship with a retailer requires negotiation talent and an understanding of business requirements. For a company, it is typically a cost to support a research study. Practical hurdles are the use of branding and the licensing

of intellectual property such as designs and product images or descriptions. When granted permission, these elements will be used to re-create a real web shop in the laboratory (Fig. 14.2).

14.4.4 Maintaining the Institutional Separation Between University and Corporate Representation in the Laboratory

Participants must be supported in understanding that an experiment features two separate institutions, a university by members of which and on the premises of which the session is administered, and a company that is their transaction partner. This separation is crucial as participants will react differently to their information requests. The university benefits from a positive trust bias when collecting and using personal information that should not spill over to the company, or the validity of the results would suffer. The two institutions (university versus company) can be separated by experiment phases (Fig. 14.4). Whilst the company collects personal details during a voluntary transaction, university researchers ask participants to complete an exit-questionnaire. Truthful responses in the latter hinge on confidentiality towards the company. For instance, buyers may be asked how they rate the shopping experience. Without institutional separation, this question would be subject to a social desirability bias.

The logical separation between the university and the company corresponds to an administrative separation of the inner and outer phases of the experiment session (Fig. 14.4). A visual break can help participants: parts of the experiment related to the university may feature a different visual language, different fonts, or colors, or it might address the participant differently (e.g., John vs. Mr. Doe, "participant" vs. "dear customer").

An exit-questionnaire records socio-demographic key indicators, such as age, gender, income, and education level. All participants take this questionnaire, whether or not they decided to transact. It allows controlling for potential decision drivers such as computer literacy and past experiences with the chosen company or with cyber-crime. Personality traits such as materialism, reciprocity, risk-aversion, and indeed privacy concerns are measured using instruments with pre-established reliability [31].

14.4.5 Deploying the Experiment: The Relative Merits of the Laboratory, the Field, and Online Platforms

Researchers have the choice between three channels for deploying their experiments: the laboratory, the field, and online platforms (e.g., Amazon Mechanical

Turk—mTurk). Each of these platforms has their own advantages and disadvantages. Dual deployment or hybrids between the channels promise more robust findings.

Laboratory sessions are the traditional way of running experiments in behavioral economics. Their main advantage is the full control that researchers have over the experimental design and the deployment. The lab creates an isolated realm, which allows controlled manipulation of the stimulus under investigation. Possible confounds can be minimized or ruled out entirely. The laboratory also features synchronized face-to-face administration of the experiment, allowing the pairing of participants without them having to wait for one another. Rich apparatus such as eye-tracking or biometric sensors can be used. In privacy experiments, tracking the gaze of participants allows the experimenter to check whether the subjects have actually read a web site's privacy policy, seen any available discounts, or all the potentially sensitive data items requested on the checkout form. Experimenters are also able to verify personal information on site, for instance, through direct observation or with an identity check; verifying someone's name over the Internet is often prohibitively difficult. The drawback of lab experimentation is their over-reliance on student participants and on educated subjects from rich Western societies [63], which may come at the expense of generalizability. Furthermore, it is difficult to scale laboratory experiments beyond a few hundred participants, as subject pools deplete.

Field experiments give access to a potentially unlimited population, although one typically restricts recruiting to a single country for practical reasons, such as language localization or compliance with national regulation. In a field study, researchers create a public facing web site or team up with an existing company to bring the interaction from the laboratory into the wild. Often, it is no longer obvious that the web site is part of a research study. The main advantage is that the experiment is no longer pushed onto participants; instead customers come to the web site self-motivated and task-driven. Pull engagement has the advantage of capturing the consumers when and how they want to transact online. This brings new challenges for recruitment: laboratory subject pools may be invited to join the field study, although it bears the risk of contaminating the natural interaction, which has otherwise no connection to a university or research institution. To recruit for field studies, advertising campaigns may be necessary, resulting in recruiting costs that might be higher than for the laboratory. Logging referrers to the fielded web site is essential. The main cost driver for field experiments remains the requirement to create an instrumentation that survives in the wild. There is a higher bar for design and visual appeal, for security, and any fielded materials face regulatory exposure, as the web site created for research purposes now enters the competitive market. A mentality shift is required for researchers, from administering an experiment session to delivering customer support. Exit-questionnaires and follow-ups are more difficult to administer in a field study; the diversity of the sampled population may go unnoticed. Researchers should be prepared that the field gives noisier data than the laboratory.

Crowdsourcing platforms have started a new wave of studies with human participants. They allow researchers to collect data more quickly and cheaply than through laboratory studies. For many computer scientists, the first experiment with human participants they ever run will be online. Amazon Mechanical Turk (mTurk) is the best-known platform, although numerous crowdsourcing platforms are now available. Started as a labor market for large numbers of small, tedious tasks such as transcribing business cards, mTurk has been seized by researchers who need to conduct experiments and deploy surveys. Comprehensive guidance is available to researchers to on-board with the platform [64]. Cost savings and timely turnaround are the two main advantages of online experiments. Payments to participants can be lower by one order of magnitude and experiments can be run around the clock with minimal supervision. The major difficulty is the introduction of a new sampling bias towards a population that goes after pennies and is recruited in a task-focused mind-set. Cheaters and spammers are common on mTurk; many of them have previously participated in psychology experiments involving deception, so their behavior may be distorted [61]. Finally, platform operators such as Amazon impose strict guidelines on what is allowed on the platform. The main hindrance for privacy researchers is the proscription to collect personal information [65], and the resulting inability to create a real invasion of privacy. Whilst crowdsourcing lends itself to many experiment procedures, researchers should refrain from retrofitting their research question or experiment design. Despite their pragmatic appeal, platforms such as mTurk are often unable to accommodate for the requirements of research into privacy economics.

14.5 Conclusions and Future Challenges

14.5.1 Principles for Empirically Studying Privacy Behaviors

Privacy is top of mind for corporate executives, regulators, and policymakers. Since the Web has brought mass-personalization to every aspect of online consumption, privacy advocates have argued how ubiquitous web tracking poses a threat to users' informational self-determination. Today, we know the reality of planet-scale government surveillance, and Big Data companies demonstrate how personal information can be monetized. This revived interest in improving the protection of consumers' personal information suffers from a serious knowledge gap into consumers' privacy concerns and behaviors. When public opinion polls repeatedly diagnose high levels of privacy concern, it seems paradoxical that consumers keep enjoying privacy-invasive services. There is surprisingly little knowledge on how consumers make privacy/price/convenience trade-offs and about the value they attach to their personal information. Reliable and valid evidence is needed to develop privacy-enhancing technologies that meet the consumers' needs.

Behavioral economics provide the methodological toolkit to explore consumers' privacy choices.

Well-crafted experiments in the laboratory or in the field put participants in real-world decision-making scenarios that allow observation of their privacy choices with predictive power. Applied to the study of privacy in electronic retailing, for instance, this means offering participants the ability to voluntarily purchase goods or services; the transaction is fulfilled with exchanges of real products, money, and personal data. Conversely, the lack of commitment and incentive-compatibility makes surveys, hypothetical choice scenarios, or studies involving deception fail to deliver actionable insight.

14.5.2 Future Challenges

In this chapter, I have outlined the principles of conducting empirical research into consumers' privacy consumption behaviors. For researchers, practitioners, and policymakers more challenges lie ahead.

Challenges for researchers include the development of new measurement instruments for privacy concerns and behaviors. On the one hand, we witness the emergence of new kinds of personal information, collected through the proliferation of sensors in mobile devices and public spaces: real-time location data, biometrics collected from eye-tracking, video surveillance, and health sensors. Big Data is not just more of the same, but introduces challenges of a new type [66, 67]. On the other hand, well-conducted experiments are time and resource consuming to a point where knowledge production has difficulties keeping up. This calls for an experiment infrastructure to conduct empirical studies at a faster pace and with lower investments, and it also calls for reliable and valid low-cost survey instruments.

Challenges for companies lie in the diversity of consumers' privacy preferences. How can a company implement superior privacy practices, when customers are diverse in how they balance the trade-off between convenience and data minimization? How can business models succeed beyond the monetization of personal data when the majority of buyers choose cheap prices over good privacy? Privacy negotiations allow companies to offer their customers the choice between different privacy regimes where the current one-size-fits-all approach of inflexible privacy statements fails [68].

Challenges for regulators include the unification of consumer protection and data protection. Two different enforcement regimes need to be combined as market power is redefined in the digital economy. Barriers to entry are no longer capital investments but access to large quantities of historical data; the demand-side network effects in data-intensive products can quickly turn a successful firm into a dominant firm. A mark-up on prices is the traditional symptom of monopolies, but how does market concentration manifest when products and services are offered free of charge. Ultimately, regulators aim to create an environment where privacy-friendly products and companies will thrive.

For companies making sense of big personal data without alienating their customers, for regulators upholding privacy norms, and for researchers envisioning new data protection technologies, it is key to understand consumers' privacy concerns and behaviors. In this chapter, I have shown how laboratory experiments and field studies can observe consumers making real-world privacy choices and thereby provide decision makers with the reliable and valid empirical evidence they need.

Acknowldgements Kat Krol (University College London) provided helpful comments on earlier versions of the chapter.

References

1. The Boston Consulting Group (2012) The digital manifesto. How companies and countries can win in the digital economy
2. The Boston Consulting Group (2010) The connected kingdom. How the internet is transforming the U.K. Economy
3. Office for National Statistics (2014) Retail Sales, May 2014
4. Nadella S (2014) A data culture for everyone. http://blogs.technet.com/b/microsoft_blog/archive/2014/04/15/a-data-culture-for-everyone.aspx
5. Kuneva M (2009) Keynote Speech: roundtable on online data collection, targeting and profiling. In: European Commission
6. Personalization consortium, personalization & privacy survey, 2000, 2005, via Internet Archive, 2014
7. Kobsa A (2007) Privacy-enhanced web personalization. In: Brusilovsky P, Kobsa A, Nejdl W (eds) Privacy-enhanced web personalization, vol 4321. Springer, Berlin, pp 628–670
8. Hess T, Schreiner M (2012) Ökonomie der Privatsphäre. Datenschutz und Datensicherheit—DuD 36:105–109
9. Consumer Commerce Barometer, February 2012. http://www.consumerbarometer.eu/chartlink?id=T1332485008
10. Hamilton Consultants, Inc. (2009) Economic value of the advertising-supported internet ecosystem. IAB
11. World Economic Forum (2011) Personal data: the emergence of a new asset class
12. European Data Protection Supervisor (2014) Privacy and competitiveness in the age of big data. European Data Protection Supervisor
13. European Commission/TNS Opinion & Social (2011) Attitudes on data protection and electronic identity in the European Union
14. Preibusch S (2015) Privacy behaviours after Snowden: the brief impact of exposed state surveillance. Commun ACM 58(5):48–55
15. Anderson R, Moore T (2006) The economics of information security. Science 314(5799):610–613
16. Heller C (2011) Post-Privacy. C.H. Beck, Prima leben ohne Privatsphäre
17. Feijóo C, Gómez-Barroso JL, Voigt P (2014) Exploring the economic value of personal information from firms' financial statements. Int J Inf Manage 34(2):248–256
18. Novotny A, Spiekermann S (2013) Personal information markets and privacy: a new model to solve the controversy. In: 11 Internationale Tagung Wirtschaftsinformatik, Leipzig
19. Department for Business, Innovation & Skills, The midata vision of consumer empowerment
20. The Gallup Organization (2008) Data protection in the European Union. Citizens' perceptions (analytical report). In: European Commission

21. Acquisti A, Gross R (2006) Imagined communities: awareness, information sharing, and privacy on the Facebook. In: Privacy enhancing technologies
22. Barnes SB (2006) A privacy paradox: social networking in the United States. First Monday 11(9)
23. Spiekermann S, Grossklags J, Berendt B (2001) E-privacy in 2nd generation e-commerce: privacy preferences versus actual behavior. In: Proceedings of the 3rd ACM conference on electronic commerce, EC '01, New York, NY, USA
24. Norberg P, Horne D, Horne D (2007) The privacy paradox: personal information disclosure intentions versus behaviors. J Consum Affairs 41(1):100–126
25. Grossklags J, Acquisti A (2007) When 25 cents is too much: an experiment on willingness-to-sell and willingness-to-protect personal information. In: Workshop on the economics of information security (WEIS)
26. Preibusch S (2013) The value of privacy in Web search. In: Workshop on the economics of information security (WEIS)
27. Rivenbark DR (2010) Experimentally elicited beliefs explain privacy behavior. University of Central Florida—College of Business Administration
28. Microsoft, Trustworthy computing—data privacy day. http://www.microsoft.com/en-us/twc/privacy/data-privacy-day.aspx
29. Connelly K, Khalil A, Liu Y (2007) Do I do what I say?: Observed versus stated privacy preferences. In: Proceedings of the 11th IFIP TC 13 international conference on human-computer interaction, INTERACT'07
30. Preibusch S, Krol K, Beresford AR (2012) The privacy economics of voluntary over-disclosure in Web forms. In: Workshop on the economics of information security (WEIS)
31. Preibusch S (2013) Guide to measuring privacy concern: Review of survey and observational instruments. Int J Hum Comput Stud 71(12):1133–1143
32. Wathieu L, Friedman A (2005) An empirical approach to understanding privacy valuation. In: Workshop on the economics of information security (WEIS)
33. Ward S, Bridges K, Chitty B (2005) Do incentives matter? an examination of online privacy concerns and willingness to provide personal and financial information. J Mark Commun 11 (1):21–40
34. Castañeda JA, Montoro F (2007) The effect of Internet general privacy concern on customer behavior. Electr Comm Res 7:117–141
35. Baumer DL, Poindexter JC, Earp JB (2006) An experimental economics approach toward quantifying online privacy choices. Inf Syst Front 8:363–374
36. Hann I-H, Hui K-L, Lee S-YT, Png IP (2007) Analyzing online information privacy concerns: An information processing theory approach. In: 40th Annual Hawaii international conference on system sciences (HICSS 2007)
37. Hann I-H, Hui K-L, Lee S-YT, Png IP (2002) Online information privacy: measuring the cost-benefit trade-off. In: 23rd international conference on information systems (ICIS)
38. Jensen C, Potts C, Jensen C (2005) Privacy practices of internet users: self-reports versus observed behavior. Int J Hum Comput Stud 63(1–2):203–227
39. Huberman B, Adar E, Fine L (2005) Valuating privacy. IEEE Secur Priv 3(5):22–25
40. Acquisti A, John LK, Loewenstein G (2013) What is privacy worth? J Legal Stud 42(2):249–274
41. Cvrcek D, Kumpost M, Matyas V, Danezis G (2006) A study on the value of location privacy. In: Proceedings of the 5th ACM workshop on privacy in electronic society, WPES '06, New York, NY, USA
42. Kai-Lung H, Hai TH, Tom LS-Y (2007) The value of privacy assurance: an exploratory field experiment. MIS Q 31(1):19–33
43. Carrascal J, Riederer C, Erramilli V, Cherubini M, de Oliveira R (2011) Your browsing behavior for a big mac: economics of personal information online. In: Proceedings of the 22nd international conference on World Wide Web
44. K-L Hui, Lee SYT, Teo HH (2007) The value of privacy assurance: an exploratory field experiment. MIS Q 31(1):19–33

45. White TB (2004) Consumer disclosure and disclosure avoidance: a motivational framework. J Consum Psychol 14(1/2):41–51
46. Kobsa A, Teltzrow M (2005) Impacts of contextualized communication of privacy practices and personalization benefits on purchase behavior and perceived quality of recommendation
47. Ellison NB, Steinfield C, Lampe C (2007) The benefits of Facebook "friends:" social capital and college students' use of online social network sites. J Comput-Mediat Commun 12 (4):1143–1168
48. Taddicken M, Jers C (2011) "The uses of privacy online: trading a loss of privacy for social web gratifications? In: Privacy Online. Springer, Berlin, pp 143–156
49. Tsai J, Egelman S, Cranor L, Acquisti A (2009) The impact of privacy indicators on search engine browsing patterns. In: Symposium on usable privacy and security (SOUPS), New York, NY, USA
50. Preibusch S, Kübler D, Beresford AR (2013) Price versus privacy: an experiment into the competitive advantage of collecting less personal information. Electron Comm Res 13(4):423–455
51. Jentzsch N, Preibusch S, Harasser A (2012) Study on monetising privacy. An economic model for pricing personal information. ENISA
52. Beresford A, Kübler D, Preibusch S (2012) Unwillingness to pay for privacy. Econ Lett 117 (1):25–27
53. Gideon J, Cranor L, Egelman S, Acquisti A (2006) Power strips, prophylactics, and privacy, oh my!. In: Symposium on usable privacy and security (SOUPS), New York, NY
54. Jentzsch N, Giannetti C (2011) Disclosure of personal information under risk of privacy shocks. SSRN
55. Preibusch S (2008) Economic aspects of privacy negotiations., Berlin, Germany: Technische Universität Berlin/Fachgebiet Volkswirtschaftslehre [Technical University Berlin/Institute of Economics]
56. Eurostat (2006/2011) Percentage of persons who have attended the cinema at least once in the last 12 months by gender and age group
57. van Eimeren B, Frees B (2012) 76 Prozent der Deutschen online—neue Nutzungssituationen durch mobile Endgeräte. MEDIA PERSPEKTIVEN, pp 362–379
58. Plott CR, Zeiler K (2005) The willingness to pay-willingness to accept gap, the "Endowment Effect", subject misconceptions, and experimental procedures for eliciting valuations. Am Econ Rev 95(3):530–545
59. Hertwig R, Ortmann A (2001) Experimental practices in economics: a methodological challenge for psychologists? Behav Brain Sci 24(3):383–403
60. Exadaktylos F, Espín AM, Brañas-Garza P (2013) Experimental subjects are not different. Sci Rep 3(1213)
61. Horton JJ, Rand DG, Zeckhauser RJ (2011) The online laboratory: conducting experiments in a real labor market. Exp Econ 14(3):399–425
62. Jamison J, Karlan D, Schechter L (2008) To deceive or not to deceive: the effect of deception on behavior in future laboratory experiments. J Econ Behav Organ 68(3–4):477–488
63. Henrich J, Heine SJ, Norenzayan A (2010) The weirdest people in the world? Behav Brain Sci 33(2–3):61–83
64. Mason W, Suri S (2012) Conducting behavioral research on Amazon's mechanical Turk. Behav Res Methods 44(1):1–23
65. Amazon.com, Inc. (2013) FAQs | Help | Requester | Amazon Mechanical Turk. https://requester.mturk.com/help/faq#restrictions_use_mturk
66. Executive Office of the President (2014) Big data: seizing opportunities, preserving values
67. Crawford K, Schultz J (2014) Big data and due process: toward a framework to redress predictive privacy harms. Boston Coll Law Rev 55(93)
68. Preibusch S (2006) Implementing privacy negotiations in E-commerce. In: Frontiers of WWW research and development—APWeb 2006, Berlin Heidelberg