

Caliendo, Marco; Clement, Michel; Papies, Dominik; Scheel-Kopeinig, Sabine

Working Paper

The cost impact of spam filters: measuring the effect of information system technologies in organizations

IZA Discussion Papers, No. 3755

Provided in Cooperation with:

IZA – Institute of Labor Economics

Suggested Citation: Caliendo, Marco; Clement, Michel; Papies, Dominik; Scheel-Kopeinig, Sabine (2008) : The cost impact of spam filters: measuring the effect of information system technologies in organizations, IZA Discussion Papers, No. 3755, Institute for the Study of Labor (IZA), Bonn, <https://nbn-resolving.de/urn:nbn:de:101:1-2008112645>

This Version is available at:

<https://hdl.handle.net/10419/35436>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

IZA DP No. 3755

The Cost Impact of Spam Filters: Measuring the Effect of Information System Technologies in Organizations

Marco Caliendo
Michel Clement
Dominik Papies
Sabine Scheel-Kopeinig

October 2008

The Cost Impact of Spam Filters: Measuring the Effect of Information System Technologies in Organizations

Marco Caliendo
IZA

Michel Clement
University of Hamburg

Dominik Papies
University of Hamburg

Sabine Scheel-Kopeinig
University of Cologne

Discussion Paper No. 3755
October 2008

IZA

P.O. Box 7240
53072 Bonn
Germany

Phone: +49-228-3894-0
Fax: +49-228-3894-180
E-mail: iza@iza.org

Any opinions expressed here are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but the institute itself takes no institutional policy positions.

The Institute for the Study of Labor (IZA) in Bonn is a local and virtual international research center and a place of communication between science, politics and business. IZA is an independent nonprofit organization supported by Deutsche Post World Net. The center is associated with the University of Bonn and offers a stimulating research environment through its international network, workshops and conferences, data service, project support, research visits and doctoral program. IZA engages in (i) original and internationally competitive research in all fields of labor economics, (ii) development of policy concepts, and (iii) dissemination of research results and concepts to the interested public.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ABSTRACT

The Cost Impact of Spam Filters: Measuring the Effect of Information System Technologies in Organizations*

More than 70% of global e-mail traffic consists of unsolicited and commercial direct marketing, also known as spam. Dealing with spam incurs high costs for organizations, prompting efforts to try to reduce spam-related costs by installing spam filters. Using modern econometric methods to reduce the selection bias of installing a spam filter, we deploy a unique data setting implemented at a German university to measure the costs associated with spam and the costs savings of spam filters. The applied methodological framework can easily be transferred to estimate the effect of other IS technologies (e.g., SAP) implemented in organizations. Our findings indicate that central IT costs are of little relevance since the majority of spam costs stem from employees who spend working time identifying and deleting spam. The working time losses caused by spam are approximately 1,200 minutes per employee per year; these costs could be reduced by roughly 35% through the installation of a spam filter mechanism. The individual efficiency of a spam filter installation depends on the amount of spam that is received and on the level of knowledge about spam.

JEL Classification: M12, M15

Keywords: spam, spam filter, treatment effects, propensity score matching

Corresponding author:

Michel Clement
University of Hamburg
Institute for Marketing and Media
Von-Melle-Park 5
D-20146 Hamburg
Germany
E-mail: michel@michelclement.com

* The authors would like to thank Felix Eggers, Mario Farsky, Karen Gedenk, Kissan Joseph, and Franziska Völckner for their valuable comments on previous versions of this paper.

Introduction

Spam¹ is defined as the use of electronic communication channels to send unsolicited bulk messages with commercial content indiscriminately to recipients (OECD 2005). However, whether an e-mail is perceived as spam depends on the preferences of the user; some believe that the opportunity offered is a good deal and will purchase the promoted product. Thus, spamming continues to be a profitable business. Even with low purchasing probabilities, spamming can be economically viable because the variable operating costs of using online communication channels are close to zero. Furthermore, already low market entry costs continue to decline because of strong price competition among e-mail list providers, which offer millions of validated e-mail addresses for prices less than US\$50 (Sipior, Ward, and Bonner 2004; Hann et al. 2006).

Spammers rely on various business models, either selling their (validated) e-mail address lists to other spammers, or directly promoting their own or third-party products. Mostly, they advertise third-party products using a revenue share model. The profit depends on the product, timing of the campaign, opening rate, and purchase probability, which is influenced by the quality of the spam mail. Potentially high profits and low market entry barriers continuously attract spammers, even though legal actions against spam have been initiated by legislation (Zhang 2005). Since October 2003, more than 50% of the total global mail traffic has been classified as spam, and in January 2008, the quota was greater than 75% (Messagelabs 2008).

Spamming is accompanied by negative externalities, with the largest share of the costs associated with sending unsolicited bulk messages being borne by e-mail service providers and those

¹ The association between unsolicited bulk e-mail and the word “spam” derives from a comedy sketch by the British comedy group Monty Python, in which the name of Hormel Foods’ meat product “spam” gets mentioned about a hundred times within just a few minutes.

who receive spam (Melville et al. 2006; Goodman, Cormack, and Heckerman 2007). As a result, large e-mail providers must handle billions of “abusive mails” per day (MAAWG 2007). Organizations complain about the costs of spam, which reduces their employees’ productivity by forcing employees to allocate their limited attention resources to the messages (Falkinger 2007). Thus, the spamming phenomenon affects individuals (in organizations) as well as economies on a global scale; in response, technical, market, and legal actions seek to reduce the costs of spam (OECD 2005; Joseph and Thevaranjan 2008).

Despite the strong interest in reducing the costs of spam, we find only limited academic research that addresses the effectiveness of anti-spam actions on costs (Joseph and Thevaranjan 2008; Pavlov, Melville, and Plice 2008). A variety of management-related studies conducted by consulting companies (e.g., Vircom 2004) measure a few indices on a corporate level and then generalize these costs or break them down to an individual level. However, precise measures of individual or corporate costs (e.g., from data centers) that assess the magnitude of the problem have not been subject to academic consideration. Most spam-related research focuses instead on mechanisms to reduce spam or its impact. This type of literature can be grouped in four streams: First, theoretical papers by economists analyze and model market mechanisms to overcome the externalities of spam by increasing the costs for the sender (e.g., Kraut et al. 2005). Second, significant research focuses on legal issues to increase the risk (and cost) for spammers (e.g., Zhang 2005). Third, another stream of literature addresses user perceptions of spam (e.g., Morimoto and Chang 2006). Fourth, literature found in the field of information systems and computer science contributes ways to enhance filtering technologies by identifying, marking, and filtering unsolicited e-mails (e.g., Cormack and Lynam 2007). The effect of market mechanisms and legal actions are long term and typically beyond the control of IS managers, whereas most individuals and organizations use filter technologies to reduce the amount of spam in inboxes. However, de-

spite the widespread use of spam filters, no study—to our knowledge—empirically addresses the question whether or not spam filters or other countermeasures really reduce costs.

This dearth of research is especially surprising for two reasons: First, with regard to the corporate level, company investments in information technology involve high risk and unclear outcomes (Dewan, Shi, and Gurbaxani 2007; Dewan and Ren 2007). Second, although spam filters are designed to reduce the spam burden and thus spam-related costs, this desired effect does not necessarily occur; the net benefit of the adoption and usage of spam filters depends on the costs of installation and the sum of time losses due to updating and training the filter, as well as for checking the filter results regarding potential misclassifications (false positives/negatives). Thus, although spam filters help users identify spam, they also lead to new costs due to two reasons: (1) filter technology may not be sufficient enough to justify the recurring costs after installation (e.g., updates and training of the filter), and (2) substantial misclassifications of relevant e-mails occur because filters are not sufficiently trained by individuals. Therefore, the individual net cost effect remains an open issue, and it remains unclear whether spam filters do indeed reduce costs.

We address this research gap and pursue two major research aims. First, we measure the central and individual costs of spam in an organization to evaluate the magnitude of the spam problem. Second, we take a first step towards evaluating the efficiency of countermeasures by focusing on spam filters. Without a controlled experimental setting, evaluating the cost effects of spam filters in cross-sectional samples requires rigorous control for the presence of a selection bias, which could arise because respondents using a spam filter might have different characteristics than those who do not. For example: In our data we find that users of spam filters have *higher* spam costs than those who do not use a filter. This raises the question of causality: Do users install filters because of high spam costs, or do filters cause rather than reduce costs as noted above? Ignoring this selection effect would lead to biased results pertaining to the effect of spam

filters with a typical cross-sectional sample. Therefore, we reduce this selection bias by drawing on a propensity score-matching procedure (Imbens 2004; Smith and Todd 2005).

To measure the effect of spam filters on the individual costs of spam, we collected data from a German university and differentiate between the individual costs of the employees and the costs to the university's central data center. Our data set comprises information regarding 1,000 employees.

Our research contributes to IS literature in two dimensions: First, the nature of our data set (which includes individual and organizational costs) enables us to provide an indication of the costs of spam on both individual and corporate levels, as well as to show the impact of spam filters on individual costs. Second, because of the likely presence of a sample selection bias in many IS research settings focusing on the impact of implemented IT interventions in organizations (e.g., the introduction of SAP), we demonstrate the application of propensity score matching, and provide guidelines and implications for its extension to other research questions. Thus, our methodological framework can be applied to many other settings that focus on estimating the success of IS interventions when experiments are not feasible.

The remainder of this article is organized as follows: In the next section, we present a brief overview of related research. We then introduce the method utilized to quantify the cost-saving effects of spam filters by correcting for selection bias. Section 4 summarizes the structure of the field experiment we conducted to estimate the costs of spam, and presents some descriptive statistics. We discuss our estimation results in section 5 and conclude with implications for theory and practice.

Related Research

Contrary to the public presence of the spam phenomenon, academic research has devoted minimal attention to researching spam-related costs. Previous publications include reports by the OECD (2005) and the European Union (2004). The latter categorizes costs incurred by spam into direct and indirect costs, distinguishing between five different cost components. *Direct costs* include (1) losses of working time and productivity caused by the need to delete spam and install and train filters, and (2) central costs that accrue in data centers and IT departments as a result of the installation of countermeasures. Furthermore, (3) direct costs may arise if Internet service providers (ISP) must adapt their capacities to respond to increased spam. *Indirect costs* refer to the effect of spam on e-mail usage: (4) e-mails can be erroneously identified as spam (false positives) or (5) might contain viruses or other potentially harmful features (European Union 2004). Several of these cost components are directly caused by the decision to adopt a spam filter: the installation and training of filter mechanisms, the central costs, and the control of the spam filter especially for false positives. Only if these costs do not exceed the cost savings achieved by the filter will a spam filter mechanism have the desired effect.

Despite these numerous consequences, academic research invests little effort into quantifying spam-related costs, especially with regard to possible cost-saving (or cost-causing) effects associated with the widespread use of spam filters. Existing research primarily addresses ways to reduce the amount of spam and can be grouped into four streams. The first major stream addresses a key characteristic of electronic communication, namely, the low marginal costs of e-mail distribution. This characteristic represents the primary reason for the existence of spam, because in the offline world, the sender-pays system prevents advertisers from engaging in heavy spamming. In the online world however, the sender does not bear significant costs for excessive mailings; in-

stead, those costs are externalized to ISPs and recipients, which characterizes e-mailing as a digital commons (Melville et al. 2006). Economists attempt to reduce the potential for externalizing the costs associated with e-mailing by confronting spammers with greater e-mailing costs. This would make spam less attractive, thus changing the economics of e-mailing. For example, one approach involves e-mail postage (Kraut et al. 2005), and another proposes a bonded sender program that requires senders to deposit a certain amount of money if not listed on a white list. If the recipient declares the e-mail to be spam, the deposit gets retained (Joseph and Thevaranjan 2008). Despite their theoretical efficiency, these and other comparable approaches could not yet be implemented. The considerable coordination effort associated with these measures makes it doubtful whether this gap will change in the near future.

The second cluster of literature deals with legal measures against spam. However, these measures have the inherent limitation that, in many cases, spammers set the pace for technical developments (Melville et al. 2006). Furthermore, legal measures are sustainable only if they are coordinated on a global basis. Since such coordination is rare, it cannot be expected that legal countermeasures will have a sustainable impact on costs (Zhang 2005).

A strong basis in the third literature stream suggests the implementation of technical measures against spam, whether implemented centrally (e.g., IT department or ISP) or decentralized (on each user's computer). Research focuses on technical issues such as blocking IP numbers, filtering e-mails, or authentication mechanisms (e.g., Sahami et al. 1998; Park and Deshpande 2006; Duan, Dong, and Gopalan 2007; Cormack and Lynam 2007).

A nascent, fourth stream of literature explores user perceptions of the growing spam burden by measuring the attitude or inconvenience costs of spam (Yoo, Shin, and Kwak 2006; Morimoto and Chang 2006). However, none of these publications addresses the effect of technical measures

on the development of spam-related costs. Our research relates to the last two groups of publications, in that we focus on both the costs and user perceptions of filter mechanisms.

Method

Our research goal is to shed light on the unresolved issue of whether the installation of a spam filter reduces working time losses experienced by employees. A suitable framework to address such a question is the potential outcome approach also known as the Roy-Rubin model (Roy 1951; Rubin 1974). To measure the *individual causal effect* of a spam filter we seek to compare working times – with and without a spam filter – for the same user at the same time. Let the potential outcomes – here working time losses in minutes – be defined as $Y_i(D_i)$ for each individual i and let D_i denote the treatment indicator – here: installation of a spam filter. The *individual causal effect* is simply the difference between both potential outcomes, hence: $Y_i(1) - Y_i(0)$.² Unfortunately, one of these potential outcomes is unobservable or counterfactual, so instead, we consider comparing the mean working time losses of employees before and after they install the spam filter. However, relying solely on this approach would also be problematic, because employees could change their behavior, e.g., by making an initial online purchase or subscribing to a newsletter. Furthermore, external circumstances might change as well, such as an overall increase in Internet usage due to new and faster connections. Therefore, solely comparing a situation today (installed spam filter) with a situation in the past (no spam filter) could be very misleading as a result of unobserved effects over time.

Another approach is to compare the mean working times of those employees who have installed a spam filter (i.e., the “treatment group” hereafter) and those who have not (“control

² See e.g. Heckman, LaLonde, Smith (1999) or Imbens (2004) for a detailed discussion of the evaluation framework.

group”). In an experimental setting, i.e., where the installation of spam filters is randomly assigned, this would be a feasible strategy (Harrison and List 2004). However, such an experimental setting is not an option in the case of spam filters or comparable technological investments for most companies. Further, our empirical data (see Tables 3-5 in the following Section) indicates that users with a spam filter differ in more aspects than just the decision to install a filter, such that simply comparing mean working times of the treatment and the control group would yield a biased estimate of the *average treatment effect of the treated (ATT)* denoted by τ_{ATT} :

$$E[Y(1) | D = 1] - E[Y(0) | D = 1] = \tau_{ATT} + E[Y(0) | D = 1] - E[Y(0) | D = 0]. \quad (1)$$

The difference between the left-hand expression and the τ_{ATT} can be called the “*self-selection bias*”. It is reasonable to assume that users with high spam-related costs have a higher propensity to install spam filters. For example, those who use their e-mail frequently probably differ from those who use it irregularly. To make meaningful comparisons and estimate the causal effects of the spam filter, we must find a proper substitute for the unobservable component $E[Y(0) | D = 0]$.

To address the self-selection problem, we assume that potential working time losses are independent of the installation decision, given a set of relevant, observable variables X (*Conditional Independence Assumption*).³ By referring to *relevant* variables, we mean that they simultaneously influence the decision to install a spam filter and the outcome variable. So selection must be solely due to observable variables, which is generally a strong assumption. Even though we are confident that our data covers the crucial X variables to justify this assumption, we subsequently test the sensitivity of our results to this assumption. With a large set of variables X and a small sample

³ See Imbens (2004) for an overview of different nonparametric estimation approaches to average treatment effects under exogeneity.

of users (as is the case in our data), it is difficult to find users from the control group who have exactly the same variable values as each user in the treatment group. Therefore, we follow Rosenbaum and Rubin (1983) who have shown that it is sufficient to use the propensity score $P(X)$ – here: the probability of installing a spam filter – instead of the whole set of observed characteristics X in order to balance the distribution of covariates between both groups.

The basic idea behind propensity score matching (PSM) is to approximate the counterfactual working times of individuals in the treatment group by finding similar users in terms of their propensity score values in the control group. Formally, the PSM estimator can be written as:

$$\tau_{ATT}^{PSM} = E_{P(X)|D=1}[Y(1) | D = 1, P(X)] - E[Y(0) | D = 0, P(X)]. \quad (2)$$

It is simply the mean difference in outcomes over the common support, appropriately weighted by the propensity score of the treatment group individuals. Restricting the comparison to individuals who fall inside the region of common support ensures that only comparable individuals, whose propensity score values overlap, are used to estimate the treatment effect. To ensure that users from the control group have a positive probability of belonging to the treatment group, we further assume that $P(D=1|X) < 1$ (*Overlap Assumption*).

Based on these two assumptions we can use PSM to estimate the average treatment effect of the treated (Heckman et al. 1998). The most straightforward matching estimator is nearest neighbor (NN) matching. For each user from the treatment group, we choose one user from the control group who is closest in terms of the propensity score. In addition, NN matching can be done with or without replacement. In the former case, a user from the control group can appear more than once as a matching partner; in the latter case a user is considered only once. Whereas NN matching algorithms use only a few observations from the control group to construct the counterfactual outcome for each treated individual, Kernel matching (KM) is a nonparametric matching estima-

tor that uses weighted averages of (nearly) all individuals from the control group. Because KM estimators use more of the available information to estimate causal effects, lower variances are achieved. However, this method also employs users from the control group, though they are potentially poor matches reflecting the trade-off between bias and efficiency. It should be clear that the imposition of the common support condition is very important for KM. Different criteria for imposing common support are available⁴; we use the “MinMax” criterion, according to which users from the treatment group whose PS is higher (smaller) than the maximum (minimum) PS in the control group are dropped from the analysis. This highlights the great advantage of PSM compared with ordinary least squares regression (OLS). With PSM and the common support condition, the treatment effects are estimated only within the common support, so individuals for individuals without comparable matches are excluded. The estimated treatment effect must then be interpreted over the region of common support.

Below we briefly describe the data that was gathered to determine the cost-saving effects of the individual decision to install a spam filter.

Research Design and Data

We conduct our research at a German university with approximately 8,000 employees (including the university’s hospitals), whose size resembles that of a medium-sized company.⁵ A unique advantage of this setting results from the integrated structure of the institution, which combines most sources for spam-related costs into one organization: the university serves as an e-mail provider, operates its own data center, and employs a sufficiently high number of computer users. Our data collection can thus be restricted to two points of measurement within the universi-

⁴ See Smith and Todd (2005) and Lechner (2002) for an overview.

⁵ Universities previously have served as research settings for spam research (e.g., Melville et al. 2006).

ty: (1) We measure all central costs incurred by the university data center, where all computer-related tasks are centralized, and (2) we contact 5,000 university employees (excluding the university's hospitals) to measure the individual costs using an online questionnaire that also contains a set of covariates.

Central Costs

We collect information about the expenses incurred by the provider through interviews with IT experts from the university data center. The data center had reacted to its increasing spam burden by setting up an infrastructure based on free open source software (SpamAssassin) that checks all incoming e-mails for spam properties before labeling them in accordance with their spam probability and forwarding them to recipients. On average, the data center processes 170,000 e-mails per day from outside the university; in 2005 and 2006, the spam quota was about 90%.

This infrastructure creates expenses in four categories. Prior to the installation, organizational and administrative tasks had to be performed to obtain clearance from all relevant organizational units. The expenses for hard- and software were accompanied by labor costs for installing the system and training of the staff. In addition, the initiative against spam generates recurring costs, because the data center regularly provides support to the university's e-mail users. Furthermore, all anti-spam measures must be controlled for efficiency and are subject to constant further development. These activities together create expenses on a regular basis. For the purpose of comparability, we aggregate these regular costs for a period of one year. Therefore, the sum of all costs (Table 1) equals Euro 15,120 for the first year after and including installation.

>> Insert Table 1 about here <<

Individual Measures

Our second measure pertains to the individual costs to the employees of the university, whom we contacted by e-mail in the winter season 2004/2005 to invite them to fill out an online questionnaire. We contacted 5,000 employees and received exactly 1,000 responses (20%). We note that 52.7% of the 1,000 respondents already had adopted a spam filter.

Adoption research (e.g., theory of reasoned action, technology acceptance model) suggests that the perceived utility of an innovation drives adoption behavior—in this case, spam filter installation. We adopt this view and assume that the propensity to install a spam filter is strongly driven by the amount of spam an individual receives. We use this measure as a proxy for the perceived utility of a countermeasure. Furthermore, we measure demographic variables and the individual's own spam-prevention mechanisms. We control for psychographic variables related to the user's reactance to spam and measure the level of distribution of his or her e-mail address to others. We rely on these variables to analyze each individual's decision to install a spam filter.

Costs. On the basis of interviews we conducted prior to launching the questionnaire, we separate recurring from nonrecurring costs. *Recurring* costs refer to the daily time involved in deleting spam and controlling spam filter results for false positives. For the 1,000 respondents, these activities take up an average of 4.87 minutes per day. The *nonrecurring* costs include inquiries about the spam filter and installation by the respondent or an assistant. We aggregate all time expenditures into an index that covers the costs for one year, assuming that the average employee in Germany works approximately 250 days. Table 2 reports the respective time expenditures, indicating an average time loss of more than 1,200 minutes per year. Direct monetary costs on the individual level are not included, since the installation of the university spam filter did not require any payments by the respondents.

Table 2 also shows significantly greater time losses (1,597 minutes per year) for the 527 participants that installed a spam filter ($D = 1$) compared with the loss of 858 minutes for the 473 participants without one ($D = 0$). This finding suggests that either a spam filter causes working time losses, or that users with high costs have a stronger inclination to install a filter. The latter explanation would suggest a sample selection bias. Thus, to measure the potential cost savings of spam filters for the 1,000 users, we must reduce the selection bias using appropriate methods.

>> Insert Table 2 about here <<

Individual Factors. We assume that the most important driver of the adoption decision is the amount of spam that each person is confronted with. We further assume heterogeneity exists in individuals' approaches for dealing with spam. We therefore include variables that capture actions related to spam prevention, such as whether the e-mail address gets used strategically to prevent spam. Furthermore, we ask how respondents identify an e-mail as spam (Table 3). The descriptive analysis shows significant differences between users with and without spam filters in several dimensions. For example, significant differences in age and gender emerge between the two groups; we also identify strong differences in the quantity of spam mails received, such that those with a spam filter receive more than three times as much spam as those without. Spam filter users have a higher propensity to use alternative e-mail addresses and request to be removed from mailing lists more often. Finally, spam filter users rely more on inspecting the subject line before they open e-mails.

>> Insert Table 3 about here <<

Reactance. Spam not only reduces productivity in the workplace, but can also be perceived as intrusive. Prior research associates unsolicited mails with a perceived loss of control that can lead to the psychological effect of reactance (Morimoto and Chang 2006). If spam is perceived to

be intrusive, thus leading to a sense of loss of control, the recipient might be inclined to take actions to restore his or her original state, which in this case would be a spam-free environment. Therefore, we control for behavioral changes caused by reactance; we also measure individual spam sensitivity to control for possible influences and account for the perceived degree of loss of control (Table 4). This measure therefore includes properties that, according to the individual, constitute the bothersome factors of spam.

The descriptive results indicate that particularly users with spam filters feel that they have wasted significantly more time with, and lost confidence in, the e-mail medium, which then increases reactance because they feel bothered by spam. We also find that filter users are significantly more sensitive to what they perceive as spam (e.g., fun mail, ads from business partners, large attachments). Furthermore, we detect significant differences in the perceptions of spam properties, such that users without spam filters note significantly higher fears about the potential hazards of spam.

>> Insert Table 4 about here <<

Distribution. The last group of variables controls for the usage habits associated with the e-mail address. We capture the degree to which the e-mail address has been distributed to known or unknown contacts, with the assumption that a “generous” distribution of the e-mail address leads to a considerably higher spam load, which in turn increases the probability of spam filter installation (Table 5). The descriptive measures indicate no differences in the distribution of the e-mail address to known contacts, but we find significantly higher distribution levels to unknown contacts by those who have installed spam filters.

>> Insert Table 5 about here <<

Cost Analysis

Group Characteristics before Matching

The average working time loss for users with spam filters is 1,597 minutes; for those without it is 858 minutes. A simple mean comparison suggests that a spam filter increases working time losses by 739 minutes. However, the groups differ in their observed characteristics, as previously stated. Users who have installed a spam filter appear to be significantly older and more likely to be men. They receive more e-mails (solicited and unsolicited) and exhibit a higher information level about spam (see Table 3). As Table 4 reveals, they also have a greater perception that the amount of spam mails, wasted time to control spam mails, and mail handling time have risen in general. Furthermore, they are less inclined to believe that spam mails can damage their personal computers. Finally, these users publish their e-mail addresses on Web sites, in online directories, or in online forums more frequently (see Table 5). The large differences in observed characteristics indicate that a simple mean comparison of working time losses between both groups cannot yield an unbiased estimate that answers the question of how much a spam filter can reduce working time losses.⁶

Propensity Score Estimation⁷

The first step of PSM is to estimate the propensity score, which we do using a binary logit model.⁸ Our dependent variable equals 1 for users that installed a spam filter and 0 otherwise.

⁶ Although it might seem to be a natural solution to control for other factors contributing to the installation decision by running a regression analysis to determine the cost effects of spam filters, this will not solve the issue: standard regression analysis does not implement a common support condition, so users with diverse characteristics are compared, and estimates are extrapolated even to regions in which common support (and number of observations) is low. Because of this disadvantage, we implement propensity score matching.

⁷ See Caliendo and Kopeinig (2008) for practical guidance on how to implement propensity score matching.

Table 6 presents the results of the propensity score estimation, including the previously derived explanatory variables.⁹ We drop those respondents for whom we lack information about the key information spam properties from the analysis, reducing the number of observations to 520 users who have installed a spam filter and 440 users have not.

>> Insert Table 6 about here <<

With this logit specification, we achieve a hit rate of 73.1%.¹⁰ However, the aim of the propensity score estimation is not to maximize the hit rate but rather to balance the covariates between both groups, which we subsequently test by calculating standardized biases. Age, gender¹¹, and the number of e-mails and spam mails received all significantly affect the decision to install a spam filter. Furthermore, factors that we relate to reactance, such as the increase in time expenditures for handling e-mails, the perceived increase of received spam mails, or the perception that spam mails are harmful to personal computers, increase the probability of installing a spam filter. The level of information about spam also positively affects installation decision.

The propensity score distribution obtained from the logit estimation is depicted in Figure 1, which indicates that the PS distribution differs considerably between the treatment and the control group. Hence, NN matching algorithms without replacement would create poor matches due to the high-score users from the treatment group, which likely get matched to low-score users from the control group. The PS interval of treated (untreated) users is [0.033 – 0.999] ([0.048 –

⁸ In the case of a binary treatment, the estimation with either a logit or a probit model should yield similar results. We also estimate the propensity score using a probit model and obtain similar values.

⁹ We also test for multicollinearity. All variables have variance inflation factor values < 10.

¹⁰ Hit rates are computed as follows: If the estimated propensity score is greater than the sample proportion of users that have installed a spam filter (i.e. $\hat{p}(X) > \bar{P}$), observations are classified as 1. If $\hat{p}(X) \leq \bar{P}$, observations are classified as 0.

¹¹ Note that the sign for gender is now negative (as opposed to the bivariate analysis, Table 3), indicating that male users have a lower propensity to install a filter if a multivariate analysis is deployed that controls for relevant factors (Table 6).

0.931]). Hence, the common support (based on the MinMax criterion) lies between 0.048 and 0.931; consequently, 74 treated users (treated off support) had to be dropped from our analysis.

>> Insert Figure 1 about here <<

Matching Results

We present three different matching estimators in our analysis: single NN matching (matching estimator A hereafter), single NN matching with common support condition (matching estimator B), and KM with common support condition (matching estimator C). For the KM estimator we use an Epanechnikov kernel function with a bandwidth parameter, according to Silverman (1986), of 0.06.¹² In Table 7, we present our matching results. All estimated effects are negative, which indicates that the installation of a spam filter lowers average working time losses, regardless of the algorithm chosen. However, the absolute effects differ between the matching algorithms.

Matching estimator A (NN matching) does not impose the common support condition, resulting in an effect of -814.48 minutes. That is, the effect of installing a spam filter reduces working time losses by roughly 800 minutes. However, this result must be treated with caution for two reasons: First, no individuals are dropped from the analysis, so that even treated individuals that cannot be properly compared with untreated users are used to measure the effect, and some control individuals appear repeatedly; for example, one member of the control group gets used 72 times. Second, any interpretation of these results should be preceded by an evaluation of matching quality. To determine whether the matching procedure balances the distribution of covariates

¹² We have tested several different bandwidths and distributions that all yield similar results. Detailed estimation results for all matching algorithms are available on request from the authors.

between both groups, Rosenbaum and Rubin (1985) propose using standardized biases (SB).

Standardized biases before and after matching are defined as follows:

$$Bias_{bef} = \frac{(\bar{X}_1 - \bar{X}_0)}{\sqrt{0.5 \cdot (V_1(X) + V_0(X))}}; Bias_{aft} = \frac{(\bar{X}_{1M} - \bar{X}_{0M})}{\sqrt{0.5 \cdot (V_{1M}(X) + V_{0M}(X))}}, \quad (4)$$

where $\bar{X}_{1|01}(V_{1|01}(X))$ is the mean (variance) in the treatment (control) group before matching, and $\bar{X}_{1|01M}(V_{1|01M}(X))$ the corresponding values after matching. The SB after matching for estimator A is highest at 15.64%. Even though this level represents a reduction compared with the situation before matching (20.12%), it is clearly not sufficient. In general, it is suggested that standardized differences should be below 5% (Sianesi 2004; Caliendo and Kopeinig 2008). Therefore, matching estimator A is not satisfactory, and we turn to the next two approaches.

For estimators B and C, we impose common support conditions and drop 74 users from the treatment group (*offsup*). These estimators balance the covariate differences between the groups, and, through the matching procedure, more than 60% of the covariate differences are removed. In addition, Sianesi (2004) suggests re-estimating propensity scores on the matched sample, and comparing the pseudo- R^2 values before and after matching. The pseudo- R^2 after matching should be lower, because systematic differences in the distribution of the covariates between groups should have been removed by matching. In our analysis, we achieve pseudo- R^2 values of 0.247 before and 0.062 (estimator B) and 0.024 (estimator C) after matching.

>> Insert Table 7 about here <<

Table 7 shows that the use of matching estimator B does not balance the covariate distribution satisfactorily (bias aft > 5%), whereas using estimator C does. Consequently, we rely on estimator C as the appropriate measure, and find that the causal effect of a spam filter installation

on working time losses equals approximately -439.52 minutes and is significant.¹³ Therefore, in our sample, the installation of a spam filter is beneficial and decreases working time losses by more than 400 minutes per year. Although savings of approximately seven hours per year might not sound too impressive, it becomes more so if viewed within the organizational context. Consider, for example, an average wage of 30 Euro per hour, and assume that 1,000 employees work for a company; the seven hours saved accumulate to a considerable sum that clearly exceeds the central costs associated with installing a spam filter mechanism.

Effect Heterogeneity

As we noted above, we observe considerable heterogeneity with regard to variables that characterize the usage intensity of e-mail communication, and we find that the decision to install a filter is strongly influenced by the intensity of e-mail communication. This notion implies that a spam filter might not be a necessary and efficient option for all users. Hence, we conducted group-specific matching procedures in order to uncover underlying factors that account for heterogeneity in the magnitude of the treatment effect. The group-specific results in Table 8 show that the desired cost saving effects of a spam filter installation do not occur in any case, rather that the size and the direction of the effect depends on user characteristics.

First, the number of spam mails received by an individual plays a central role in the cost effect of a spam filter. We see that a spam filter only saves costs for those users who are bothered by a large spam burden. For users who receive less than 10 spam mails per day, the cost effect of a spam filter is even positive. This implies that a spam filter does not save, but rather induces costs for users who only receive few spam mails; in these cases a manual identification and eli-

¹³ To draw inferences about the significance of the effect, we report bootstrapped standard errors (*s.e.*) with 200 replications. Heckman et al. (1998) show that bootstrapping is valid for drawing inferences for kernel matching methods.

mination will probably be more efficient. A likely reason is that the costs associated with the installation, training, and control of the filter exceed the beneficial effect of saving time through classification of e-mails.

>> Insert Table 8 about here <<

Second, since the efficiency of spam filter training or manual handling of spam mails is likely to depend on the individual's know-how, we include the level of proficiency in dealing with spam in the group-specific analysis. If the cost effects of a spam filter are analyzed conditionally on how well informed the user is, we observe that significant cost-saving effects only occur when the user is not well informed concerning spam. If a user does not have a profound knowledge about spam, a manual inspection appears to be less efficient than an automatic classification; in this case, the filter uses information that can only be readily substituted by visual inspection as is the case of experienced users. For these well-informed users, the cost-saving effects are present but fail to be significant, indicating that an experienced user might as well rely on his or her proficiency to manually classify e-mails.

Sensitivity to Unobserved Heterogeneity

The validity of our estimates depends on the conditional independence assumption. For this assumption to be fulfilled, we must observe all variables that simultaneously influence the propensity to install a spam filter, and the outcome variable. Because of this very strong assumption, we validate whether unobserved heterogeneity might alter our results by applying the bounding approach proposed by Rosenbaum (2002). The basic idea of this approach is to determine how strong an unobserved variable must be to influence the decision to install a spam filter and to change our matching results. Becker and Caliendo (2007), and DiPrete and Gangl (2004) provide

guidance for implementing this bounding approach in the case of a discrete or metric outcome variable. As a starting point, we assume that the propensity score is influenced not only by observed variables X , but also by unobserved variables U , such that $P(D = 1 | \beta X + \gamma U)$. If the selection is based solely on observable variables X , the study is free of hidden bias, γ will be 0, and the installation probability will be determined solely by X . However, if a hidden bias exists, two individuals with the same observed covariates X will have differing chances of installing a spam filter.

>> Insert Table 9 about here <<

By varying the influence of γ , we can examine the sensitivity of our results to two different scenarios. First, we consider a situation in which we underestimate (t-hat-) the true treatment effect; second, we address a situation in which we overestimate (t-hat+) the true treatment effect. For both scenarios, we re-estimate the test statistics (see Table 9) and check the significance of the coefficients. Given the negative estimated treatment effect, the bounds that emerge under the assumption that we have overestimated the true treatment effect are of less importance. The effect is significant at $\gamma = 1$ and becomes even more significant for increasing values of γ if we have overestimated the true treatment effect. However, the bounds under the assumption that we have underestimated the treatment effect reveal that even high levels of γ would not alter the significance of the results. To be more specific, at a value of $\gamma = 1.8$, the result remains significant at the 5% level; at $\gamma = 1.9$, it would be still significant at the 10% level. Only at a γ -value of 2.0 do the results become insignificant. However, $\gamma = 2$ implies that the unobserved component in $P(D = 1 | \beta X + \gamma U)$ would have to be as strong as the observed component. Given the informative data at hand, this is rather unlikely; therefore, we can state that only very high levels of unobserved heterogeneity would alter our results.

Conclusion and Limitations

Our analysis shows that the existence of spam confronts organizations with significant expenses, primarily in the form of working time losses. Every year, employees waste an average of 1,200 minutes—or two working days—dealing with spam.

When an organization decides to react by setting up a spam filter mechanism, it incurs further expenses, and the cost-saving effects have been unclear thus far. We clarify this situation by showing that spam filters can reduce individual spam-related costs. The effect is strong; cost savings accumulate to 439 minutes per person per year, and our findings are significant and insensitive to unobserved heterogeneity. The magnitude of the different cost components also suggests that the primary concern in organizations should be the effectiveness of filter mechanisms on the individual level rather than central costs caused by spam. Due to the fact that cost-saving effects only occur for those users with an excessive spam load, those with little knowledge about spam, or those lacking adequate countermeasures, companies should primarily address these users in order to reduce costs through spam filters. For these users the installation of a spam filter will lead to the desired effect. If a user is well informed or is not strongly affected by spam, a company should not encourage the implementation of technical countermeasures. In this case, manual inspections appear to be more efficient than filter mechanisms that tend to increase overall spam costs.

We derive our conclusions by applying an econometric matching approach that controls for selection bias. It is unlikely that the selection bias is unique to our sample; rather, a selection bias probably poses a problem for a multitude of other research questions in IS. Within organizations, this effect might arise in evaluations of the effectiveness of optional IT innovations, for which an experimental setting is not available. Consider, for example, a mobile e-mail solution offered to

employees. To evaluate its effectiveness, the company cannot use a simple cross-sectional approach because the estimation cannot distinguish whether the outcome measure (e.g., efficiency) causes or is affected by adoption. A similar case might be made for adoptions of antivirus software, optional SAP modules, and hardware (e.g., Blackberry).

Comparisons between several organizations encounter a similar problem. Consider the introduction of a new accounting software system. An efficiency evaluation cannot occur without correcting for a sample selection bias because cross-sectional estimation procedures cannot distinguish whether efficient companies tend to be early adopters of new software, or whether the adoption of new software enhances their efficiency. Thus, when experimental settings are infeasible, we recommend the application of a quasi-experimental setting that draws on matching procedures and thus provides a viable and efficient way to correct for sample selection bias.

Finally, we note some limitations to our study. First, our research focuses on a single German university. Studying different organizations (companies) with different organizational settings and in different countries would certainly yield deeper insights into this important matter. Second, the data we use is gathered through self-reported measures, which is common practice in research and provides generally accepted validity. However, a comparison with observed measures, perhaps in an experimental setting, might enhance generalizability. Third, though we demonstrate the positive effects of spam filters on the individual and organizational levels, we cannot extend our findings to general welfare implications because it remains unclear whether spam filters represent the most efficient way to deal with spam in the long term. Some theoretical evidence suggests that the widespread use of filters might even increase the overall amount of spam (Melville et al. 2006). This question therefore should be addressed by research in order to derive long-term recommendations about spam policy.

References

- Becker, S.O. and M. Caliendo. 2004. Sensitivity analysis for average treatment effects. *Stata Journal*. **7**(1) 71-83.
- Caliendo, M and S. Kopeinig. 2008. Some practical guidance for the implementation of propensity score matching. *Journal of Economic Surveys*. **22**(1) 31-72.
- Cormack, G.V. and T.R. Lynam. 2007. Online supervised spam filter evaluation. *ACM Transactions on Information Systems*,. **25**(3) 1-31.
- Dewan, S. and F. Ren. 2007. Risk and return of information technology initiatives: Evidence from electronic commerce announcements. *Information Systems Research*. **18**(4) 370-394.
- Dewan, S., C. Shi, and V. Gurbaxani. 2007. Investigating the risk-return relationship of information technology investment: Firm-level empirical analysis. *Management Science*. **53**(12) 1829-1842.
- DiPrete, T. and M. Gangl. 2004. Assessing bias in the estimation of causal effects: Rosenbaum bounds on matching estimators and instrumental variables estimation with imperfect instruments. *Sociological Methodology*. **34**(1) 271-310.
- Duan, Z., Y. Dong, and K. Gopalan. 2007. DMTP: Controlling spam through message delivery differentiation. *Computer Networks*. **51** 2616-2630.
- Falkinger, J. 2007. Attention economies. *Journal of Economic Theory*. **133** 266-294.
- Goodman, J., G.V. Cormack, and D. Heckerman. 2007. Spam and the ongoing battle for the inbox. *Communications of the ACM*. **50**(2) 25-31.
- Hann, I.-H., K.-L. Hui, Y.-L. Lai, S.Y.T. Lee, and I.P.L. Png. 2006. Who gets spammed? *Communications of ACM*. **49**(10) 83-87.

- Harrison, G.W. and J.A. List. 2004. Field experiments. *Journal of Economic Literature*. **42**(4) 1009-1055.
- Heckman, J., R. LaLonde, and J. Smith. 1999. *The economics and econometrics of active labor market programs*. in O. Ashenfelter and D. Card (eds), *Handbook of Labor Economics*. Vol III 1865–2097. Elsevier, Amsterdam.
- Heckman, J.J., H. Ichimura, J. Smith, and P. Todd. 1998. Characterizing selection bias using experimental data. *Econometrica*. **66**(5) 1017-1098.
- Imbens, G.W. 2004. Nonparametric estimation of average treatment effects under exogeneity: A review. *Review of Economics and Statistics*. **86**(1) 4-29.
- Joseph, K. and A. Thevaranjan. 2008. Investigating pricing solutions to combat spam: Postage stamp and bonded senders. *Journal of Interactive Marketing*. Forthcoming.
- Kraut, R.E., S. Sunder, R. Telang, and J. Morris. 2005. Pricing electronic mail to solve the problem of spam. *Human-Computer Interaction*. **20** 195-223.
- Lechner, M. 2002. Some practical issues in the evaluation of heterogeneous labour market programmes by matching methods. *Journal of the Royal Statistical Society, A*. **165** 59-82.
- MAAWG. 2007. *Email metrics program - the network operators' perspective, report #5 - first quarter 2007*. (accessed Jan. 12, 2008), [available at http://www.maawg.org/about/MAAWG20071Q_Metrics_Report.pdf].
- Melville, N., A. Stevens, R.K. Plice, and O.V. Pavlov. 2006. Unsolicited commercial e-mail: Empirical analysis of a digital commons. *International Journal of Electronic Commerce*. **10**(4) 143-168.
- MessageLabs. 2008. *2007 annual security report*. (Accessed Jan. 12, 2008), [available at http://www.messageLabs.com/mlireport/MLI_2007_Annual_Security_Report.pdf].

- Morimoto, M. and S. Chang. 2006. Consumers' attitudes toward unsolicited commercial e-mail and postal direct mail marketing methods: Intrusiveness, perceived loss of control, and irritation. *Journal of Interactive Advertising*. **7**(1) 8-20.
- OECD. 2005. *Spam issues in developing countries*, <http://www.oecd.org/dataoecd/5/47/34935342.pdf>.
- Park, J.S. and A. Deshpande. 2006. Spam detection: Increasing accuracy with a hybrid solution. *Information Systems Management*. **23**(1) 57-67.
- Pavlov, O., N. Melville, and R. Plice. 2008. Toward a Sustainable Email Marketing Infrastructure. *Journal of Business Research*. Forthcoming
- Rosenbaum, P.R. 2002. *Observational studies* Springer, New York.
- Rosenbaum, P.R. and D.B. Rubin. 1983. The central role of the propensity score in observational studies for causal effects. *Biometrika*. **70**(1) 41-50.
- Rosenbaum, P.R. and D.B. Rubin. 1985. Constructing a control group using multivariate matched sampling methods that incorporate the propensity score. *The American Statistician*. **39**(1) 33-38.
- Roy, A. 1951. Some thoughts on the distribution of earnings. *Oxford Economic Papers*. **3**(2) 135-145.
- Rubin, D.B. 1974. Estimating causal effects to treatments in randomised and nonrandomised studies. *Journal of Educational Psychology*. **66**(5) 688-701.
- Sahami, M., S. Dumais, D. Heckerman, and E. Horvitz (1998), "A Bayesian approach to filtering junk e-mail," in AAAI'98 Workshop on Learning for Text Categorization. Madison, Wisconsin.
- Sianesi, B. 2004. An evaluation of swedish system of active labour market programmes in the 1990s. *The Review of Economic and Statistics*. **86**(1) 133-155.

- Sipior, J.C., B.T. Ward, and P.G. Bonner. 2004. Should spam be on the menu? *Communications of the ACM*. **47**(6) 59-63.
- Smith, J.A. and P.E. Todd. 2005. Does matching overcome LaLonde's critique of nonexperimental estimators? *Journal of Econometrics*. **125**(1-2) 305-353.
- Union, E. 2004. *Communication from the commission to the european parliament, the council, the european economic and social committee and the committee of the regions on unsolicited commercial communications or 'spam'*. (accessed January 19, 2008), [available at http://eur-lex.europa.eu/LexUriServ/site/en/com/2004/com2004_0028en01.pdf].
- Vircom. 2004. *Why spammers spam*, White Paper at www.vircom.com.
- Yoo, S.-H., C.-O. Shin, and S.-J. Kwak. 2006. Inconvenience cost of spam mail: A contingent valuation study. *Applied Economics Letters*. **13**(14) 933-936.
- Zhang, L. 2005. The CAN-SPAM act: An insufficient response to the growing spam problem. *Berkeley Technology Law Journal*. **20** 301-332.

Tables and Figures

Table 1: Central costs

Nonrecurring tasks	Time Expenditure	Costs
Administrative tasks (e. g., coordination of privacy concerns, legal and organizational clearance).	58 hours	Euro 1,740
Acquisition costs for hard- and software	-	Euro 3,000
Setup of infrastructure (e. g., installation, training)	78 hours	Euro 2,340
Recurring costs		
Maintenance and further development, support, training	268 hours	Euro 8,040
Total costs for provider		Euro 15,120

Table 2: Spam-induced working time losses

	<i>Symbol</i>	Total	D=1 spam filter	D=0 no filter
Number of observations		1000	527	473
Time expenditure per year Index = $(w_1 \cdot 250) + w_2 + w_3 + w_4 + w_5 + w_6 + w_7$	<i>Spamtime</i>	1247.22 (1686.33)	1,596.90 ^{***} (1,813.70)	857.63 (1,436.78)
(w ₁) Time expenditure for daily spam treatment ^a	<i>Daily</i>	4.87 (6.70)		
(w ₂) Time expenditure for finding university spam filter ^a	<i>Research1</i>	17.03 (20.11)		
(w ₃) Time expenditure for university filter installation ^a	<i>Installation1</i>	15.68 (18.87)		
(w ₄) Time expenditure for filter installation by assistant ^a	<i>Installation2</i>	11.73 (11.63)		
(w ₅) Time expenditure for finding alternative filter ^a	<i>Research2</i>	53.34 (54.83)		
(w ₆) Time expenditure for installation of alternative filter ^a	<i>Installation3</i>	34.40 (72.23)		
(w ₇) Time expenditure for installation of alternative filter by assistant ^a	<i>Installation4</i>	23.37 (21.24)		

***/**/* Statistical difference (two-sided t-test) from D = 0 at the 1%, 5%, and 10% levels, respectively.

^aVariable measured in minutes.

Notes: Numbers in brackets are standard deviations.

Table 3: Individual factors

	<i>Symbol</i>	Total	D=1 spam filter	D=0 no filter
Number of observations		1000	527	473
<i>Age</i>	<i>Age</i>	40.80 (11.22)	42.98 ^{***} (11.27)	38.37 (10.66)
<i>Gender (2=male)</i>	<i>Gender</i>	1.57 (0.50)	1.61 ^{***} (0.53)	1.47 (0.52)
<i>Spam quantity</i> Spam = $(w_1 \cdot w_2) / 100$	<i>Quantity</i>	16.880 (28.397)	25.62 ^{***} (32.66)	7.14 (18.37)
(w_1) Number of e-mail per day	<i>Email</i>	23.378 (31.615)	8.14 ^{***} (13.18)	4.66 (4.45)
(w_2) Spam share (in %)	<i>Spamshare</i>	44.846 (36.038)	57.18 ^{***} (34.52)	31.33 (32.59)
<i>Spam prevention</i>				
(w_1) Avoided to publish e-mail address on website ^a	<i>Website</i>	2.942 (1.692)	2.92 (1.70)	2.96 (1.68)
(w_2) Avoided transfer of e-mail address ^a	<i>Transfer</i>	2.539 (1.458)	2.57 (1.49)	2.51 (1.42)
(w_3) Alternative e-mail address was used ^a	<i>Alternative</i>	2.723 (1.833)	2.83 ^{**} (1.85)	2.60 (1.81)
(w_4) Uncommon e-mail address was used ^a	<i>Uncommon</i>	1.257 (0.796)	1.28 (0.84)	1.23 (0.75)
(w_5) Requested removal from e-mail-lists ^a	<i>Removal</i>	1.821 (1.437)	1.90 [*] (1.47)	1.73 (1.40)
<i>Spam control</i>				
(w_1) Inspection of sender ^a	<i>Sender</i>	4.733 (0.768)	4.75 (0.71)	4.71 (0.83)
(w_2) Inspection of subject ^a	<i>Subject</i>	4.721 (0.783)	4.77 ^{**} (0.69)	4.67 (0.87)
(w_3) Inspection by opening e-mail ^a	<i>Open</i>	1.670 (0.961)	1.64 (0.90)	1.71 (1.02)
<i>Level of information on spam^a</i>	<i>Infolevel</i>	2.699 (1.121)	3.07 ^{***} (1.12)	2.28 (0.97)

*** / ** / * Statistical difference (two-sided t-test) from D = 0 at the 1%, 5%, and 10% levels, respectively.

^aVariable measured on a 5-point Likert-type scale (1 = disagree – 5 = agree).

Notes: Numbers in brackets are standard deviations.

Table 4: Reactance

	<i>Symbol</i>	Total	D=1 spam filter	D=0 no filter
Number of observations		1000	527	473
<i>Behavioral changes</i>				
(w ₁) Usage of e-mail was reduced ^a	<i>Reduction</i>	1.276 (0.775)	1.32* (0.86)	1.23 (0.67)
(w ₂) Time expenditures for e-mails have been increased ^a	<i>Timeincrease</i>	2.807 (1.499)	3.20*** (1.48)	2.37 (1.39)
(w ₃) Confidence in e-mail was reduced ^a	<i>Confidence</i>	2.294 (1.282)	2.42*** (1.30)	2.16 (1.24)
<i>Spam sensitivity</i>				
(w ₁) Advertising e-mail from unknown sender ^a	<i>Unknown</i>	4.809 (0.617)	4.83 (0.59)	4.78 (0.64)
(w ₂) Unsolicited e-mail by political or other organization ^a	<i>Advertise1</i>	4.568 (0.870)	4.58 (0.86)	4.56 (0.89)
(w ₃) Unsolicited e-mail by non-commercial organization ^a	<i>Advertise2</i>	4.410 (1.010)	4.45 (0.94)	4.36 (1.08)
(w ₄) Fun e-mails ^a	<i>Fun</i>	2.594 (1.455)	2.70*** (1.49)	2.48 (1.40)
(w ₅) Advertising e-mail from business partners ^a	<i>Advertise3</i>	2.561 (1.346)	2.66*** (1.37)	2.45 (1.31)
(w ₆) E-mail with large attachment ^a	<i>Attachment</i>	2.141 (1.334)	2.23* (1.39)	2.04 (1.26)
(w ₇) Solicited commercial e-mail ^a	<i>Solicited</i>	1.711 (1.061)	1.70 (1.08)	1.72 (1.04)
<i>Spam properties</i>				
(w ₁) Spam e-mail is unsolicited ^a	<i>Unsolicited</i>	4.520 (0.925)	4.57*** (0.87)	4.46 (0.99)
(w ₂) Spam surge cannot be stopped ^a	<i>Nonstop</i>	4.209 (1.122)	4.25 (1.08)	4.16 (1.16)
(w ₃) Spam e-mail is potentially harmful for own computer ^a	<i>Damage</i>	4.185 (1.254)	4.00*** (1.37)	4.40 (1.06)
(w ₄) Perceived amount of spam received ^a	<i>PercAmount</i>	3.835 (1.353)	4.14*** (1.24)	3.48 (1.40)
(w ₅) Perceived magnitude of time expenses for spam ^a	<i>PercTime</i>	3.399 (1.385)	3.60*** (1.35)	3.16 (1.39)

***/**/* Statistical difference (two-sided t-test) from D = 0 at the 1%, 5%, and 10% levels, respectively.

^aVariable measured on a 5-point Likert-type scale (1 = disagree – 5 = agree).

Notes: Numbers in brackets are standard deviations.

Table 5: Distribution of e-mail addresses

	<i>Symbol</i>	Total	D=1 spam filter	D=0 no filter
Number of observation		1000	527	473
<i>Distribution of e-mail address to known contacts</i>				
(w ₁) E-mail address was distributed to colleagues / business partners ^b	<i>Colleagues</i>	0.916 (0.278)	0.92 (0.28)	0.92 (0.28)
(w ₂) E-mail address was distributed to friends / acquaintances ^b	<i>Friends</i>	0.544 (0.498)	0.53 (0.50)	0.56 (0.50)
<i>Distribution of e-mail address to unknown contacts</i>				
(w ₁) E-mail address was published on websites ^b	<i>Publwebsites</i>	0.721 (0.449)	0.79*** (0.41)	0.65 (0.48)
(w ₂) E-mail address was published in online directories ^b	<i>Directories</i>	0.673 (0.469)	0.73*** (0.44)	0.61 (0.49)
(w ₃) E-mail address was published in online forums ^b	<i>Forum</i>	0.069 (0.254)	0.10*** (0.30)	0.04 (0.19)
(w ₄) E-mail address was used when signing up for newsletters or webpages ^b	<i>Newsletter</i>	0.317 (0.466)	0.33 (0.47)	0.30 (0.46)

***/**/* Statistical difference (two-sided t-test) from D = 0 at the 1%, 5%, and 10% levels, respectively.

^b Variable measured as binary variable (0/1).

Notes: Numbers in brackets are standard deviations.

Table 6: Estimation results of the Logit model

<i>Independent Variables</i>	<i>Coef.</i>	<i>s.e.</i>	<i>P> z </i>
<i>Age</i>	.049 ***	.0089	0.000
<i>Gender</i>	-.334 **	.1667	0.045
<i>Quantity</i>	.016 ***	.0044	0.000
<i>E-mail</i>	.040 ***	.0156	0.010
<i>Website</i>	.070	.0597	0.244
<i>Transfer</i>	.023	.0651	0.719
<i>Alternative</i>	.006	.0529	0.912
<i>Uncommon</i>	.009	.1021	0.932
<i>Removal</i>	.024	.0565	0.668
<i>Sender</i>	.001	.1425	0.992
<i>Subject</i>	.189	.1412	0.182
<i>Open</i>	-.135	.0839	0.107
<i>Infolevel</i>	.794 ***	.0915	0.000
<i>Reduction</i>	.083	.1126	0.462
<i>Timeincrease</i>	.115 *	.0712	0.105
<i>Confidence</i>	-.008	.0715	0.912
<i>Unknown</i>	.031	.1433	0.827
<i>Advertising1</i>	.090	.0653	0.168
<i>Fun</i>	.003	.0701	0.966
<i>Attachment</i>	.015	.0767	0.843
<i>Advertise2</i>	.207 *	.1348	0.125
<i>Advertise3</i>	-.236	.1526	0.122
<i>Solicited</i>	.003	.0805	0.965
<i>PercAmount</i>	.165 *	.0851	0.053
<i>PercTime</i>	-.037	.0814	0.651
<i>Unsolicited</i>	.145	.0987	0.142
<i>Damage</i>	-.172 **	.0732	0.019
<i>Nonstop</i>	-.053	.0850	0.531
<i>Friends</i>	-.082	.1688	0.628
<i>Colleagues</i>	-.165	.2941	0.575
<i>Publwebsites</i>	.241	.2017	0.233
<i>Directories</i>	.236 *	.1721	.0170
<i>Forum</i>	.583	.3720	0.117
<i>Newsletter</i>	.027	.1796	0.883
<i>Const.</i>	-5.96 ***	1.224	0.000

Notes: Number of observations: 960; Pseudo-R²=0.247.

***/**/* Statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 7: Matching results

<i>Est.</i>	<i>Effect</i>	<i>s.e.</i>	<i>t-value</i>	<i>offsup</i>	<i>biasbef</i>	<i>biasaft</i>	<i>R²after</i>
A	-814.48	379.34	-2.15	0	20.12	15.64	0.119
B	-468.37	275.04	-1.70	74	20.12	8.30	0.062
C	-439.52	225.06	-1.95	74	20.12	4.90	0.024

offsup: Number of users outside common support region.
biasbef: Mean standardized bias (over all variables used in PS-specification) before matching.
bias aft: Mean standardized bias after matching.
Est.: A(B): Nearest Neighbor Matching without (with) common support condition,
C: Kernel matching (Epanechnikow kernel function, bandwidth parameter: 0.06) with common support condition.
s.e.: Standard errors based on 200 bootstrap replications.

Table 8: Group analysis: matching results (group-specific scores)

<i>Obs.</i>	<i>Effect</i>	<i>s.e.</i>	<i>t-value</i>	<i>offsup</i>	<i>biasbef</i>	<i>biasaft</i>	<i>R²after</i>
Number of spam mails <=10							
n(D=1): 232	156.96	86.14	1.82	14	14.84	4.35	0.019
n(D=0): 357							
Number of spam mails >10							
n(D=1): 288	-688.73	397.40	-1.73	62	16.25	9.75	0.069
n(D=0): 83							
Level of information on spam <3							
n(D=1): 152	-528.92	301.67	-1.75	3	17.32	5.9	0.028
n(D=0): 271							
Level of information on spam >=3							
n(D=1): 368	-160.92	296.65	-0.54	71	20.20	7.35	0.044
n(D=0): 169							


offsup: Number of users outside common support region.
biasbef: Mean standardized bias (over all variables used in PS-specification) before matching.
bias aft: Mean standardized bias after matching.
Estimator C: Kernel matching (Epanechnikow kernel function, bandwidth parameter: 0.06) with common support condition.
s.e.: Standard errors based on 200 bootstrap replications.

Table 9: Sensitivity analysis, unobserved heterogeneity

<i>Gamma</i>	<i>t-hat+</i> (<i>Sig+</i>)	<i>t-hat-</i> (<i>Sig-</i>)	<i>CI+</i>	<i>CI-</i>
1	-484.037 (4.0e-13)	-484.037 (4.0e-13)	-663.714	-317.766
1.1	-561.568 (4.4e-16)	-401.175 (1.5e-10)	-737.963	-258.395
1.2	-636.654 (0)	-338.821 (1.7e-08)	-790.884	-211.95
1.3	-700.058 (0)	-288.26 (7.3e-07)	-833.384	-173.577
1.4	-754.04 (0)	-245.048 (.000015)	-872.107	-140.304
1.5	-793.397 (0)	-209.412 (.000176)	-908.88	-105.836
1.6	-827.199 (0)	-178.921 (.001264)	-948.688	-71.8773
1.7	-858.578 (0)	-152.968 (.006165)	-990.629	-37.0085
1.8	-887.141 (0)	-126.06 (.02185)	-1023.3	-3.449
1.9	-915.892 (0)	-99.5326 (.059483)	-1053.24	27.0441
2	-946.357 (0)	-73.8106 (.130113)	-1083.04	54.8713

Gamma: log odds of differential assignment due to unobserved factors
Sig+: upper bound significance level
Sig-: lower bound significance level
t-hat+: upper bound Hodges-Lehmann point estimate
t-hat-: lower bound Hodges-Lehmann point estimate
CI+: upper bound confidence interval (a= .95)
CI-: lower bound confidence interval (a= .95)

Figure 1: Distribution of the Propensity Score. Common Support



Note: This figure shows the distribution of the propensity score for individuals who installed a spam filter (upper half) and those who did not (lower half). According to the MinMax-criterion 74 users from the first group must be excluded from the analysis (Treated off support), because their propensity score values lie outside the region of common support.