Januar 14, 2009

ESMT WORKING PAPER

ESMT No. 09-001

An Empirical Approach to Understanding Privacy Concerns

Luc Wathieu, ESMT
Allan Friedman, John F. Kennedy School, Harvard University

ISSN 1866-3494

Abstract

An Empirical Approach to Understanding Privacy Concerns

Authors*: Luc Wathieu, ESMT
Allan Friedman, John F. Kennedy School, Harvard University

This paper shows that privacy concerns in commercial contexts are not solely driven by a desire to control the transmission of personal information or to avoid intrusive direct marketing campaigns. When they express privacy concerns, consumers anticipate indirect economic consequences of data use, such as price discrimination. Our general hypothesis is that consumers are capable of expressing differentiated levels of concerns in the presence of changes that suggest indirect consequences of information transmission. We suggest that there is a homo economicus behind privacy concerns, not simply a primal fear. This hypothesis is tested in a large-scale experiment evoking the context of affinity-based direct marketing of insurances, which relies on data transmitted by alumni associations. Because opt-in and opt-out choices offered by firms to consumers usually capture non-situational preferences about data transmission, their ability to enact privacy concerns is questioned by our findings.

Keywords: privacy, opt-in/opt-out, insurance

JEL Classification: D8, M38

* Contact: Luc Watthieu; ESMT, Schlossplatz 1, 10178 Berlin, Germany; Tel: +49 (0)30 212 31-1284
Email: watthieu@esmt.org.

1. INTRODUCTION

There are many perspectives on how privacy sentiments manifest themselves and should be addressed in society. Recently, with the rise of electronic commerce and database marketing, researchers have increasingly interpreted the demand for privacy in economic terms. Some authors and commentators suggest that consumers view targeted marketing communications as a costly annoyance (Hann et al. 2008), while other researchers posit that finely informed firms might take actions detrimental to the consumer surplus of some market segments (Hermalin and Kast 2006; Odlyzko 2004; Taylor 2004; Wathieu 2006). While such research has generated many realistic predictions, the connection between its basic assumptions and the motives behind ordinary privacy concerns remains largely un-documented. The purpose of this paper is precisely to suggest and test a set of behavioral hypotheses concerning the sensitivity of consumers towards the potential economic consequences associated with the dissemination of personal information.

Many observers have noted the existence of a "privacy paradox" in that consumers at the same time (1) routinely declare that they value their privacy highly (Westin 1998, Hann et al. 2007), but (2) do not seem to actively incorporate privacy concerns in their transactions (Berendt, Gunther, and Spiekermann 2005). This paradox might suggest that consumers are too unsophisticated to envision (or even to sense) the economic consequences of transmitting their personal information. It also suggests that consumers inconsistently engage in market behaviors that they oppose in principle. Such interpretation would be bad news for economic theories of privacy demands. However, an alternative interpretation of the privacy paradox could be that, accounting for the complexity of the anticipated consequences and the lack of means at their disposal, consumers currently feel unable to enact their privacy preferences (Shostak and Syverson, 2004). In particular, the use of "opt-in" and "opt-out" choices to allow

consumers control over their information assumes that privacy concerns are tightly associated with a consumer's ability to control the *transmission* of personal information to self-interested actors. This paper will argue that opt-in and opt-out choices do not allow consumers to completely enact their privacy preferences, because privacy preferences are not tightly linked to own information dissemination.

In order to determine whether consumers are sensitive to the economic implications of privacy beyond data dissemination, this paper proposes an experiment based on a real-world situation involving the transmission of personal information in a commercial context. Much of the earlier research aimed at highlighting and understanding demands for privacy has been survey-based, focusing on non-situational antecedents of privacy demands (Cranor, Ackerman and Reagle 2000, Westin 1998). An experimental approach is suitable to assess the impact of context modifications on the privacy sentiment.

There have been a limited number of experiments attempting to study privacy valuation. Huberman, Adar, and Fine (2005) suggest that privacy valuation is a function of perceived deviance. While this finding helps clarify the strength of some individuals' preferences, deviance cannot explain all preferences, particularly when privacy applies to data that does not fit in a normal/deviant framework (e.g., name and address on a mailing list). Rational privacy protection behavior has been isolated in a study with very explicit information on risks and rewards (Poindexter, Earp, and Baumer 2006) which does not tell us whether people's natural notion of privacy usually encompass such consequences. A series of detailed, interactive surveys by Acquisti and Grossklags (2005a) questioned the model of a rational privacy-protecting consumer, in an analysis that included a broad range of privacy lifestyle choices but did not directly induce trade-offs between information transmission and economic benefits. Such trade-offs can be studied by watching user behavior (Acquisti and Grossklags 2005b, Chellappa and Sin 2005, Berendt, Gunther, and Spiekermann 2005) or in conjoint analyses to derive the perceived

value of resolving privacy concerns (Hann et al. 2007). While these research approaches provide an important understanding of privacy sentiments in a specific context, or a useful dollar value, it is difficult to apply them in a broader context where the implications of information transmissions can be complex and indirect.

After a discussion of hypotheses and conceptual issues in the next section, Section 3 outlines the experimental data generating process. In Section 4, key empirical results are described, and Section 5 concludes.

2. ASSESSING THE NATURE OF PRIVACY CONCERNS

2.1. General Framework

We assume a general framework in which firms seek to gain consumer information in anticipation of a profitable course of action. This course of action could take many forms, including internal systems development and improvement, targeted marketing (Milne and Rohm 2000), loyalty programs (Deighton, 2000) or maximizing profits through price discrimination (Hermalin and Katz 2006). In most such cases, the firm is the driving actor to collect and/or use personal information. Whether this raises a privacy problem depends on the consumer's reaction. In extreme cases, the consumer impact is obviously positive (e.g., when an emergency medical practitioner obtains life-saving information) or negative (e.g., when it results in unwanted telephone solicitation), and predicting consumer reaction is trivial. What is less understood is the reaction of consumers when trade-offs are more balanced, with subtler benefits or less obvious cost or risk factors implied.

In particular, this paper differentiates between a *direct* privacy concern and an *indirect* privacy concern, and argues that the second form, while subtler, constitutes a measurable influence on ordinary privacy demands in commercial contexts. A direct privacy concern is motivated by harm from information release that is immediately perceived by the offended party. For instance, a fear of impersonation fraud (identity theft) or a dislike of direct marketing represent disutilities that are directly attached to the transmission of consumer information. An indirect privacy concern, in contrast, is predicated on multiple steps between the transmission of personal information and the resulting impact in terms of other variables that affect the consumer's well being. These consequences could appear in terms of access to low prices, product variety and quality, or even economic growth. Indirect privacy concerns are necessarily harder to isolate, as the impacted variables are a function of many other decisions and data beyond the collection of personal information.

Privacy concerns are both direct and indirect in nature. For example, when thinking about government surveillance and airport security, a direct privacy concern would be the fear of a stranger intruding your intimate space, but citizens have also expressed annoyance at the delays resulting from *others* receiving such treatment, or complained from the fact that *everyone* loses rights when even a few people are unfairly treated as suspects. Interestingly, as noted by Wathieu (2006), such indirect privacy effects are not necessarily attached to the individual's personal information disclosure, they are related to harm incurred by the overall system when people's privacy is restricted.

In connection with marketing information, one can also highlight the distinction between direct and indirect privacy concerns, with greater doubt cast on the empirical relevance of the latter. Indeed, it should be fairly straightforward to show that individual consumers are concerned about exposing personal peccadilloes to marketers. Similarly, if revealing a telephone number or email address leads to the annoyance of telemarketing or spam, a theory claiming reluctance to reveal information should be

uncontroversial. In contrast, the fear of price discrimination or market dysfunctions that might result from consumer exposure (Hermalin and Kast 2006; Odlyzko 2004; Taylor 2004; Wathieu 2006) requires the consumer to understand (or at least to sense) the fact that personal harm can accrue from the collection of everyone's information to gauge demand. Indirect effects on variables of interest for consumers are critical for an interpretation of privacy debates from an economic perspective, but it remains to be shown whether and how such indirect effects relate to the ordinary experience of privacy concerns by real-life consumers.

Our general hypothesis is that *consumers are capable of expressing differentiated levels of concerns in the presence of changes that suggest indirect consequences of information transmission*. In other words, we suggest that there is a *homo economicus* behind privacy concerns, not simply a primal fear.

If consumers do not, in fact, have a sophisticated understanding of indirect privacy effects, then they will not be concerned with subtle factors in a given context, nor will they appreciate factors that only affect the indirect concern without triggering a direct, immediate privacy concern.

2.2. Hypotheses on the Causes of Privacy Concerns

If a direct utility of privacy were driving privacy concerns, information dissemination would be the critical trigger of concern for consumers, with more information transmission causing a greater concern. In contrast, if consumers are more sophisticated, so that privacy concerns anticipate possible indirect consequences that occur when firms acquire finer-grained consumer knowledge, we should observe relative indifference towards mere information dissemination across databases when it is clearly inconsequential. As well, consumers should be able to develop privacy concerns when their own data is not transmitted, while some *other* consumers have had their data disseminated in a way that can indirectly affect everyone else's deal on the market (this is similar to opposing searches at airport or the

monitoring of phone calls, even when you know that you are unlikely to be searched or monitored). Under the conventional thinking regarding privacy, there should be no privacy concern if personal information is not transmitted at all. By including indirect privacy effects, a concern might arise even when the consumer's own personal information is protected. For instance, the fact that other consumers transmit their information might lead to structural changes (e.g., in terms of monopolistic positions, or of the amount of variety available) that affect a non-transmitter and should cause a reaction in defense of privacy (as in Wathieu 2006). This logic even applies in the case of a privacy concern motivated by impersonation fraud, as consumers absorb the added costs of the misuse of others' identifiers.

In sum, this discussion suggests that *an increase in data transmission is neither necessary nor sufficient to cause a privacy concern.* This notion leads to the following more formal and testable hypotheses, which contradict the ordinary intuition that privacy concerns are a purely a matter of individual control on the transmission of one's own information (e.g., through opt-in and opt-out choices). How these hypotheses can actually be tested will become evident in Section 3.

H1 (Use Without Transmission): *The use of personal information can cause a significant privacy concern even if it is dissociated from data transmission.*

H2 (Extended Transmission With or Without Change in Use): *The transmission of an enlarged set of data is causing a larger privacy concern when the additional information transmitted is more useful (relevant) to the firm.*

H3 (Determinants of Privacy Concerns): *Indirect economic implications of information transmission are relatively more significant determinants of privacy concerns and behavior than dissemination itself.*

Hypothesis 2 raises the question of data relevance. Naïve approaches to the privacy concern would seek protection of any kind of personal data (with perhaps an emphasis on personally identifying pieces of data). The results in Hann et al. (2007) even suggest that personal valuation is independent of personal context. But if the privacy concern is driven by indirect consequences of data usage by marketers, privacy demands should, in principle, be greater towards data that is more likely to be used by the firm who collects it (i.e., data that can be leveraged more profitably) when this use could be to the subject's detriment. For instance, if a consumer is worried about obtaining health insurance in a given context, then sharing family medical history should cause concern, while if the consumer is confident in his future health coverage, sharing the history is less of a concern.

2.3. Enactment of Privacy Concerns

Another aspect of consumer behavior that can help assess whether consumers perceive the indirect implications of personal data transmission is their response to policy solutions in response to privacy concerns. If consumers think of privacy only in terms of direct disutility upon disclosure of information about themselves, we can expect that control levers such as the ability to opt-in and opt-out will be deemed attractive and sufficient. In contrast, a consumer's distinct call for regulation or intermediation (broadly speaking: any collective intervention to limit the transfer of data concerning a group of people) can only be understood in light of a perception of interdependence of individual (and indirect) consequences. In particular if H1 is true, perceived harm from consumer exposure can occur whether or not the individual can control his or her own individual participation in the data transfer. While personal participation preferences may not be strongly applicable in situations that suggest indirect effects, the role of the social planner (or of a representative body) becomes more important. If individuals are

affected by the actions of the group, then individuals should sense that the solution lies with group (or intermediated) action. This gives our final hypothesis:

H4 (Limits of Personal Control): *Even when allowed to control their privacy preferences through opt-in and opt-out choices, consumers value group decisions that regulate information use by firms.*

3. RESEARCH DESIGN AND DATA

To better understand how consumers treat information privacy in a complex environment and test the above hypotheses, we presented participants with a realistic scenario involving the dissemination of personal information in a commercial context, and measured their response through a brief survey. A scenario-based experiment was deemed appropriate because it would allow experimental manipulations while evoking a relevant, relatively natural, relatively complex situation.

There were twelve experimental conditions, each involving a specific modification of the same baseline scenario. A manipulation check questionnaire was also applied to verify that respondents and researchers shared the same interpretation of the various scenarios.

Respondents were 646 randomly selected members of a subject pool maintained by the research center of a business school in the United States. This subject pool features over 10,000 members diverse in background and gender, including business and undergraduate students who accounts for 45% of the population. Respondents were recruited by email and participation was voluntary, with a $5 payment upon completion. The experiment was administered via a website. It was made clear to respondents that there was no right or wrong answer. The average experimental group included 54 respondents, with no group having fewer than 48 respondents.

3.1. Control Scenario

To evaluate the theoretical hypotheses presented above, we looked for a realistic and intuitive situation where consumer data were disseminated in a way that might (1) allow the consumer to access advantageous offers, (2) expose the consumer to marketing hassle, and (3) have likely indirect consequences in terms of the consumer's welfare.

Affinity-based direct marketing of car insurance contracts provided such a context. This marketing process, documented in a case study by Wathieu and Morris (2004), uses the membership databases of trusted associations (such as alumni associations) to channel targeted deals to their members, through direct communications means that blend direct mail and telemarketing. When associations negotiate such deals, often for considerable fees, they have an interest in minimizing potential hassle for their members, and they also seek to minimize the possibility that marketers discriminate among different types of members, in order to maintain membership cohesiveness. Governments, on the other hand, monitor the impact of these arrangements on competition and the industrial structure. The control scenario, which serves as a baseline for our analyses, is evoking one such arrangement between an alumni association and a car insurance company:

> *As a service to its members your college alumni association has negotiated a special deal with a well-known car insurance company.*
>
> *The insurance company will use data (including members' name and contact information) on a one-time basis to offer alumni (via a mail and phone marketing campaign) an alumni association-endorsed deal featuring first-class service levels and a 30% discount on annual insurance premiums.*
>
> *Based on certain parameters specified by the insurance company, data for 20% of the alumni have been transmitted to the insurance company and all of these alumni are about to be offered the deal. At this point it is still unknown whether you are among the beneficiaries of this deal.*

The underlined parts of the scenario are those privacy-sensitive aspects that will be modified in experimental conditions.

3.2. Behavioral Measures

The scenario itself did not explicitly offer the respondents a choice. However, after reading the scenario, respondents were asked 10 questions (7 answers were on 7-point Likert scales, the last 3 answers were Boolean, emphasis added here for readability only):

- *This is an example of a situation in which **I am concerned about privacy**.*
- ***How happy are you** that this deal was struck between your alumni association and the car insurance company?*
- *In this instance, how **fairly** do you feel your alumni association is treating you?*
- *Are you **fearful** that this kind of activity in the insurance market might ultimately reduce your access to a low-premium contract?*
- *Alumni **should be given an opportunity to opt-out** (withdraw) from this program before their data is transmitted.*
- *Alumni should be included in this program **only if they specifically sign up** before their data is transmitted.*
- *I would like this kind of initiative to be reviewed and voted on (either banned or explicitly authorized) by **the Board of Alumni**.*
- *Given the opportunity to **opt-out** of (withdraw from) this program before your name is actively considered for this deal, would you do so?*
- *If it were necessary but easy to **opt-in** (sign up) to have your name actively considered for this deal, would you do so?*
- *If you were on the Board of Alumni and were requested to **vote for or against** this initiative, what would be your inclination?*

3.3. Experimental Conditions

Experimental conditions changed the baseline scenario by inserting one or more of the following five modification:

Dissemination. Instead of assuming that the alumni association would transmit data parsimoniously (underlined part of scenario starting with “Based on certain parameters specified by the insurance

company data for 20% of the alumni have been transmitted…"), some participants were told that "Data for all the alumni have been transmitted to the insurance company and, based on certain parameters certified by the insurance company, 20% of the alumni are about to be offered the deal." As a result of this manipulation the likelihood of data transmission has increased from 20% to 100%, while no other significant change is taking place,[1] an example of what could be called mere dissemination.

More relevant data. This modification implies increasing transmitted data to include education and occupational data, commonly viewed (based on conversations with professionals) as relevant for an insurance company trying to assess client risk. "Name and contact information" is accordingly replaced by "name, contact information, degree obtained and year, honor student status, GPA, and current occupation." Manipulation-check respondents rated each of these elements as highly useful to predict whether a person is a safe driver or not.

More irrelevant data. This scenario modification increases transmitted data to include data that is personally meaningful, but less likely to be used by an insurance company assessing client risk: "name, contact information, membership in college associations, city of birth, and city of residence at college registration time." These additional elements were seen as least relevant as predictors of safe driving in the manipulation check questionnaire.

Priming. To increase the salience of a risk of discriminative practices by better-informed insurance companies, the following paragraph was sometimes inserted before the baseline scenario's last paragraph: "Some have wondered whether the premium paid by ordinary drivers can stay low if car insurance companies continue to use databases to offer special deals to consumers predicted to be 'safe

[1] Responses to a manipulation check questionnaire confirmed that respondents in the target population reliably agreed with this interpretation of the manipulation.

drivers.'" Manipulation checks used 7-point scales to verify that respondents found this statement both clear and legitimate.

No personal benefit. Some respondents were told that they were not beneficiaries of the deal. The last phrase of the baseline scenario was replaced by "it has become clear that you are not among the beneficiaries of the deal." By implication, these respondents were excluded both from data transmission and from the benefits of the deal.

For a parsimonious test of the individual impact of each modification against the control condition we only needed five experimental conditions in addition to the control. Four additional experimental conditions were added to measure, in the presence of the "dissemination" modification, the impact of each of the other four modifications. Finally, to further scrutinize the potential role of fairness, the condition that combined (priming indirect concern, dissemination, no personal benefit) was also included in the experiment, leading to a total of 12 experimental groups.

4. RESULTS

4.1. General Observations

While the deal proposed here to alumni offers a clear benefit to 20% of the members and does not have any specific negative implication for the other members, it turns out that only 37% of respondent were happy (score of 5, 6, or 7 on the happiness scale) and 36% were unhappy (score of 1, 2, or 3 on the happiness scale, where 1 meant "extremely unhappy") that this deal was struck. In terms of privacy response, 64% of the respondents have at least chosen some form of reactance; 40% would not opt-in, 46% would opt-out, and 51% would lean towards a ban of the proposed deal. 62% of respondents have given a relatively high privacy concern score of 5, 6, or 7 (on the seven point scale).

Table 1 gives the mean privacy responses for each of the twelve experimental conditions, with indication of significance when the response obtained is statistically different from the response in the control group. Dichotomization of the 7-point scale of relative sentiment to a simple yes/no Boolean variable was occasionally used, to characterize "concerned consumers" as those who gave a rating of 6 or 7.

Table 1: Mean Response (Privacy Concern)

Conditions		**Privacy Concern** (1-7 scale)	(% w/6 or 7)
C1 (n = 50)	Control	4.16	34%
C2 (n = 57)	Dissemination	4.86*	38.6%
C3 (n = 50)	More relevant data	5.26***	50%*
C4 (n = 59)	More relevant data + Dissemination	4.95**	47.5%*
C5 (n = 53)	More Irrelevant Data	4.70	43.4%
C6 (n = 53)	More Irrelevant Data + Dissemination	4.70	43.4%
C7 (n = 54)	Priming	4.48	42.6%
C8 (n = 48)	Priming + Dissemination	4.77	37.5%
C9 (n = 58)	No Personal Benefit	4.43	36.2%
C10 (n = 52)	No Personal Benefit + Dissemination	4.77	38.5%
C11 (n = 55)	Priming + No Personal Benefit	4.76	32.7%
C12 (n = 57)	Priming + No Pers. Ben. + Dissemin.	5.05**	42.1%

*Significance of difference w.r.t. C1: *** = (p < .01), ** = (p < .05), * = (p < .1)*

A few notes on the control group's response are in order. With 2/3 of the respondents placing their level of concern at 4 or higher out of 7 (1 meaning "Not at all concerned" and 7 "extremely concerned"), the control group already appears concerned about privacy. While the respondents were concerned, they were not dissatisfied with the offer made to them: over 80% recorded a 6 or a 7 when asked if they were happy that a deal was struck between their alumni association and the car insurance company. The control group reveals a concerned population that is nonetheless open to making a trade-off between

personal data dissemination and direct marketing interruption on the one hand, and an opportunity to access a better deal on the other hand.

4.2. Use Without Transmission

It is remarkable to observe that a significant privacy concern is registered in conditions C9 and C11, where respondents are told that they have no personal involvement whatsoever in the deal being struck between the alumni association and the insurance company. No data is being transmitted, no direct marketing will take place, no advantageous condition will be offered. Note, however, that the privacy concern is not without a logical basis: the respondent's data was *used* (but not transmitted or associated with any specific costly or beneficial action) to the extent that it was looked at to determine whether a deal would be offered. Thus, this is a confirmation of H1, whereby the use of personal information can cause a privacy concern even if it is dissociated from data transmission. In fact, the results show that privacy concerns are independent of whether a piece of data is changing hands. There is no significant difference between the privacy concern measures (ratings or proportion of concerned consumers) between C9 and C1, nor between C11 and C7.

4.3. Transmission With or Without Change in Use

It appears from the comparison of privacy concerns in C3 vs. C1 and in C5 vs. C1 that the impact on privacy concerns associated with the transmission of a larger set of data causes greater increase in concern when the additional data is relevant for use by the recipient, which corroborates hypothesis 2. Remarkably, the transmission of additional but irrelevant personal data (C5 vs. C1) does not cause a significant increase in privacy concern, indicating that mere data transmission without implied change in the recipient's behavior is disconnected from privacy concerns.

To further this discussion, one can focus on the mere dissemination conditions: all even-numbered conditions, in which respondents faced the same situation where 20% of the alumni will be offered a deal on car insurance, but everyone was told that their data would be transmitted to the insurance company. Consistent with H2, holding everything else constant, going from a 20% chance of having ones' data disseminated to a 100% certainty of having ones' data disseminated, which was inconsequential, did *not* cause an increase of the privacy concern. Table 2 shows the impact of mere dissemination on privacy concerns for each condition. The control condition comes closest to significance with a *p*-value of .0516. This may suggest that mere dissemination is raising eyebrows in scenarios where consumers are not distracted by another potential source of privacy concern.

Table 2: Mere Data Dissemination Impact on Privacy Concern, Depending on Baseline Conditions

Baseline condition	Impact on privacy concern (1-7 scale)
Control (C1 → C2)	-0.702, $p = 0.0516$
More relevant data (C3 → C4)	0.31, $p = 0.3777$
More irrelevant data (C5 → C6)	0, $p = 1$
Priming (C7 → C8)	-0.28, $p = 0.4617$
No personal benefit (C9 → C10)	-0.33, $p = 0.3394$
Priming + No personal benefit (C11 → C12)	-0.28, $p = 0.373$

4.4. Determinants of Privacy Concerns

More evidence of the relative insensitivity of consumers towards inconsequential dissemination transpires from regression analyses of the attitudinal and behavioral measures of privacy concern across the entire data set, reported in Table 3. Fairness judgments and the perception that economic side effects might occur (based on the third and fourth questions asked to respondents) are significant determinants

of privacy concerns and associated behaviors, but mere dissemination, or the transmission of irrelevant data, are not, consistent with hypothesis 3. It is interesting to observe that when consumers are told that they are excluded from the offered deal ("No Personal Benefit"), there reaction is to opt-out less, opt-in more, and vote more in favor of the deal. Thus, instead of despising direct marketing campaigns that do not include them, respondents in this study are open to a deal, as long as they are fairly treated and don't fear price discrimination. Another potentially interesting finding is that priming (highlighting the potential economic drawbacks in terms of price discrimination) has a significant effect only when it comes to a respondent's desire to support or ban the intrusive direct marketing campaign through a vote. It may well be that the way consumers are invited to express their privacy concern (e.g., through a vote vs. through opting out) frames their sensitivity towards the individual vs. social consequences of their decisions.

Table 3: Regression results: determinants of privacy attitude and behaviors

(n=646)	*Privacy Concern*	*Request opt-out*	*Request opt-in*	*Request board dec.*	*Would opt-out* ^	*Would not opt-in* ^	*Would vote against* ^
Intercept	4.28***	6.57***	5.98***	4.76***	0.40	0.78**	0.96**
Dissemination	0.16	0.03	-0.04	0.09	0.03	0.06	-0.29*
More relevant data	0.38*	0.02	0.17	0.13	-0.03	-0.32	0.16
More irrelevant data	0.21	0.05	0.25	0.08	-0.08	-0.39	0.16
Priming	0.11	-0.05	-0.08	-0.05	0.16	0.09	0.46**
No Personal Benefit	-0.22	-0.01	0.03	0.11	-0.50**	-0.74***	-0.52**
Fairly treated	-0.23***	-0.07**	-0.06**	-0.01	-0.31***	-0.32***	-0.43***
Economic concern	0.29***	0.01	0.06**	0.15	0.16***	0.05	0.18***

^ *: Logistic regressions*

4.5. Enactment of Privacy Concerns

Our results tend to dissociate privacy concerns from information transmission, thus emphasizing the potential importance of indirect privacy concerns. This poses a problem with regard to the remediation

of privacy concerns: an non-situational preference for or against information transmission, expressed in the format of opt-in or opt-out is unlikely to resolve privacy anxieties. Table 4 gives all correlations among the dependent measures collected from the respondents to our study. The correlations suggest that there is a significant (and correctly signed) link between privacy concerns and the proposed ways to enact them.

It can also be seen that the requests for privacy expression (variables 5 to 7) are specifically motivated by privacy concerns (as opposed to other feelings such as fairness or fear of economic discrimination), more so than the acts of opting-in and opting-out themselves.

Table 4: Correlations Between Measurements (Across All Conditions)

	1	2	3	4	5	6	7	8	9	10
1 Happy with Deal	1									
2 Fairly treated	0.67	1								
3 Economic concern	-0.25	-0.3	1							
4 Privacy concern	**-0.3**	**-0.26**	**0.31**	1						
5 Requests opt-out option	(n.s.)	-0.1	(n.s.)	**0.23**	1					
6 Requests opt-in option	(n.s.)	-0.08	0.08	**0.37**	0.56	1				
7 Requests board decision	(n.s.)	-0.06	0.17	**0.27**	0.21	0.25	1			
8 Would opt-out	-0.35	-0.25	0.18	**0.36**	0.08	0.18	0.07	1		
9 Would not opt-in	-0.35	-0.22	0.09	**0.32**	0.07	0.14	(n.s.)	0.61	1	
10 Would vote against	-0.44	-0.33	0.21	**0.39**	0.15	0.18	0.1	0.53	0.46	1

The expressions of privacy concerns through opting out or not opting-in appear to be equivalent in our data (similar popularity, high correlation between both, see also Table 5). One specificity of our experiment is that not opting-in and opting-out were immediately juxtaposed decisions, and the ease of opting-in was emphasized. The usual empirical difference between the two measures (Bellman et al.

2001, Johnson and Goldstein, 2003) therefore seems to emerge from psychological costs arising from the departure from a default. We have no evidence to suggest that not opting-in and opting-out express different specific aspects of privacy concerns.

Table 5 offers an analysis of the imperfect enactment of privacy concerns through the various options offered to our respondents. There are concerned consumers who nevertheless opt-in, attracted by the cost-benefit equation associated with the offered insurance deal, and there are not-so-concerned consumers who opt-out or oppose the deal, presumably on the basis of fairness concerns that they dissociate from a privacy sentiment.

Table 5: Privacy Behavior as a Function of Privacy Concern (Across All Conditions)

	Propensity of concerned consumers ♦	*Propensity of not-so-concerned consumers* ♦	*Difference in propensity* ↔	*Diagnosticity* ♠
Request opt-out option♦	97%	85%	12%	43%
Request opt-in option♦	97%	61%	36%	46%
Request board decision♦	83%	71%	12%	46%
Would opt-out	63%	25%	38%	57%
Would not opt-in	57%	21%	36%	58%
Would vote against	73%	30%	43%	58%
n	263	244		

♦ *rated 6 or 7 on the 1-7 scale*
♦ *rated 1-4 on the 1-7 scale*
♦ *all differences are significant, and so are all differences between differences ($p<0.01$)*
♠ *proportion of concerned participants among those who exhibit this behavior*

If we wanted to infer privacy attitudes from privacy behaviors and demands, two findings stand out, based on Table 5. First, actions speak louder than words: the diagnosticity of privacy decisions is stronger than the diagnosticity of opinions about privacy tools (last column). Second, voting to ban the direct marketing practice in the association's board is most revealing of the depth of privacy concern (see before-last column of Table 5). Eleven percent of respondents would opt-in, would not opt-out but would like this kind of deal to be banned. Twelve percent of the respondents would opt-out or at least

not opt-in, while a the same time voting in favor of allowing the proposed deal. All these elements confirm, consistent with hypothesis 4, that there is a distinct contribution of social decision making with regard to privacy concerns, above and beyond decentralized opt-in and opt-out choices.

4.3. *General Discussion*

The indirect harm that applies across experimental conditions is the fear of price discrimination in the car insurance market. Factors such as support for the deal and feelings of fairness vary along these conditions, but the privacy concern remains constant whether the individual's personal information is involved or not. Privacy concern is heightened by the transmission of market-relevant data (hypothesis 2). The nature of desired intervention, in some cases, is consistent with a demand for protection against indirect harms beyond personal information dissemination. Taken together, the data suggests a concern about a privacy externality, which is not a function of data collection from individuals, but rather data use by those who collect any data.

The experimental design specifically targets this distinction between data collection and data use. We presented the participants with a tradeoff situation, rather than a generalized survey, then recorded their sentiments about privacy. Instead of attempting to directly measure the value of privacy, which would be entangled in valuations of other experiment-specific variables, we only focused on how respondents felt. Measuring the presence and relative strength of feelings across independent groups allowed us to capture feelings of utility while controlling for the anticipated benefit.

One could argue that, in showing that privacy concern did not vary much across conditions, we have only detected a constant, latent privacy sentiment. In such a case, distinctions may or may not exist, but the participants failed to discern the relative importance of different treatments. This could be because they did not understand the privacy issues at stake to begin with. Alternatively, the experimental treatment differences were too subtle. However, treatments were administered to independent groups,

and the use of other vehicles such as happiness and fairness metrics allow us to be fairly confident that we have measured valid responses for different treatment groups.

This design offers some external validity to the findings. The experiment does not use specific sets of rewards, and avoids specifying any explicit harm. Privacy sentiments are all relative, so that they can be scaled to other situations. A realistic situation was used to help prompt realistic responses, but nothing about the scenario offered implies that a similar set of incentives in a different context would produce different results.

5. *CONCLUSION*

The nascent field of the economics of privacy requires more empirical information about what consumers value and why. Against a null hypothesis of a consumer that was concerned about all aspects of personal information sharing or, alternatively, was focused exclusively on direct, explicit harms from privacy violations, we proposed a set of hypotheses arguing that consumers think about context and indirect effects. We demonstrated support for our hypotheses, and found that consumers are sensitive to context and indirect effects, rather than data collection itself. There are several implications of these results.

First, the privacy concern does not revolve around unitary "atoms" of personal data. This contradicts the assumptions of some models, which assume that units of personal information have intrinsic value. If, as we show above, privacy concerns are the same whether information is shared or not. This has broader implications for privacy regulation paradigms. If the flow of personal information is not the root of how people think about privacy, then policy solutions that rely on market mechanisms

(Rust, Kannan, and Peng, 2002) to efficiently control that flow will not function properly. Moreover, the idea of a privacy externality that introduces concerns based on dissemination of other people's information means that personal use of privacy-enhancing technologies will not eliminate the privacy concern. More broadly, information protection regimes should not treat all data as equal.

In fact, the above findings suggest that focusing on the data itself does not address the source of the privacy concern: data *use*. While we can draw no conclusions based on any specific mechanisms of society-wide control, we do find evidence that there is consumer demand for some social control, and that control should focus on data use. This emphasis seems more aligned with approaches like the OECD's Guidelines, which advocate the principles of purpose specification and data limitation (OECD, 1980). While such principles could be made manifest in the private market, many proposed mechanisms have fallen short (Greenstadt and Smith, 2005).

To understand and model privacy, more information is needed about consumer preferences, beyond "people want privacy." More sophisticated privacy models require evidence of a sophisticated, economically aware consumer. We have presented evidence from an experiment that people do behave somewhat rationally when considering realistic privacy situations. We find evidence of a sophisticated consumer that cares about economic context and indirect economic effects.

REFERENCES

Acquisti, Alessandro and Jens Grossklags (2005a), "Privacy and Rationality in Individual Decision-Making," IEEE Security and Privacy, 3(1), 26-33.

--- and --- (2005b), “Uncertainty, Ambiguity and Privacy,” Proceedings of the Fourth Workshop on the Economics of Information Security, Harvard University (June 2-3).

Bellman, Steven, Erik J. Johnson, and Gerald L. Lohse (2001), “On Site: to Opt-In or Opt-Out?: It Depends on the Question,” Communications of the ACM, 44 (February), 25-27.

Berendt, Bettina, Oliver Gunther, and Sarah Spiekermann (2005), “Privacy in E-Commerce: Stated Preferences vs. Actual Behavior,” Communications of the ACM, 48 (April), 101-106.

Chellappa, Ramnath. K. and Raymond G. Sin (2005), “Personalization versus Privacy: An Empirical Examination of the Online Consumer’s Dilemma,” Information Technology and Management, 6 (2-3), 181-202.

Cranor, Lorrie F., Mark S. Ackerman, and Joseph M. Reagle (2000), “Beyond Concern: Understanding Net Users' Attitudes About Online Privacy,” in The Internet Upheaval: Raising Questions, Seeking Answers in Communications Policy, Ingo Vogelsang and Benjamin M. Compaine, eds. Cambridge, MA: The MIT Press, 47-70.

Deighton, John A. (2000), “Frequency Programs in Service Industries,” in Handbook of Services Marketing and Management, Dawn Iacobucci and Teresa Swartz, eds. Thousand Oaks, CA: Sage, 401-407.

Greenstadt, Rachel and Michael D. Smith, “Protecting Personal Information: Obstacles and Directions,” Proceedings of the Fourth Workshop on the Economics of Information Security, Harvard University (June 2-3).

Hann, Il Horn, Kai-Lung Hui, Sang-Yong T. Lee and Ivan P.L. Png (2007), “Analyzing Online Information Privacy Concerns: An Information Processing Theory Approach,” Journal of Management Information Systems, 24 (Fall), 13-42.

---, ---, --- and --- (2008), “Consumer Privacy and Marketing Avoidance: A Static Model,” Management Science, 54 (June), 1094-1103.

Hermalin, Benjamin E. and Michael L. Katz (2006), “Privacy, Property Rights and Efficiency: The Economics of Privacy as Secrecy,” Quantitative Marketing and Economics 4, 209-239.

Huberman, Bernardo A., Eytan Adar, and Leslie R. Fine (2005), “Valuating Privacy,” IEEE Security & Privacy, 3 (September), 22-25.

Johnson, Eric J. and Daniel Goldstein (2003), “Do Defaults Save Lives?,” Science, 302 (21 November), 1338-1339.

Milne, George R. and Andrew J. Rohm (2000), “Consumer Privacy and Name Removal Across Direct Marketing Channels: Exploring Opt-in and Opt-out Alternatives,” Journal of Public Policy and Marketing, 19(2), 238-249.

Odlyzko, Andrew M. (2004), “Privacy, economics, and price discrimination on the Internet,” in Economics of Information Security, L. Jean Camp and Stephen Lewis, eds. Norwell, MA: Kluwer, 187-212.

OECD (1980), Guidelines on the Protection of Privacy and Transborder Flows of Personal Data.

Poindexter, J. C., Julie B. Earp and David L. Baumer (2006), “An Experimental Economics Approach Toward Quantifying Online Privacy Choices.” Information Systems Frontiers, 8(5), 363-374.

Rust, Roland T., P.K. Kannan, and Na Peng (2002), “The Customer Economics of Internet Privacy,” Journal of the Academy of Marketing Science, 30 (4), 451-460.

Shostak, Adam and Paul Syverson (2004), “What Price Privacy?” in Economics of Information Security, L. Jean Camp and Stephen Lewis, eds. Norwell, MA: Kluwer, 129-142.

Taylor, Curtis R. (2004), “Consumer Privacy and the Market for Customer Information,” Rand Journal of Economics, 35 (Winter), 631-651.

Wathieu, Luc and Kevin Morris (2004), “Meloche Monnex,” case study nr. 504 008, Harvard Business School, Boston.

--- (2006), “Marketing and Privacy Concerns,” Working Paper, Harvard Business School, Boston.

Westin, Allan P. (1998), “Privacy Concerns & Consumer Choice.” technical report, Louis Harris & Associates (December).

Recent ESMT Working Papers

	ESMT No.	Competence Center
An Empirical Approach to Understanding Privacy Concerns Luc Wathieu, ESMT Allan Friedman, John F. Kennedy School, Harvard University	09-001	Management and Technology
Cosmopolitanism, Assignment Duration, and Expatriate Adjustment: The Trade-Off between Well-Being and Performance Luc Wathieu, ESMT Amir Grinstein, Guilford Glazer School of Business and Management, Ben Gurion University of the Negev	08-011	Leadership
Trust and Creativity: Identifying the Role of Trust in Creativity-oriented Joint-developments Francis Bidault, ESMT Alessio Castello, Georgia Tech France	08-010	Management and Technology
Career Entrepreneurship Konstantin Korotov, ESMT Svetlana Khapova, ESMT Visiting Professor and Associate Professor at VU University Amsterdam Michael B. Arthur, Sawyer School of Management, Suffolk University	08-009	Leadership
Technology Commercialization Strategy in a Dynamic Context: Complementary Assets, Hybrid Contracts, and Experiential Learning Simon Wakeman, ESMT	08-008	Management and Technology
Organizational Redesign, Information Technologies and Workplace Productivity Benoit Dostie, HEC Rajshri Jayaraman, ESMT	08-007	Management and Technology
Resource and Revenue Management in Nonprofit Operations Francis de Véricourt, ESMT Miguel Sousa Lobo, Duke University	08-006	Management and Technology
Nurse-To-Patient Ratios in Hospital Staffing: A Queueing Perspective Francis de Véricourt, ESMT Otis B. Jennings, Duke University	08-005	Management and Technology
Critical Mass Michał Grajek, ESMT Tobias Kretschmer, Ludwig-Maximilians-Universität München	08-004	Management and Technology
The Rhythm of the Deal: Negotiation as a Dance Erik H. Schlie Mark A. Young, Rational Games, Inc.	08-003	Leadership
Legacy Effects in Radical Innovation: A Study of European Internet Banking Erik H. Schlie, ESMT Jaideep C. Prabhu, Tanaka Business School, Imperial College London Rajesh K. Chandy, Carlson School of Management, University of Minnesota	08-002	Management and Technology
Upsetting Events and Career Investments in the Russian Context Konstantin Korotov, ESMT Svetlana Khapova, ESMT Visiting Professor and Assistant Professor at VU University Amsterdam	08-001	Leadership
Ambiguity Aversion and the Power of Established Brands A. V. Muthukrishnan, Hong Kong University of Science and Technology Luc Wathieu, ESMT	07-005	Management and Technology

ESMT
European School of Management
and Technology GmbH

Schlossplatz 1
10178 Berlin
+49 (0) 30 212 31-0
+49 (0) 30 212 31-1279

www.esmt.org