THE GEORGE WASHINGTON UNIVERSITY

WASHINGTON, DC

Working Paper[1]

November 6, 2017

Measuring Costs and Benefits of Privacy Controls

Conceptual Issues and Empirical Estimates

Joseph J. Cordes, Co-Director[2]

Daniel R. Pérez, Policy Analyst[3]

The George Washington University Regulatory Studies Center

Introduction

As more and more items of personal information become potentially available to internet providers, the government, and employers, a lively debate has emerged about the role of public policy in ensuring a proper balance between the various parties who may benefit from greater access to information, and the protection of individual rights to privacy. A recent example is legislation passed in the Congress repealing a regulation that "would have required Internet service providers—like Comcast, Verizon and Charter—to get consumers' permission before selling their data."[4] As Hahn and Layne-Farrar[5] and Adam Thierer[6] have noted, it is desirable

[1] This working paper reflects the views of the author, and does not represent an official position of the GW Regulatory Studies Center or the George Washington University. The Center's policy on research integrity is available at http://regulatorystudies.columbian.gwu.edu/policy-research-integrity.

[2] Professor of Economics, Public Policy and Public Administration, and International Affairs, Trachtenberg School of Public Policy and Public Administration, the George Washington University. Co-Director of the George Washington University Regulatory Studies Center. PhD University of Wisconsin-Madison Economics (1977).

[3] PhD Student, Public Policy and Public Administration, Trachtenberg School of Public Policy, the George Washington University. Policy Analyst at the George Washington University Regulatory Studies Center.

[4] Brian Naylor, *Congress Overturns Internet Privacy Regulation*, NPR (2017), *available at* http://www.npr.org/2017/03/28/521831393/congress-overturns-internet-privacy-regulation.

that this debate be informed by a formal benefit-cost analysis based on empirical measures of benefits and costs.

Additionally, emerging technologies such as highly automated vehicles (HAVs or "driverless cars") and unmanned aircraft systems (UAS or "drones") bring privacy concerns to the forefront—particularly regarding the proper role of federal regulatory agencies. Accordingly, agencies such as the National Highway Traffic and Safety Administration (NHTSA) and the Federal Aviation Administration (FAA) currently face the difficult task of balancing their objectives of issuing sensible regulations that offer protections to consumers, with allowing continued innovation and use of these emerging technologies.

It is worth noting that the regulatory process "incorporates significant requirements regarding the collection, use and accessibility of data that differ from other policymaking processes."[7] Statutes such as the Administrative Procedure Act of 1946[8] (APA) require agencies to "justify most regulatory decisions based on the data, analyses, and other information collected and made part of a publically available record."[9] Data and other evidence used by agencies to justify rulemaking become part of the public record and are particularly relevant in the case of judicial review—where regulations can be vacated if courts determine agency actions to be "arbitrary and capricious."[10] The APA is only one of numerous mandates that constrain and guide the rulemaking process.[11]

Usable estimates of consumer privacy are of particular benefit to federal regulatory agencies considering existing analytical requirements concerning the collection of information such as the Paperwork Reduction Act (PRA).[12] The PRA requires agencies "to justify any collection of

[5] *See* Robert W. Hahn & Anne Layne-Farrar, *The Benefits and Costs of Online Privacy Legislation*, 54 ADMIN. L. REV. 85, (2002).

[6] *See* Adam D. Thierer, *A Framework for Benefit-Cost Analysis in Digital Privacy Debates*, 20 Geo. Mason L. Rev. 1055, 1056-57 (2013).

[7] Marcus C. Peacock, Sofie E. Miller & Daniel R. Pérez, *Public Interest Comment to the Commission on Evidence-Based Policymaking* 2 (2016), *available at* https://regulatorystudies.columbian.gwu.edu/public-comment-commission-evidence-based-policymaking. (detailing a framework for producing evidence-based regulation structured around the three main phases of regulating: design, decision-making, and retrospective review).

[8] PUB. L. NO. 79-404, 60 Stat. 237.

[9] Peacock, Miller & Pérez, *supra* note 7, at 2.

[10] 5 U.S.C. § 706(2)(A).

[11] *See, e.g.,* The Privacy Act of 1974, 5 U.S.C. § 552(a). *See generally* Susan E. Dudley & Jerry Brito, THE MERCATUS CTR. AND THE GEO. WASH. UNI. REG. STUDIES CTR., *Regulation: A Primer,* 45-7 (2d ed. 2012) (for a thorough list of laws and executive orders affecting regulatory policymaking). *See also* Susan E. Dudley, *Putting a Cap on Regulation*, 42 REG. LAW NEWS, AMERICAN BAR ASSOCIATION 4-6 (2017) (for a detailed explanation of executive orders affecting the rulemaking process signed by President Trump which include: EXEC ORDER NO. 13,771, 82 FED. REG. 9339 (February 3, 2017) and EXEC ORDER NO. 13777, 82 FED. REG. 12285 (March 1, 2017).

[12] 44 U.S.C. §§3501-3520

information from the public by establishing the need and intended use of…information…and showing that the collection is the least burdensome way to gather the information."[13]Agencies must receive approval from the Office of Information and Regulatory Affairs (OIRA) before initiating any information collection from ten or more people.[14]

In short, these mandates require agencies to base their rulemaking on a thorough analysis of regulatory benefits and costs—with added requirements to conduct retrospective (*ex post*) review of regulatory impacts. As the U.S. economy continues to be exponentially reliant on data generated via the collection of individuals' personally identifiable information (PII), regulatory agencies will need empirical measures of consumer valuations of privacy.

Our paper hopes to contribute to the development and greater use of such empirical measures. Drawing on the economics of privacy literature, we summarize the insights that this literature has to offer about how the benefits and costs of privacy controls should be measured *in principle*. We then discuss attempts that have made been to measure the benefits and costs of privacy control. Finally, we synthesize the various findings to advance promising practices for generating useful estimates of U.S. consumers' valuation of privacy.

The Basic Framework for Benefit-cost Analysis

As noted in a widely used textbook on benefit-cost analysis[15] the two foundational measures of benefit and cost in benefit cost analysis are (1) *willingness to pay* (WTP) as measures of benefit to individuals who gain from a policy or as costs to individuals who are harmed; and (2) the *social opportunity costs* of inputs used to implement the policy.

Willingness to pay, in turn can be measured in principle by the compensating variation (or in some cases equivalent variation) of a policy change, where the compensating variation equals the maximum amount of income a beneficiary of a policy would be willing to give up in order to have the policy implemented. Conversely, the compensating variation of someone harmed by the policy would equal the minimum amount of income that would need to be paid to someone harmed by the policy to leave them no worse than before the policy change. An alternative measure of willingness to pay, *equivalent variation*, equals the minimum amount of income that would need to be paid to a beneficiary of a policy in lieu of implementing the policy, or the maximum amount of income that someone harmed by a policy would be prepared to pay to prevent the implementation of the policy.

[13] Maeve P. Carey, *Cost-Benefit and Other Analysis Requirements in the Rulemaking Process*, CONGRESSIONAL RESEARCH SERVICE, CRS REPORT R41974 (2014) 1-31 at 14-5.

[14] *Id.* at 15.

[15] *See* Anthony Boardman, David Greenberg, Aidan Vining, and David Weimer. *Cost-Benefit Analysis*, 4th Edition. Pearson Economic Series.

Defining the social opportunity cost of a policy is somewhat more straightforward. Namely, it is the value to society in their next best use of the resources that are used up in implementing a policy.

These measurement building blocks also apply to defining the benefits and costs of privacy controls, with appropriate adjustment for the somewhat distinctive nature of privacy markets and/or property rights.

A Simple Model of the Valuation of Online Privacy

To help organize the discussion, we begin by summarizing the main features of an economic model of privacy formulated by Savage and Waldman in their article attempting to estimate U.S. consumers' WTP to protect their personal information when using applications on their smartphones (apps).[16] Consumer use of apps is particularly instructive since a substantial quantity of market data exist regarding actual transactions conducted in the market.

In the Savage and Waldman model the individual is assumed to maximize a utility function which has as its arguments consumption (c), leisure (L), and privacy (P), which in turn is a declining function of the number of apps, a, so that P = P(a).

Thus, the consumer's maximization problem can be stated as:

$$max_{h,a}U(c, L, P(a))s.t.\ c = y + wh - p \cdot a;\ and\ L = T - h - T(a, e)$$

where *y* = unearned income, *w* is the wage rate, *h* is hours of work, and *p* is the per unit price of an app. The function T(*a*,*e*) represents the impact of using apps on the amount of time the consumer uses for essential activities (essential time), which depends both on the number of apps used (*a*), and the individual's experience in using apps, *e*. Holding *e* constant, increased use of apps is assumed to decrease the amount of essential time (e.g. result in essential time savings).

A key result of the model is that the rational consumer will acquire additional apps up to the point where:

$$-\mathrm{wT_a} = \mathrm{p} + \left(\frac{\mathrm{U_P}}{\mathrm{U_c}}\right) \cdot \mathrm{P_a}$$

In expression (2), the left-hand side is the marginal value of essential savings of the marginal app purchased: the right hand side equals the marginal cost of the marginal app which in turn is comprised of the per-unit app price, p, plus the marginal value of the privacy lost by purchasing

[16] *See* Scott Savage & Donald M. Waldman, *The Value of Online Privacy*, SSRN (2013), available at https://ssrn.com/abstract=2341311.

an additional app, $\left(\frac{U_P}{U_c}\right) \cdot P_a$.

The term $\left(\frac{U_P}{U_c}\right) \cdot P_a$ represents the marginal value to the consumer of giving up an additional "unit" of privacy at the margin, and hence represents the consumer's marginal valuation or willingness to pay for privacy.

Empirical Estimates of the Willingness to Pay for Privacy

There are several ways of estimating the willingness to pay for privacy. One can attempt to estimate the marginal willingness to pay directly using data from choices that consumers are observed to make in the market place. Alternatively, one can use data from choices that consumers are observed to make in experimental settings, or in surveys. Inferences can also be made from analogous markets, such as those that provide protections of consumer privacy. In this section, we summarize the results of such efforts.

It is worth noting in advance that the literature review on privacy provides considerable evidence that consumer privacy preferences vary substantially across different characteristics of interest. For example, studies generally find that females have higher valuations for privacy protection than males. However, females also tend to value particular kinds of privacy protections over others (e.g. location data collected via a smartphone's GPS). In contrast, males tend to value concealing their browsing history more highly than hiding their location data.

Generating valid measures of consumer privacy is also made more difficult due to the so-called "privacy paradox" which notes that consumers' stated preferences for privacy protection are often completely uncorrelated with their behavior (i.e. what they actually pay to protect their PII).[17]

1. Savage and Waldman, *The Value of Online Privacy*[18]

Savage and Waldman estimated U.S. consumers' WTP to conceal various types of personal information from companies and third parties when downloading smartphone apps. The authors posed two primary research questions: 1) what is the value of online privacy for adults in the U.S., and 2) to what extent do these valuations vary with user experience? They operationalized the concept of privacy by estimating U.S. consumers' WTP for smartphone apps in 2013. Data on downloads of apps are generally useful for informing privacy valuations because consumers are required to relinquish various kinds of private information to app developers and third parties—in addition to the actual cost of the app—to benefit from using these applications on their

[17] *See generally* Patricia A. Norberg, Daniel R. Horne & David A. Horne, *The Privacy Paradox: Personal Information Disclosure Intentions versus Behaviors*, 41 J. CONSUMER AFF. 100-126 (2007).

[18] Savage and Waldman, *supra* note 13.

smartphones.

Methodology

The research design involved administering an in-person survey to consumers, with a pre- and post-test, either in their home or a public place during the summer of 2013. Interviewers used the pre-test partly to classify participants as either "experienced" or "inexperienced" users. Interviewers then showed participants an app on the interviewer's phone that was available to download in the marketplace at the time.

Participants were told that the app developer was considering several alternate versions that were identical with the exception of different privacy permissions, prices, and whether they included advertisements. They were also told that they would have the opportunity to purchase the alternate version of their choice, which would soon be available in the market. Interviewers asked respondents two questions: 1) which app do you prefer, and 2) do you intend to download the app once it is available?

The post-test consisted of revealing to participants that the survey was conducted for research purposes only (there were no alternative apps being developed) and asking exit questions to determine how likely participants were to follow-through with their stated preference (in cases where they indicated they were doing to download an alternate version of the app once it was available).

Primary Findings

The survey data indicated that the representative U.S. consumer is willing to pay $2.28 to conceal their browser history, $4.05 to conceal their list of contacts, $1.19 to conceal their location, $1.75 to conceal their phone's ID number, $3.58 to conceal the contents of their text messages, and $2.12 to eliminate advertising per app downloaded on their smartphone. Table 1 contains a detailed list of findings, but it is worth noting here that the authors found the following characteristics to have significant effects on privacy valuations: gender, age, level of user experience, and education.

2. Acquisti, John & Loewenstein, *What is Privacy Worth?*[19]

Most of the studies summarized in the literature review attempt to generate specific estimates for consumers' WTP for privacy. However, Acquisti, John & Loewenstein focused their efforts on investigating the extent to which contextual, nonnormative factors affect estimates for privacy

[19] Alessandro Acquisti, Leslie K. John & George Loewenstein, *What is Privacy Worth?*, 42 J. LEGAL STUD. 249, 249-74 (2013).

preferences. The authors note that findings from behavioral economics and decision research frame their assumption that consumer preferences for privacy are not as consistent or easy to measure as assumed by traditional economic theorists. In short, they generally question the validity of estimates for consumers' WTP generated by research designs that tend to rely only on a single method of data collection.

Methodology

The authors conducted a field experiment that involved offering two types of Visa gift cards to female consumers shopping at a mall in the U.S. exchange for participating in a survey. The subjects were offered (under various configurations) the option of choosing between a $10 "anonymous" gift card—for which purchases would not be linked to (PII)—and a $12 "identified" gift card—for which purchases made would be tracked under their name and additional identifying information. The authors provide a summary of the five conditions used to offer gift cards to subjects; conditions one and two test for endowment effects, conditions three and four check for order effects, and condition five is a rationality check control condition (offering a $12 anonymous card or a $10 identified card) to see if participants understood the trade-offs being presented:

1. **$10 endowed**: Keep the anonymous $10 card or exchange it for an identified $12 card.
2. **$12 endowed**: Keep the identified $12 card or exchange it for an anonymous $10 card.
3. **$10 choice**: Choose between an anonymous $10 card and an identified $12 card.
4. **$12 choice**: Choose between an identified $12 card and an anonymous $10 card.
5. **Rationality check control condition**: choose between a $10 identified card or a $12 anonymous card.

Primary Findings

Over half of the participants endowed with the anonymous $10 card rejected an offer of $2 to reveal their future purchase data while over 90% of the participants endowed with the identified $12 card refused to pay $2 to protect their privacy (e.g. not accepting the offer to switch to the $10 gift card). These findings indicate that consumers' willingness to accept (WTA) is greater than or equal to $2 while consumers' WTP is less than $2. The findings of this study raise substantial validity concerns for research that does not take into account insights from behavior economics including order and endowment effects into the design of the study.

3. Beresford, Kübler & Preibusch, *Unwillingness to Pay for Privacy: A Field Experiment*[20]

The authors conducted a field experiment in the form of a revealed preference test to estimate consumers' WTP pay for privacy protection—the disclosure of their monthly income—during business transactions requiring the disclosure of PII. The findings of this article are a substantial outlier relative to the other articles included in the review.

Methodology

The experiment involved 225 participants who were students at the Technical University of Berlin; 74 of the participants providing data via the option to purchase a DVD from one of two online stores. The authors partnered with Amazon to create fictitious branches for two different retail stores that were ostensibly part of a known multichannel retailer of DVDs in Germany (SilverDisc). Both stores were set up with different privacy disclosure requirements. The treatments consisted of: 1) a scenario where both stores offered the same selection of DVDs for the same price and 2) a scenario where one store offered the same selection of DVDs but at a discount of one Euro. The store offering the one Euro discount required consumers to disclose their monthly income in exchange.

Primary Findings

The authors indicate that consumers are generally unwilling to pay for privacy. When faced with a trade-off between providing less sensitive, private information and a modest discount in price, approximately 92% of participants chose the discount. Interestingly, the experiment also seems to indicate that varying privacy disclosure requirements without varying price results in no significant effect on consumer decision making. However, it is worth reiterating here that this study's findings are a substantial outlier in the privacy literature's estimates for consumer valuations of privacy. This is likely not only a result of sample selection bias (college students) but also a result of the way that the authors chose to operationalize the concept of privacy (disclosure of monthly income).

4. Hann, Hui, Lee & Png, *Overcoming Online Information Privacy Concerns: An Information-Processing Theory Approach*[21]

Hann et al. administered a survey to estimate consumers' WTP to protect their PII during online

[20] Alastair R. Beresford, Dorothea Kübler & Sören Preibusch, *Unwillingness to pay for privacy: A field experiment*, 117 ECON. LETTERS (2012).

[21] Il-Horn Hann, Kai-Lung Hui, Sang-Yong Tom Lee & Ivan P.L. Png, *Overcoming Online Information Privacy Concerns: An Information-Processing Theory Approach*, 24 J. MGMT. INFO. SYS. 13, 13-42 (2007).

transactions. The authors administered the survey to university students from both the U.S. and Singapore. The survey questions asked as part of the pre-test are motivated by the authors' choice to conduct a conjoint analysis of the data based on the expectancy theory of motivation.[22] The pre-test involves asking participants to rate their reasons for valuing privacy across several dimensions. Answers from the pre-test were later compared to results of stated valuations to determine if there were any significant drivers that motivate participants to prefer more or less privacy under different contexts.

Methodology

The authors administered a survey to undergraduate students in both the U.S. and Singapore. The students were first asked to rank their level of concern for privacy (generally) and then asked to rank specific reasons that motivated that belief. The participants then made a series of choices concerning the use of websites that facilitated transactions for different industries (financial, health care, and travel).

The websites presented to participants varied in two ways: 1) cost and 2) the ability given to users to manage the private information they would be required to disclose to websites in order to use them. Privacy management was broken down into three areas: 1) users' ability to review (and correct) private information disclosed to websites, 2) the ability to restrict private information against improper, third-party access, and 3) the ability to prevent private information from being used for secondary uses (e.g. by someone other than the website for marketing purposes).

Primary Findings

The authors found U.S. participants' privacy to be worth between $30.49 and $44.62 (annually/person) while participants from Singapore valued their privacy at an average value of $57.11. Additionally, based on their pre-test questions, the authors claimed to have identified three distinct segments of Internet users: privacy guardians, information sellers, and convenience seekers.[23] The author's breakdown of their survey results is as follows:

Value of Privacy (in U.S. dollars)[24]

Web site privacy policy	United States	Singapore
Review for error	$11.18-16.36	$10.45
Restriction against improper access	$11.33-16.58	$19.73
Secondary use not allowed	$7.98-11.68	$26.93

[22] *Id.* at 21-4.

[23] *Id.* at 30.

[24] *Id.* (modified from authors' Table 3).

5. Schreiner & Hess, Why Are Consumers Willing to Pay for Privacy? An Application of the Privacy-freemium Model to Media Companies[25]

In addition to estimating a WTP for privacy protection, this study applied the theory of planned behavior (TPB) to explain the necessary conditions under which consumers would be willing to pay for additional privacy protection when using online content platforms (e.g. Facebook). Similar to the study by Acquisti, John & Loewenstein, the authors generated valuable evidence regarding the contexts that shape consumer's preferences for privacy protection. Additionally, they provide a model—a framework based on TPB—that is useful for conceptualizing the various drivers and forces shaping consumer preferences to pay for privacy protection.

Methodology

The authors administered an online survey to 553 Facebook users in Germany. The survey involved deceiving participants into believing that they were being offered the opportunity to bid on a soon-to-be-released premium version of Facebook with additional privacy control features in return for paying a monthly fee. The auction and deception components were valuable for estimating WTP via revealed (rather than stated) preferences.

A pre-test was administered to operationalize participants' motivations for privacy preferences across seven potential drivers: 1) attitude, 2) intention, 3) perceived behavioral control, 4) perceived internet privacy risk, 5) perceived usefulness, 6) subjective norms, and 7) trust. Measures for each area were estimated via respondent answers to questions within each category on a seven-point Likert scale. The authors used these results to describe the potential drivers of their consumer WTP estimate. The following is an illustration of their research model:[26]

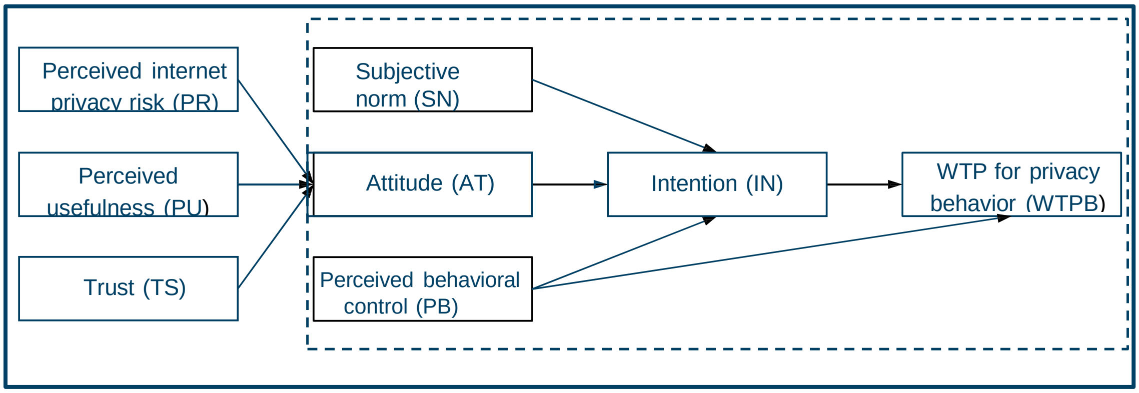
Perceived internet privacy risk (PR)
Perceived usefulness (PU)
Trust (TS)
Subjective norm (SN)
Attitude (AT)
Perceived behavioral control (PB)
Intention (IN)
WTP for privacy behavior (WTPB)

[25] Michel Schreiner & Thomas Hess, *Why Are Consumers Willing to Pay for Privacy? An Application of the Privacy-freemium Model to Media Companies*, 164 ECIS COMPLETED RESEARCH PAPERS 1, 1-15 (2015).

[26] *Id*. at 6.

Primary Findings

The authors estimated a consumer WTP for additional privacy protection when using online content platforms like Facebook of 0.63 Euros per month. They also performed an analysis on the various motivational coefficients (captured during the pre-test) to determine their relationship to revealed consumer WTP estimates. The authors found that participants' perceived usefulness (PU) and level of Trust (TS) in the fictitious premium version of Facebook on offer significantly affected consumers' attitude (AT) about subscribing. In this model, PU is a measure of the extent to which users believe that the privacy solutions on offer via the premium version of Facebook are likely to address their privacy concern. Trust is a measure of the degree to which users believe Facebook is a trustworthy company. Attitude is a more direct measure of participants' perception of actually subscribing to the premium version.

Interestingly, the authors did not find a significant relationship between consumers' level of perceived internet risk (PR) and their attitude towards the premium version of Facebook. Overall, PR, PU, and TS explained 52% of the observed variance in AT (under the causal assumptions of the model). Finally, it is worth noting that Subjective norms (SN) were estimated to also have a significant effect on Intention (IN).

6. Cvrcek, Matyas, Kumpost, & Danezis, *A Study on the Value of Location Privacy*[27]

The authors conducted a survey to estimate the value of privacy—defined as participants' willingness to accept payment in exchange for the use of their mobile phone data to track their location/movement on a daily basis for a month. The auction involved the use of deception to convince participants that they were submitting bids to receive actual compensation in exchange for disclosure of their location data. The experiment indicated that several consumer attributes may drive consumers' WTA payment for location data including: gender, nationality, the use of data (academic vs. commercial), and the duration of collection.

Methodology

The study involved surveying 1,200 people from five different countries: Belgium, the Czech Republic, Germany, Greece, and Slovenia. The survey was structured using three separate auctions: 1) a one-month study with tracking data to be used for academic purposes only, 2) a one-month study where tracking data would also be used for commercial purposes, and 3) a year-long study that extended the conditions of the second case. It is worth noting briefly that the authors re-calculate the values of bids submitted across different countries using a "value of

[27] Dan Cvrcek, Vashek Matyas, Marek Kumpost & George Danezis, *A Study on the Value of Location Privacy*, PROCEEDINGS OF THE 5TH ACM WORKSHOP ON PRIVACY IN ELECTRONIC SOCIETY, 109, 109-18 (2006).

money" coefficient, computed as a ratio of average salaries and price levels within a particular country.

Primary Findings

The median bid (using exchange rates in August 2006) was 43 Euros under the condition where participants chose to disclose their location data for academic purposes during a period of one month. A breakdown of the data illustrates substantial variation among participants with certain characteristics. For instance, females' bids for the first condition were similar to males, but were 1.4 times higher for commercial use and 1.8 times higher for extending the study from one month to one year. Finally, participant nationality also affected bids. For example, German and Slovak bids were five times the median bid.

Synthesis of Article Findings

A look across the findings within the articles reviewed yields valuable information for future research to generate estimates for consumer valuations of their privacy. Table 1 contains a detailed list of findings for each study. The following are several key takeaways.

Operationalizing Privacy is Highly Context Dependent

The studies demonstrate that, although "privacy" is a complex concept, thoughtful research designs can generate useful estimates of consumers' WTA or WTP for privacy. However, these estimates are most valid given a context-specific definition of the privacy issue in question. For example, it would be of questionable validity to say with any certainty that an individual's privacy (broadly speaking) is worth $X to them; it is substantially more plausible to state that male consumers in the U.S. using social media platforms online are willing to pay $X every month to prevent private companies from sharing their browsing information with third parties.

Interestingly, since the use of various technologies require almost identical kinds of privacy disclosures (e.g. location tracking), estimates that are sufficiently well specified (i.e. location tracking provided to whom for what duration, etc.) are transferable across conditions. This is particularly valuable for regulatory agencies—all of which work under considerable time constraints—as it prevents them from having to "reinvent the wheel" to find estimates for the benefits and costs of consumer privacy that can be used as evidence to support their rulemaking.[28] This applies even to cases where agencies are considering regulation of emerging technology.

[28] *See* Ray Pawson, *The Science of Evaluation: A Realist Manifesto* (2013) (discussing the transferability of knowledge generated across fields and institutions).

Privacy Valuations are Not Necessarily Stable

Acquisti, John & Loewenstein point out that most studies within the privacy literature operate under the assumption that a rational consumer's WTP and WTA should be equal. In fact, the literature review indicates that consumer valuations fluctuate substantially under different conditions and are highly dependent on certain decisions made in the research design. For example, researchers should play close attention to the role that endowment effects can have in driving estimates of consumer privacy. This applies to both stated and revealed preference studies. Do participants begin with a default expectation of privacy? Are consumers paying for a benefit they don't currently have or are they being offered money in exchange for disclosure of their PII?[29]

Even estimates of well-specified privacy conditions can vary with minor differences—such as changes in the recipient of PII (even within the same industry). For instance, consumers might state or reveal certain WTP to protect their data from a company they consider "trustworthy" but may be willing to pay substantially more to protect their PII from a company they (personally) consider "untrustworthy."[30]

Improving the Validity of Privacy Estimates

The most convincing research efforts seem to make use of auctions and/or deception to more closely approximate actual consumer market behavior. Assuming the privacy paradox remains valid, research designs generating estimates using participants' stated preferences are not likely to yield valid results. Designs that either: 1) require participants to make purchases with their own money or 2) successfully deceive participants into believing that they are receiving payment (or studies that actually do pay participants) using an auction system are more likely to generate more useful estimates of consumer valuations of their PII.

Consumer Characteristics Matter

Finally, research designs that treat consumers as a homogenous group are unlikely to produce useful estimates. Almost all of the studies covered by this review indicate that consumer characteristics are highly correlated with their valuations of particular kinds of privacy. For example, gender may affect WTP for certain privacy areas (location) but not others (duration).[31] Country (a proxy for admittedly difficult to conceptualize cultural differences) also affects privacy valuations. This is worth noting, in particular, because it presents substantial limits on the estimates that U.S. regulatory agencies can use to support their rulemaking (i.e. consumer

[29] *See* Acquisti, John & Loewenstein, *supra* note 16.

[30] Cvrcek, Matyas, Kumpost & Danezis, *supra* note 24.

[31] *Id.*

valuations of privacy in Singapore or Germany are likely to vary considerably relative to consumer valuations of privacy in the U.S.).[32]

Conclusion

Adam Thierer,[33] who argues for the need for benefit-cost balancing in evaluating privacy regulations, also notes that the empirical data needed for such balancing may be difficult to gather. Our survey offers a somewhat more optimistic view.

Although estimating the economic value of privacy is challenging, it is not impossible. Estimation of the social costs of implementing privacy regulations are comparable in difficulty to estimation of social costs in other policy areas. Not surprisingly, estimating individual willingness to pay to protect privacy is more difficult. However, both theoretical and empirical frameworks exist for doing so. Indeed, there appears to be enough of an empirical literature to provide “plug-in” values of both the social benefits and social costs of privacy regulations to be used in undertaking benefit-cost analysis. An important next step will be to adapt such estimates for the purposes of undertaking actual and proposed regulation of privacy.

[32] *Id. See also* Hang, Lee & Png, *supra* note 19.

[33] *Ibid*.

Appendix: Empirical Estimates of Consumer Privacy Valuations

Study	Country	Empirical Estimates	Additional Findings
Savage & Waldman (2013)	U.S.	U.S. consumer WTP for privacy (per app): • $2.28 to conceal browser history • $4.05 to conceal list of contacts • $1.19 to conceal location data • $1.75 to conceal unique phone ID • $3.58 to conceal text messages • $2.12 to eliminate advertising Given typical app in U.S. marketplace: • Benefit of app must be at least $5.06 • Estimated $17.08 billion benefit of app marketplace	• WTP varies substantially with level of user experience • Consumer preferences are heterogeneous and vary across race, gender, income, education, and level of technological experience.
Acquisti, John & Loewenstein (2013)	U.S.	U.S. consumer WTP ≠ WTA to conceal purchasing data: • WTA ≥ $2.00 • WTP < $2.00	• Privacy estimates generated are sensitive to framing of research design (*e.g.* endowment and order effects) and other contextual, nonnormative factors.[34]
Beresford, Kübler & Preibusch (2012)	Germany	German consumer WTP to conceal monthly income during online purchases < 1 Euro.	
Hann, Hui, Lee & Png (2007)	U.S. and Singapore	Consumer WTP to protect PII across 3 different categories during online purchases (protection against errors, improper access, and secondary use of personal information): • Between $30.49 - $44.62 in the U.S. • $57.11 in Singapore	• Participants from Singapore valued privacy more highly relative to U.S. participants. • Study identifies three distinct groups of subjects based on behavior toward privacy: privacy guardians, information sellers, and convenience seekers
Schreiner & Hess (2015)	Germany	WTP for additional privacy protection when using online content platforms like Facebook of 0.63 Euros per month.	• Consumer WTP for privacy protection highly contingent upon the *perceived trustworthiness* of the company making the offer and the *belief* that the product addresses the underlying privacy concern.
Cvrcek, Matyas, Kumpost, & Danezis (2006)	Belgium, the Czech Republic, Germany, Greece, & Slovenia.	Participants' median WTA for disclosure of location tracking data (6 months, for academic purposes) = 43 Euros.	• Consumer valuations of PII highly contingent upon *recipient* of PII (*i.e.* academic vs. commercial) and *duration* of tracking. • WTA payment for location data varies substantially across characteristics including *gender* and *nationality*.

[34] See Acquisti, John & Loewenstein, *supra* note 16 at 249.