

Pricing Electronic Mail to Solve the Problem of Spam

Robert E. Kraut

Carnegie Mellon University

Shyam Sunder

Yale University

Rahul Telang and James Morris

Carnegie Mellon University

ABSTRACT

Junk e-mail or spam is rapidly choking off e-mail as a reliable and efficient means of communication over the Internet. Although the demand for human attention increases rapidly with the volume of information and communication, the

Robert Kraut is a social psychologist with interests in the social impact of computer technologies and especially computer-mediated communication; he is a Herbert A. Simon Professor of Human Computer Interaction in the Human Computer Interaction Institute at Carnegie Mellon University. **Shyam Sunder** is an economist with interests in the Contract Theory of Accounting and Control, experimental economics, and electronic commerce; he is a James L. Frank Professor of Accounting, Economics and Finance in the School of Management at Yale University. **Rahul Telang** is an economist with interests in the economics of information security and economic models of user choices; he is an assistant professor at the Heinz School of Public Policy and Management at Carnegie Mellon University. **James Morris** is a computer scientist with an interest in computer-mediated communication; he is Dean of Carnegie Mellon University's West Coast campus and professor of Human Computer Interaction at Mellon University.

CONTENTS**1. INTRODUCTION****2. MECHANISMS TO COMBAT SPAM**

2.1. Filters

2.2. Regulation

2.3. Pricing

The Effects of Pricing

No Targeting, Flat Rate

No Targeting, Variable Rate

Targeting Recipients

Signaling

3. DESIGN OF THE EXPERIMENTS

3.1. Experiment 1: Comparing Flat-Rate and Variable-Rating Pricing

Participants and Task

Postage Regimes

Hypotheses:

Results and Discussion

3.2. Experiment 2: Adding Targeting Information

Methods

Hypotheses

Results and Discussion

4. DISCUSSION

supply of attention hardly changes. Markets are a social institution for efficiently allocating supply and demand of scarce resources. Charging a price for sending messages may help discipline senders from demanding more attention than they are willing to pay for. Price may also credibly inform recipients about the value of a message to the sender before they read it. This article examines economic approaches to the problem of spam and the results of two laboratory experiments to explore the consequences of a pricing system for electronic mail. Charging postage for e-mail causes senders to be more selective and to send fewer messages. However, recipients did not interpret the postage paid by senders as a signal of the importance of the messages. These results suggest that markets for attention have the potential for addressing the problem of spam but their design needs further development and testing.

1. INTRODUCTION

As Herbert Simon noted over 2 decades ago, “a wealth of information” that accompanies the increasing capability of computers to generate, store, and transmit information can lead to a “poverty of attention” (Simon, 1982).

At about the same time, Poole, Inose, Takasaki, and Hurwitz (1984) empirically demonstrated that the supply of information has been growing faster than our ability to consume it. Over the past 100 years, the volume of words and images available in various forms has grown exponentially, much faster than the number of recipients and the time they can devote to processing the information. This discrepancy between the two growth rates means that a higher proportion of the information produced remains unread by many who could benefit from it, and is in that sense wasted. In addition, the glut of information makes it increasingly difficult for consumers to find what is relevant, useful, or enjoyable.

This imbalance between the wealth of information and the poverty of attention is manifest in science, literature, entertainment, conventional and electronic media, Web sites, and correspondence. In this article, we consider a particularly egregious case of imbalance between information and attention—the unsolicited, bulk electronic mail known as junk e-mail or spam.

Spam is no longer a mere nuisance. It is growing rapidly and threatens to choke off e-mail as a reliable and efficient means of communication over the Internet. Credible estimates from mid-2003 suggest that spam accounts for about 45% of all e-mail sent, up from 8% in 2001. Seventy percent of all e-mail received by AOL subscribers is spam (Hansell, 2003).

Although there are many ways of looking at the spam deluge, we take the perspective that economics lies at its core. Computer technology has cut the cost of delivering messages by orders of magnitude. Given the fixed cost of hardware and software, the marginal cost of delivering an e-mail message is negligible. The marginal cost of sending a marketing message to 1 million recipients by e-mail is less than \$2,000, whereas the same solicitation sent by conventional, bulk-rate postal mail would cost \$198,000 in postage alone, not counting paper, printing, and other distribution costs (Hansell, 2003). In the face of these low costs, it is hardly surprising that experts estimate that commercial e-mail is profitable if even one recipient in 100,000 makes a purchase (Hansell, 2003).

This study explores whether an economic solution—pricing—can solve this problem rooted in the economics of communication. A postage approach was one of the alternatives proposed at Senate hearings in May 2003 to deal with the problems of unsolicited commercial e-mail (see Gross, 2003). It also received significant attention when Bill Gates argued for solutions to spam that also included stamping e-mails.¹ In the following section, we briefly review two other approaches for dealing with spam—filtering and regulation—

1. See <http://www.cnn.com/2004/TECH/internet/03/05/spam.charge.ap/> for information.

and then summarize the basic idea of a pricing model. In Section 3, we provide some intuitions for various economic interventions that may be used to restrict spam. It can be shown, in principle, that when the value of sending or receiving messages is different across recipients, per-message pricing will benefit both senders and receivers. Such pricing encourages senders to target their messages to the most interested recipients and provides a reliable signal that allows recipients to distinguish among the messages and identify those of greatest interest to them. In Section 4, we describe two laboratory experiments that test some of the core predictions of the economic model. The results of the experiments show that per-message pricing does indeed improve targeting of messages. The final section of the article considers how the insights from these simple economic and laboratory models might be translated into practical e-mail systems that can control the inefficiencies of spam.

2. MECHANISMS TO COMBAT SPAM

2.1 Filters

E-mail filters and rules allow recipients to flag messages with a priority (Cranor & LaMacchia, 1998).² Although they are used extensively in practice, few believe that filters can be a steady-state solution to spam for several reasons. First, commercial mailers can and do continually change or disguise the source of their messages to get through the filters.³ Second, effective filter rules are difficult to program. It is beyond the capabilities of filter rules based solely on parsable attributes of messages to accurately distinguish between messages that recipients would or would not want to read (MacKay et al., 1988). Even modern, research-based Bayesian filters (Sahami, Dumais, Heckerman, & Eric Horvitz, 1998) are only 92% to 95% accurate with static e-mail corpora⁴ and fail to account well for the continual evolution of sender strategies. Third, designing and deploying effective filters consumes scarce recipient resources, getting past them consumes sender resources, and the messages themselves consume network and storage resources. Fourth, not all

2. Bälter and Sidner (2002) provide a rule-based technique that can categorize and prioritize the incoming messages. This technique does not require the use of filters.

3. In a Senate hearing, the chief executive of one Internet marketing firm boasted that he could crack most sophisticated junk e-mail filters in less than 24 hr (Krim, 2003).

4. Some vendors report that their filters have 98% accuracy (see http://www.surfcontrol.com/products/email/SEF_Effectiveness_Results.pdf). But others argue that such claims can be based on misleading statistics (see www.cs.utk.edu/~moore/opinions/spam-filters.html).

senders are spammers and the recipients employ filters with little regard for the legitimate interests of the senders. Finally, the effect of filters on spam senders' behavior is only indirect; they do not directly deter such behavior.

2.2 Regulation

Government regulation to place limits on commercial e-mail exist in both Europe and the United States. The European Union's Directive (2002/58/EC) prohibits false identities on commercial e-mail and requires that recipients must explicitly elect to receive commercial e-mail before it can be sent to them. In the United States, at least 25 states currently have some form of commercial e-mail legislation.⁵ For example, 2002 Utah legislation requires that unsolicited commercial e-mail include the sender's name and physical address, opt-out instructions, and accurate routing information and that the message contain a subject-line label indicating that it is advertising.

At the federal level, the Controlling the Assault of Non-Solicited Pornography and Marketing Act (an awkward title chosen to yield the acronym Can Spam) went into effect on January 1, 2004. The Can-Spam Act of 2004 requires commercial e-mailers to truthfully identify themselves, provide valid "opt-out" options, include a business address, and identify the message as advertising. The Federal Trade Commission can impose civil fines up to \$250 per unlawful message and 1 year of jail term for sending commercial e-mail in which the header information is misleading or inaccurate.

Although it is possible for the regulatory approach to work, the highly distributed design of the Internet makes such regulation difficult to enforce. Although the European Union has had more stringent e-mail legislation for a longer time than the United States, recent field studies find that the pattern of junk e-mail messages in the United States and Europe remains similar (Jamal, Maier, & Sunder, 2003, 2005). Recent survey research in the United States shows that Americans' concerns and dissatisfaction with e-mail have grown, not declined, since the passage of the Can-Spam Act (Rainie & Fallows, 2004). In March 2004, 3 months after the passage of the Can-Spam Act, respondents were more likely to report that spam has made being online more unpleasant, made them less trusting of e-mail, and reduced their overall use of e-mail compared to the situation in June 2003, 6 months before its passage. One problem is that a large proportion of commercial e-mail originates in regulatory sanctuaries outside of the recipients' national borders. Even if the regulatory approach were to be effective within a single country, the global scope of the Internet renders enforcement of national laws across international bound-

5. See <http://www.spamlaws.com/state/index.html> for a current list.

aries ineffective. A problem of such international scope calls for an international enforcement regime.

Moreover, interest groups and direct marketers lobby to tone down the stronger aspects of antispam legislation. Such groups argue that many regulatory controls would unfairly restrict e-mail marketing and put e-commerce at a disadvantage compared to traditional commerce. According to the American Direct Marketing Association, direct marketing is a \$7.1 billion market in the United States and consumers save close to \$1.5 billion due to direct marketing (Grimes, 2003). To date, the Can-Spam legislation seems to have had little effect on reducing the volume of spam. Indeed, some critics have argued that it has merely provided legal cover for the spammers.

2.3 Pricing

Pricing is a third approach to dealing with the problem of spam. Consider one example: senders would pay or commit to pay for each message they sent an amount that would depend on characteristics of the recipients (e.g., whether addressed individually or as a group) or the quality of service (e.g., immediate or delayed delivery). This pricing regime would be analogous to the U.S. Postal Service's tariffs,⁶ where postage is charged per message depending on whether the message is sent by express, registered, first-class, second-class, or bulk mail services. Pricing of e-mail would be an example of using a market mechanism to allocate scarce resources—human attention in this case. It is based on the fundamental principle of economics—that market-determined prices can help allocate scarce resources in a Pareto-efficient manner. In a Pareto-efficient allocation, an alternative does not exist that makes one person better off without also making at least one other person worse off. Markets are social institutions that have evolved to solve difficult society-level optimization problems using information in possession of individuals (Hayek, 1945). We conjecture that a market for attention that charges senders for each message they send can, like other markets, efficiently allocate the resources through decentralized decision making. Without revealing private information, it can induce potential message senders to rationally decide if they should send a message, after considering its cost. A market for attention can also help potential readers to decide whether to receive a message, based on the value they place on their own attention, again without revealing private information.

A pricing system for e-mail would conform to the spirit of the Internet in that it (a) would not depend on a central authority regulating the message traf-

6. See <http://pe.usps.gov/cpim/ftp/notices/not123/not123.pdf> for information.

fic, (b) would not block anyone from sending a message to anyone else, (c) would allow the electronic message pricing to evolve naturally in the marketplace, and (d) could be institutionalized so that it imposes no net cost for the users, with the money paid by the senders being credited to their respective recipients.

Pricing e-mail based on volume or importance is not a new idea. Researchers in academia and industry have explored how such systems might work (e.g., Dwork, Goldberg, & Naor, 2003; Dwork & Naor, 1993; Gross, 2003; Loder, Van Alstyne, & Walsh, 2005; Malone, Grant, Turbak, Brobst, & Cohen, 1987; Zandt, 2001) and leaders in the software industry have advocated its use (Gates, 1995). The pricing policy is based in part on the senders' assessment of the importance of a message reaching a designated recipient within a specified time period.

Recently, pricing-based mechanisms have begun to emerge in practice as well. As we discuss in a subsequent section, the Daum Corporation, the largest Internet portal in Korea, created an Online Stamp Service. It charges bulk e-mailers a fee to send messages to its subscribers (see <http://onlinestamp.daum.net/>). If a company sends more than 1,000 messages a day from a single IP address or cluster of them, Daum requires that it register and pay a fee to send mail to Daum subscribers. The fee is scaled according to usage, with the maximum of 10 KRW per recipient (approximately 0.8 cents, in U.S. dollars) charged to those who send the most mail. Mailers receive a rebate, however, proportional to the fraction of subscribers who rate their bulk e-mail as "informative." Subscribers receive points redeemable for gifts for each e-mail they rate. Similarly, Goodmail Systems (see www.goodmailsystems.com) has also proposed a system in which ISPs would allow e-mail messages carrying encrypted stamps to bypass spam filters. Goodmail proposes to sell such stamps to mass-e-mailers and provide them free of charge to other e-mail users in small numbers.

Zandt (2001) developed an economic model to explain why pricing of electronic mail should be beneficial to both senders and recipients. Zandt's model can be extended to show that pricing improves communication efficiency most when the cost varies with the number of messages and the number of recipients and when the senders receive information that enables them to differentiate among the potential responsiveness of various recipients to the message. The key insight is that by charging a small price to send a message, the pricing system shifts the task of screening messages from recipients, who don't know the content of a message, to senders, who do. Pricing rewards senders for being selective in sending messages. Senders' information about the recipients' interests enables the senders to be more selective, increasing the chances of their messages being relevant to and read by the recipients. Loder et al. (2005) also argued that senders, because they know more about

their message than receivers, should bear the burden of paying if they want receivers to read their message. One key contribution of our article is to test the impact of different pricing strategies in a laboratory setting.

Imagine a world with many senders of e-mail whose messages appeal to different recipients. For example, these senders might be advertisers, who market products that appeal to particular classes of consumers. If a sender sends a message that reaches an interested recipient and the recipient reads and responds to it, both sender and recipient benefit. For example, if a plant nursery advertises a sale on perennials and its message reaches a gardening enthusiast in the market for irises, both nursery and gardener benefit. Both senders and recipients lose if the nursery fails to send out advertisements, if the advertisements fail to reach the interested gardener, or if the gardener fails to process the message because it is buried in a flood of other messages.

However, recipients have limited time and attention. In addition, despite filters and other tools, recipients cannot perfectly determine the value of a message without spending some time processing it. As a result, they devote some of their time and attention to processing irrelevant messages at the expense of relevant ones. Thus, in the gardening example, when there are too many messages, both nursery and gardener lose because relevant messages get buried in a flood of irrelevant ones.

The Effects of Pricing

All else being equal, demand for goods and services decreases as their price increases. This downward slope of the demand function is a basic proposition in economics. In a messaging economy, senders must consider not only the costs but also the benefits of sending messages, and the net consequence for senders of raising the price of messages depends on both. The cost of sending messages is a function of the price schedule and the action of senders; the benefit is a function of the strategy of recipients under the given regime.

When the marginal cost of messages is zero, individual senders send their messages to all possible recipients they can reach. Inboxes of recipients are soon flooded with more messages than they can possibly process. As a result, both senders and recipients suffer when a recipient fails to read a relevant message. Recipients also suffer when they waste time reading or processing irrelevant messages. This was illustrated in the gardening scenario.

In the following sections, we consider the impact of two pricing schemes—fixed, in which senders are charged a fixed cost for sending any messages at all; and variable, in which senders are charged more as they send to more recipients—under two environments. In one environment, senders have no information about individual recipients and therefore cannot target them selec-

tively. In the second environment, senders have such information. In this discussion, we assume that both senders and recipients benefit if a recipient reads a relevant message.

No Targeting, Flat Rate

As long as the benefits of mass mailing exceed the fixed cost of mailing, senders will enter the market. As recipients bombarded with more messages lower the fraction of messages they read, which we call their processing rate, benefits to senders decline. When the benefits drop to the level of fixed costs, there are no incentives for more senders to enter the market. But those who remain continue their mass mailings. If the recipients are already at their processing capacity and the number of senders is already at the zero-profit equilibrium, a marginal change in the fixed cost of sending a message will not change the welfare of either the senders or the recipients. This is because any benefit that results from a reduction in number of irrelevant senders will be offset by the reduction in the number of relevant ones. If the fixed cost is sufficiently high, the zero-profit equilibrium for senders may be achieved with so few senders that the recipients are able and willing to process all the messages they receive; but again, the receivers will not benefit because these senders will send too many irrelevant messages. We restrict our own experimental investigation to a range of parameters where the number of messages is beyond the processing capacity of recipients.

No Targeting, Variable Rate

When the marginal cost of sending messages is sufficiently high, it is unprofitable to send messages to all potential recipients. But, in the absence of targeting information, any senders who choose to remain in the business can send messages only to a randomly-chosen subset of recipients. Whether the recipients read all or a subset of messages received depends on the volume of messages relative to their processing capacity, and on the expected value of reading a message. But even with a variable-rate pricing scheme, receivers and senders are unlikely to benefit by changes in prices because, as before, any benefits that result from the reduction in the volume of irrelevant messages are balanced because receivers also fail to receive relevant messages.

In summary, when targeting information is not available, neither the fixed-price nor the variable-price scheme is likely to improve sender or receiver benefits.

Targeting Recipients

When senders have information to distinguish among the potential recipients of their messages, then they can restrict their messages to the relevant receivers. But it is clear that under the fixed-rate (zero marginal cost) regime, they would still send their messages to all potential recipients as long as the fixed cost is less than the total benefit of sending to all. Therefore, the fixed pricing scheme will not lead to any improvement for sender or recipients.

In summary, we would expect that a variable pricing scheme combined with information that differentiates among recipients will cause senders to target recipients, which will increase benefits to both senders and recipients. Testable hypotheses about targeting and benefits are listed in Section 3. It should be noted, however, that although variable pricing reduces the volume of communication and the chances of information overload, these benefits depend on the price charged. In the extreme, too high a price will choke off all communication, reducing benefits to both sender and recipients to zero. A sufficiently high variable price could be just as disadvantageous as a sufficiently high fixed price.

Signaling

The pricing model discussed earlier focuses on the benefits of that result when senders are selective. Prices can also signal the sender's assessment of the value of the message to the recipient and help the recipient decide which of the competing messages deserve attention. Signaling requires economic agents to take observable and costly measures to convince others of the value or quality of their products (Spence, 1973). In many situations, a sender can signal the importance of a message by paying more to deliver it. The price that the sender paid is the major reason that recipients open their express mail before their bulk mail. The credibility of the signal should increase with the cost to the sender. A "high-priority" label on a message only works if the recipient has reason to believe, from prior experience or its cost to the sender, that such labels are not used indiscriminately. In short, a legitimate sender (but not a spammer) is willing to use a "high priority" message only when (a) the signal is credible enough so that receivers are more likely to see the message and respond to it and (b) price is high enough to deter spammers (but not too high for a legitimate sender that any benefit of signaling is negated by cost). Testable hypotheses about the use of high-low priority signals are also listed in Section 3.

3. DESIGN OF THE EXPERIMENTS

We conducted two laboratory experiments to examine the effects of different postage regimes on message-sending behavior, message-receiving behavior, and the overall social welfare of the participants. The goal across the two experiments was to see if the qualitative predictions of economic intuitions described in the previous section are consistent with empirical data. In particular, we were interested in whether different types of pricing with and without information differentiating recipients would influence how much e-mail people sent and whether they targeted their messages, how many messages recipients read and whether they read the most relevant ones, and whether these changes in sending and reading behavior would influence participants' welfare. As in most experiments, our results show only what the effects of an intervention—in this case, message pricing and targeting—can be under controlled circumstances. Producing the same effects in the real world is an additional research and design challenge.

In these experiments, players earned money by completing a crossword puzzle (the foreground task). They could ask other players for help to improve their crossword scores, and earn additional money by sending answers to them. They could label their messages as high or standard priority. They engaged in these activities under two different postage regimes. Under one, they were charged a flat rate for every message they sent, regardless of the number of recipients or priority label. Under the second postage regime, they paid per recipient and paid more for sending high-priority messages. We provide the details later.

The experiments were designed to capture a situation with the following conditions:

1. Players had a foreground task—in this case, completing crossword puzzles. Their performance in this task had personal value to them. Their skill in the foreground task determined the opportunity costs of the time they spent on processing e-mail.
2. Players had an incentive to send electronic mail to other players—asking or responding to questions—which may or may not have been of value to the recipients.
3. Players received more messages than they could read in the time provided. Under these conditions, players had to allocate their time between the foreground and e-mail tasks and among the messages they received.
4. The value of any given message varied across players.

3.1. Experiment 1: Comparing Flat-Rate and Variable-Rate Pricing

Participants and Task

Four to six university undergraduates played a game that required each to fill out a unique crossword puzzle in each of five 10-min rounds. The participants received a monetary reward for their performance in the crossword puzzle task, based on two criteria. They earned \$.05 for each correct letter filled in their crossword puzzle. They also earned money for helping others. They were given clue sheets containing clues and answers for the words appearing on other players' puzzles.

Participants could send e-mail messages to one or more players requesting help or responding to such requests. To simplify data analysis, players were required to characterize the content of their messages and were allowed to include only a single query or answer in each message. On receiving a message purporting to contain help, the recipient indicated whether he or she would use the answer. If so, the sender received \$.15 per letter for providing the help.

Although exchanging messages was potentially rewarding, it also competed for the time a participant could devote to work on the puzzle. To emphasize the opportunity cost of messaging, the screen blanked out for 5 sec when participants clicked to open an incoming message.

Standard-priority messages delivered with a 20-sec delay appeared in the recipient's inbox in standard font. High-priority messages were delivered immediately and were displayed in the recipient's inbox in bold font.

To make attention a scarce resource, a server was programmed to send out e-mail messages approximately every 7 sec. The server-generated messages, all standard-priority, appeared to come from another player and contained useless text.

None of the messages contained a subject field. As a result, the message's priority level and the identity of its sender were the only clues the recipient had about its potential value. Elimination of the subject field enabled us to focus the experiment on the consequences of pricing on communication. Given the simple communication involved in the crossword puzzle task, we could not allow the subject line to become a substitute for the message itself.

Postage Regimes

We randomly assigned all participants in a session to one of two different postage regimes—an inexpensive fixed-rate regime and a more costly variable-rate regime.

In the fixed-rate postage condition, both standard-priority and high-priority messages cost \$.02 each, independent of the number of addressees. We sometimes refer to this condition as message-based pricing. In the variable-rate postage condition, standard messages cost \$.02 and high-priority messages cost \$.04 per addressee. We sometimes refer to this condition as recipient-based pricing. A high-priority message to three people cost \$.02 (\$.02 \times 1 message) in the fixed-rate condition, and \$.12 (\$.04 \times 3 addressees) in the variable-rate condition.

At the end of the sessions, participants filled out a debriefing questionnaire and received their earnings in cash (approximately \$23 for 100 min, on average).

The Customized E-Mail System

We built an experimental environment, including a customized e-mail system, for exchanging messages and keeping track of costs and earnings. Four windows appeared on participants' screens (see Figure 1). One of these (top left) was a customized version of Microsoft Excel™ containing a crossword puzzle and its clues. One window (bottom middle) was used to send e-mail messages and another (right) was used for viewing the inbox and the content of received messages. The fourth window (bottom left) was used for calculating the players' cash earnings. Values in the earnings window were updated every 30 sec to reflect the rewards from helping other players and the postage charges. At the end of each round, a computer graded the crossword puzzles and updated players' earnings.

Analysis. We expected that the postage regime under which participants played the game would directly influence their sending decisions—number of messages, number of recipients, and use of the high-priority option. These decisions would, in turn, influence recipients' behavior—reading messages and responding to help requests. Finally, these factors directly, and in interaction with the postage regime, should influence players' earnings.

Because messages were nested within players within rounds of the experiment and because players were nested within experimental session, we used hierarchical linear models to account for the nonindependence of observations (Bryk & Raudenbush, 1992). For analyses about particular messages and responses to them, the message was the unit of analysis. For analyses about communication value and performance outcome, the player within round within session was the unit of analysis. In all models, the university where the session was run, the number of players involved, and the round within session, were included as control variables.

Figure 1. Crossword Puzzle Control

The screenshot displays a crossword puzzle control interface. The main window shows a crossword puzzle grid with some letters filled in. A 'Name' field contains 'Jane Doe' and a 'Puzzle' field contains '533.xls'. A 'Score Display' window shows player information and scores. A 'Send Message' window is open, showing delivery options and help info. A 'Messages Received' window shows a list of messages with columns for Mail #, Viewed, From, Postage, and Time Sent.

Score Display

Player: Jane Doe
 Number: 3
 Puzzle: 533.xls Round: 0
 Score: 10.19
 Credit For Arriving On-Time: 4.00
 Credit For CrossWord Answers: 4.10
 Total Possible: 210
 Correct: 82
 Credit For Help You Provided: 2.85
 Total Length of Accepted Answers: 19
 Postage For Sent Emails: -0.64
 Immediate Delivery: 16
 Standard Delivery: 6
 Total: 10.19

Send Message

To: Person 1 Person 2 Person 3 Person 4 Person 5 Person 6 Everyone

Delivery: Standard (\$0.02 each) Immediate (\$0.04 each)

Help Info: This is a request for help. I am providing an answer. No requests for help or answers are in this message.

Message:

Status: Send Message

Messages Received

Mail #	Viewed	From	Postage	Time Sent
13433		Player # 5	standard	3:09:00 PM
13435		Player # 4	standard	3:09:06 PM
13437		Player # 1	standard	3:09:15 PM
13439	X	Player # 5	standard	3:09:23 PM
13441		Player # 2	standard	3:09:29 PM
13443		Player # 2	standard	3:09:35 PM
13445		Player # 1	standard	3:09:43 PM
13451	X	Player # 1	standard	3:09:48 PM
13457		Player # 2	standard	3:09:57 PM
13459		Player # 2	standard	3:10:02 PM
13463	X	Player # 1	immediate	3:10:16 PM
13461		Player # 1	standard	3:10:09 PM
13465	X	Player # 1	standard	3:10:25 PM
13467		Player # 1	standard	3:10:33 PM
13475	X	Player # 1	immediate	3:10:52 PM
13479		Player # 4	immediate	3:10:59 PM
13471		Player # 5	standard	3:10:45 PM
13473		Player # 4	standard	3:10:53 PM
13481		Player # 5	standard	3:11:07 PM
13483		Player # 4	standard	3:11:16 PM
13485		Player # 4	standard	3:11:21 PM
13487		Player # 2	standard	3:11:27 PM
13489		Player # 2	standard	3:11:35 PM
13491		Player # 4	standard	3:11:42 PM

Message: 13475 From: 1 Time Sent: 3:10:52 PM
 Whales breath through a blowhole

Measures

Sending behaviors included the number of messages sent during a round, the number of recipients per message, and the proportion of messages sent by high priority.

Attentional behaviors included the proportion of messages read and responded to and the percentage of help accepted.

Performance outcomes included a player's earnings during a round—the sum of the amounts earned from completing puzzles and offering help, less the cost of postage.

Hypotheses

We divide our testable hypotheses into three categories, encompassing sending behaviors (H1), recipients' attentional behaviors (H2), and performance outcomes for the senders and recipients (H3). All hypotheses compare behavior under variable-rate, recipient-based pricing, and fixed-rate, message-based pricing.

Regarding sending behavior

- H1a: Senders send fewer messages under variable-rate pricing (compared to fixed-rate pricing).
- H1b: Senders send each message to fewer recipients under variable-rate pricing.
- H1c: Senders use the high-priority designation less often under variable-rate pricing.

Regarding attentional behavior

- H2a: Recipients read a larger proportion of all messages they receive under variable-rate pricing.
- H2b: Recipients increase reading of high-priority messages relative to standard-priority messages under variable-rate pricing.
- H2c: Recipients reply to a larger proportion of messages under variable-rate pricing.

Regarding overall performance

- H3a: Players earn more from completing the puzzle task under variable-rate pricing.
- H3b: Players earn more from offering help under variable-rate pricing.

Results and Discussion

Descriptive Statistics

Figure 2 shows samples of the messages players exchanged. Players sent questions, asking for the words associated with the clues for crossword puzzles. In addition, they sent answers, either in response to questions addressed to them or proactively, without responding to a question. Thirty-five percent of their messages were questions and 63% were answers.⁷ They read 60% of the messages they received and replied to 20.6% of the questions they received. Forty percent of questions received at least one reply. Senders used high priority for 83% of their messages. Across all conditions, they spent \$.91 per round on postage.

7. Note that there were more answers than questions because participants proactively sent answers to other players as well as sending them in reply to specific questions. In addition, a question sent to several players could generate multiple answers.

Figure 2. Examples of Exchanges Between Players

Example	Content	Number of Recipients
Question-Answer sequences		
1	Q: egyptian river? A: nileriver A: Nile?	2 1 1
2	Q: traveling musician?—immediate delivery only—lots of money for help! A: just a guess: troubador	4 1
3	Q: marijuana? 3 letters A: pot	4 1
4	Q: 6 letter word for lodestone, starts with m A: answer is MAGNET	1 1
Unanswered questions		
5	Q: first state in the union?	5
6	Q: What is slang for a harmonica?	1
7	Q: cheese filled pasta?	4
8	Q: help me for What is an indian sauce???	1
Unsolicited answers		
9	A: unable to do things = incompetent	4
10	A: overruled or sustained OBJECT	4
11	A: ionic is a charged particle	4
12	A: "TYRANNOSAURUS means 'terrible lizard'"	4

Effects of Postage Regime on Message Sending

The postage regime influenced message sending in economically sensible ways. Figure 3 shows these effects. In the variable-rate postage condition (paying per recipient and paying more for high-priority messages), players rationed their messages. They sent a third fewer messages per round and addressed each to fewer recipients. In the fixed-price condition, where the cost of a high-priority message was the same as that of a standard message, the high-priority label became routine and players used it for almost every message; the high-priority label was used more sparingly in the variable-pricing condition. The results confirm hypotheses H1a, H1b, and H1c.

Effects of Message-Sending Behavior on Attention

The players' behavior in sending messages had large effects on the attention recipients paid to them. Recipients read a higher proportion of their mes-

Figure 3. Experiment 1: Effects of Postage Regime

Dependent Variable	Postage Regime		<i>p</i>
	Fixed	Variable	
Sending behavior			
Unique messages sent ^a	18.7	12.6	.13
Recipients per message ^b	3.3	2.4	.08
Percentage high-priority message ^b	93	69	.01
Attention behavior			
Percentage messages read ^b	70	61	.38
Percentage questions replied to ^b	20	23	.75
Percentage help accepted ^b	22	42	.15
Performance			
Total earnings ^a	\$2.73	\$1.69	.05
Puzzle completion earnings ^a	\$2.30	\$2.09	.40
Reward for help ^a	\$0.82	\$0.84	.87
Paid for postage ^a	\$0.37	\$1.22	.06

Note. N = 3888 messages nested within 55 players nested within 11 experimental sessions.

sages when fewer messages were sent ($r = -0.14$, $p < .02$) and when each original was sent to fewer recipients ($r = -0.19$, $p < .001$). They were especially likely to read messages labeled as high priority ($r = .43$, $p < .0001$). They read 77% of the high-priority messages versus only 16% of the standard-priority messages. This large difference in reading rates occurred partly because recipients could be sure that high-priority messages were sent by real players and not the spam server.

Variable-rate pricing reduced the number of messages sent and recipients per message, while increasing the use of the high-priority label. These effects cancelled each other out. As a result, recipients read approximately the same percentage of messages in both the fixed-rated and variable-rate pricing conditions ($p = .38$). Thus, there was no support for hypotheses H2a and H2c.

We had expected that high-priority messages would be read more when this designation was costly (i.e., in the variable-rate postage condition). This was not the case, however. Players were more likely to read high-priority messages than standard ones and this increase was approximately the same in both the variable-rate pricing condition (80% vs. 20%) and the fixed-rate pricing condition (74% vs. 12%, $p = .68$). Thus, hypothesis H2b was not supported.

Effects of Postage Regime on Economic Outcomes

We expected that the effects of the variable-rate postage regime—in particular reducing the volume of communication—would help players more effi-

ciently allocate their time between their foreground task (the crossword puzzles) and communication. We therefore expected them to earn more money under the variable-rate postage condition, at least before subtracting the postage costs.

This prediction was not confirmed. Participants in the variable-rate postage condition netted significantly less than those in the fixed-rate postage condition ($p = .05$). They earned insignificantly less for completing puzzles and for providing answers to fellow players but spent substantially more for sending messages ($p = .06$).

Although this result is inconsistent with the hypotheses, it is understandable. Without access to targeting information in this experiment, the increase in the cost of communication induced players to send fewer messages to fewer recipients. They sent fewer questions, reducing the probability of reaching a recipient who had the information they needed. They also sent fewer answers, reducing their chance of sending information to someone who needed it. The reduction in communication volume cancelled any gains of lower information overload.

3.2 Experiment 2: Adding Targeting Information

Experiment 2 was designed to examine messaging behavior when senders of messages have information so they can selectively target recipients. In Experiment 1, senders had no information to discriminate among the potential recipients. In Experiment 2, we introduced the possibility of targeting communication by letting players know the domains of expertise that others in the game had. With this information, a player who needed the answer to a question about science, for example, could send a question to the science expert. This was the case in Figure 2, example 4, where one player asked the science expert for a six-letter word for lodestone. Our prediction was that players would be more likely to address their questions to experts when communication was more costly.

To contrast communication that allowed targeting with communication that did not, we introduced a new type of message—advertisements. Senders received no benefit from targeting advertisements. Senders were rewarded when recipients opened advertisements and the reward they received did not depend on which recipients opened it. In addition, unlike the case of questions, which were valuable to recipients when they answered them, recipients were penalized for opening advertisements, because it left them with less time to do their crossword puzzles.

Experiment 2 was designed to examine the differential impact of charging for communication on senders' behavior when recipients were identifiably heterogeneous, as they were in respect to questions, compared to when recipients were not heterogeneous, as they were with respect to advertisements.

We predicted that expensive, variable-rate pricing would cause senders to reduce their communication more when they could target recipients (i.e., when sending questions) and that senders would take advantage of targeting information more when communication was more expensive (i.e., under variable-rate pricing).

Methods

Procedures for Experiment 2 were similar to Experiment 1, with the following differences.

Advertising

All participants could send questions and answers, as in Experiment 1. In addition, two participants in each session could send advertising messages, earning \$0.25 per recipient who opened an advertisement. Opening an advertisement covered the recipient's crossword puzzle for 5 sec.

Expertise

As in Experiment 1, answers to puzzles were distributed among players. In Experiment 2, answers were distributed so that each player was an "expert" in one of the six domains (science, law, cooking, geography, music, living things). The expert was 80% likely to have an answer, whereas the others had only a 50% chance of having an answer. All players were given a table showing the assignment of expertise across players.

Hypotheses

As in Experiment 1, the testable hypotheses for Experiment 2 also are divided into three categories: sending behaviors (H1), attentional behaviors (H2), and overall economic outcomes (H3). All hypotheses compare behavior under expensive, variable-rate, recipient-based pricing relative to inexpensive, fixed-rate, message-based pricing.

Sending Behaviors

H1a, H1b, and H1c are also applicable to Experiment 2. In addition

- H1d: Senders will target sending questions more under variable-rate pricing than under fixed-rate pricing. That is, they are more likely to send a question to a relevant expert.

- H1e: Senders will send advertising messages to more recipients per message than they do for questions, and this difference will be greater under variable-rate pricing compared to fixed-rate pricing.
- H1f: Senders will use the high-priority designation more for sending questions than for advertisements, and answers and this difference will be greater under variable-rate pricing compared to fixed-rate pricing.

Attentional Behaviors

H2a, H2b, H2c, and H2d are also applicable to Experiment 2. In addition

- H2e: Recipients are more likely to read advertising messages with fixed-rate than with variable-rate pricing.
- H2f: Recipients are less likely to open and read advertising messages than nonadvertising messages.

The overall efficiency hypotheses H3a and H3b should also apply to Experiment 2.

Results and Discussion

In Experiment 2, participants had targeting information available when asking questions but not when sending advertising messages. Participants used high priority for 56% of the messages. Across all conditions, they spent \$0.64 per round on postage. Thirty-three percent of their messages were questions, 45% were answers, and 23% were advertisements. Recipients opened 55% of all messages they received and replied to 45% of the questions.

Effects of Postage Regime on Message Sending

As in Experiment 1, costly communication influenced participants' sending behavior in economically sensible ways. Under the variable-rate postage, participants sent fewer messages, addressed each message to fewer participants, and used high priority less often (see Figure 4).

In addition, as predicted, participants in the variable-rate postage condition were more likely to target their queries to the experts (see Figure 4). Players sent most of their messages to multiple recipients and did so more for advertisements (93%) than for questions (74%). To test H1d, that variable-rate pricing along with targeting would reduce the number of recipients per message, we examined the interaction between pricing regime and message type on the percentage of players to whom they sent their messages. Players had no targeting information when they sent advertisement but did for questions.

Figure 4. Experiment 2: Effects of postage regime

Dependent Variable	Postage Regime		<i>p</i>
	Fixed	Variable	
Sending behavior			
Unique messages sent ^a	13.6	9.8	.01
Recipients per message ^b	3.6	3.0	.02
Percentage high-priority message ^b	81	41	.001
Percentage questions matching addressee's expertise ^b	32	45	.01
Attention behavior			
Percentage messages read ^b	42	45	.99
Percentage questions replied to ^b	4	5	.60
Performance			
Total Earnings ^a	\$2.07	\$1.47	.01
Puzzle Completion Earnings ^a	\$1.61	\$1.39	.02
Reward for help ^a	\$.76	\$.66	.47
Reward for advertising ^a	\$.06	\$.23	.19
Paid for postage ^a	\$.36	\$.81	.001

Note. N = 7609 messages nested within 120 players nested within 24 experimental sessions.

^aPlayer as the unit of analysis. ^bMessage as the unit of analysis.

Consistent with H1d, variable-rate postage inhibited sending to multiple recipients most when senders could differentiate among recipients. We plot the interaction in Figure 5, showing the percentage of potential recipients sent to for advertising and questions under fixed-rate and variable-rate postage. The significant interaction between message type and postage regime shows that costly communication inhibited multiple addressees more for questions than for advertising. Because no targeting information was available for advertisements, the postage regime did not affect them as much.

We also tested the interaction of high-priority messaging with targeting. We reasoned in Section 3 that when targeting information is available, users tend to send high-priority messages as opposed to standard messages in the variable-postage regime, to signal the value of their message to the recipient. Because targeting information was available only for questions and not for advertisements, we hypothesized that this use of pricing to signal quality would occur primarily for questions. This hypothesis was confirmed. As shown in Figure 6, variable-rate pricing caused players to send a lower percentage of their messages with the high-priority label but this drop was steeper for advertising, which allowed no targeting, than for questions, which did allow targeting; for the interaction, $F(1, 5332) = 3.60$, $p = .06$.

Figure 5. Percentage of all recipients addressed by message type and postage condition.

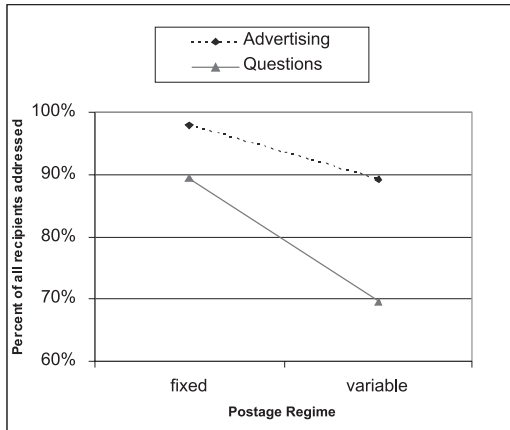



Figure 6. Percentage of high-priority messages by message type and postage condition.



Effects of Postage Regime on Attention

Players were more likely to read high-priority messages (59% of high-priority messages were read vs. 41% of standard-priority messages), highlighting the importance of signaling. Although recipients read and replied to slightly more of their messages in the variable-rate postage condition, neither effect approached statistical significance (see Figure 4). As in Experiment 1 and contrary to expectations (H2b), variable-rate postage did not enhance the signaling power of the high-priority option (for the pricing regime \times high priority

interaction, $p > .50$). Neither did the variable-rate postage increase the likelihood that players would reply to questions they received (for the pricing regime \times high priority interaction, $p = .32$).

Participants were 18% less likely to open advertising messages than those containing questions or answers ($p < .001$). Because the messages carried no subject lines, senders' identity was their only cue to content. Recipients' reluctance to open advertising messages meant that they gradually learned to open fewer messages from advertising-enabled players than from others.

Effects of Postage Regime on Economic Outcomes

In terms of outcomes, hypotheses H3a and H3b were disconfirmed. Participants in variable-rate postage sessions earned reliably less than those in fixed-rate sessions. They both earned less from completing the puzzles and paid more for postage (see Figure 4).

Although we argued that the total welfare would increase under variable-rate pricing when targeting was possible, we did not observe this result here. There are two main reasons. First, targeting information was available only for queries and not for advertisement and answers. Therefore, the information overload did not decrease sufficiently, even under variable-rate pricing. Second, the high-priority messages were probably too costly relative to the expected benefit from a message. As a result, people sent too few messages and too few high-priority messages in the variable-rate pricing condition.

4. DISCUSSION

To summarize, this article examined the economics of pricing e-mail on the behavior of e-mail senders and recipients and on the benefits they gained. Under the current regime, with practically free e-mail, it is economically rational for individual advertisers and other bulk mailers to send their messages to as large an audience as they can. The consequences of this "commons" approach to the Internet are, however, suboptimal for both advertisers and their targets. Because recipients can only read, understand, and respond to a fraction of the mail they receive, and because they can't adequately distinguish useful mail from the worthless without processing it to some degree, they therefore waste time on messages that have no value for them and ignore some of the messages that would have been valuable for both themselves and the senders. Even under conditions where senders can easily distinguish between interested and uninterested consumers, it is still economically rational to send messages to all possible consumers on the chance that some putatively uninterested ones might still respond.

Charging a fixed fee to send e-mail would not change this situation and would not increase the benefits to either senders or recipients. Under a fixed fee, some advertisers would stop sending mail altogether, but those who find it profitable to send mail would still send it to all possible recipients. Some proposals to curb spam are based on the premise that reducing the volume of communication per se would make the recipients better off and ignore the consumers' loss from the failure to receive relevant messages. Economic models predict that benefits to senders and recipients increase when senders can distinguish the interested from the uninterested consumers and when they incur a per recipient cost for sending messages. Under these conditions, senders are motivated to send only to interested recipients, increasing their own and the recipients' benefits.

The article reports the results from two experiments testing these hypotheses. The experiments examined the consequences of making electronic communication more costly, by charging per recipient and by imposing a fee for high-priority messages. In both experiments, variable-pricing schemes had predicted consequences: fewer messages were sent, each message was sent to fewer recipients, and the high-priority option was used less often. Experiment 2 examined the prediction that people would be least likely to spam (i.e., send messages indiscriminately to all available recipients) when they had information relevant for targeting potential recipients and when they incurred a variable (per-recipient) cost for communication. This prediction was confirmed. Senders were more likely to target their messages to relevant recipients under variable-rate postage, when they had information to differentiate among the recipients (i.e., for questions but not for advertisements).

Evidence is strong that the message-sending behavior influenced the frequency of recipients' reading and responding behavior. Players read a higher proportion of messages when they had fewer to read and when the messages were addressed to fewer recipients. In both experiments, they read and responded to more messages labeled high priority than standard priority. In neither experiment, however, did players read or respond to a higher proportion of messages under variable-rate pricing. In addition, in neither experiment did the cost of high-priority messages change the signaling power of the high-priority designation.

These changes made communications more efficient, reducing volume and increasing relevance. We had expected that the variable pricing would benefit both senders and recipients if senders could target messages. In Experiment 2, where senders had information to differentiate among recipients, senders indeed targeted more under variable-rate pricing. However, the net benefits to senders and recipients declined because senders sent too few messages. We suspect that the overall social welfare did not increase under variable-rate pricing because we chose parameters that made communication too


expensive relative to the rewards associated with it and did not make message recipients sufficiently differentiable. As a result, the costs associated with an undersupply of valuable communication overwhelmed the benefits associated with more efficient, targeted communication.

Both sender and recipient should benefit from relevant communication. Our economic argument is that both senders' ability to target recipients and the right pricing scheme are required for higher communication efficiency and higher benefits to senders and recipients. Similarly, Zandt (2001) hypothesized that the benefits from postage would depend on the diagnosticity of the cues senders use to address their messages. Empirical studies that vary the diagnosticity of the information, pricing, and expected value of a message are needed to test this prediction further.

If further theoretical modeling and empirical research demonstrate the potential of markets for attention, we must overcome a number of challenges before these ideas can be implemented in real-world systems. With the design of an appropriate pricing scheme, both senders and recipients could benefit from the pricing of e-mail. One hurdle in implementing such a vision is convincing people to accept having to pay for what is currently a free service. Indeed, GoodmailSystems, a company offering an e-mail stamps service, addresses this challenge directly in its "frequently-asked questions" list when it states, "Isn't it a fundamental principle of the Internet that e-mail should be free?" (GoodmailSystems, 2004). This challenge is not insurmountable. The transition from free to paid service has been accomplished in other domains, including the shift from flat-rate pricing of telephone service to metered pricing, the shift from free broadcast TV to fee-based cable and pay-per-view TV, and the shift from tax-supported, free roads to usage-based toll roads. As these examples indicate, people will pay for better service.

The pricing of e-mail need not hurt the use of e-mail by innocent nonspam senders, especially those of insufficient means. Those who have approximately balanced incoming and outgoing message traffic will have approximately balanced payments and receipts, with little net cost. Communication with acquaintances (e.g., people already in one's address book) can easily bypass the pricing system. Loder et al. (2005) proposed a pricing model that places the cost of a stamp in escrow, forcing senders to put a small amount of their money at risk of forfeiture in case their messages are considered spam by the recipients. Presumably, under such a system, recipients would return payment of the escrow for family, friends, and acquaintances with whom they have personal communication and for firms with which they do legitimate business. The Vanquish service (see <http://www.vanquish.com/news.shtml>) implements a similar scheme in a peer-to-peer e-mail system. Such a service protects nonspam senders from having to pay for e-mail.

Figure 7. Daily bulk e-mail volume before and after the initiation of Daum's Online Stamp service.



We noted in the introduction that the Daum Corporation started an Online Stamping service to curb spam. The Online Stamp service became operational April 1, 2002. Figure 7 shows the daily volume of bulk e-mail sent to Daum subscribers before and after the initiation of the online stamp service. Bulk e-mail traffic in the 3 months following the introduction of the service was 46% of the traffic in the prior 3 months. Company officials report that the quality of commercial e-mail also changed following the introduction of the Online Stamp service, to become more informative. Unfortunately, however, the company collected no systematic data about whether the introduction of the Online Stamp system changed the likelihood of e-mailers targeting particular subscribers or the likelihood of subscribers reading commercial e-mail or responding to it. As a result, this service provides no evidence about whether the reduction in bulk e-mail changed the benefits for either senders or recipients.

In the pricing scheme used in the experiments, senders paid a fee to send messages on a per-recipient basis, with the postage going to a bank (the experimenters). The Daum service uses a similar system. This is, however, only one of many possible pricing alternatives. For example, the postage could go to recipients of messages rather than to a bank, either on delivery or when read. The latter pricing scheme would not only reduce the volume of messages and improve targeting but might also induce recipients to read messages that senders consider important. Loder et al. (2005) proposed this model. In this sense, senders are directly buying the recipient's time, much as some telephone services reduce fees if a caller agrees to listen to advertisements or as broadcast TV and radio stations provide news and entertainment in exchange for the consumer's willingness to receive advertising messages. In ad-

dition, the information about the postage attached to the message could be made available to the recipient before he or she decides whether to open the message.

Although our economic arguments and experiments have treated postage for e-mail as a monetary cost, money is not intrinsic to any of our proposals. Others, for example, have examined cost functions based on the CPU processing needed to solve a puzzle (e.g., Back, 2002; Dwork & Naor, 1993) or on computer memory. We believe our analysis and empirical results would generalize to these cashless pricing mechanisms.

The pricing scheme we described in this article assumes usage-based pricing that is constant for each recipient. It would be possible to elaborate this model, so that the price charged varied with the value of the recipient's time. For example, Horvitz, Kadie, Paek, and Hovel (2003) and Hudson et al. (2003) have developed algorithms to assess an individual's interruptibility. Postage could vary with a recipient's interruptibility (Fahlman, 2002). This type of pricing model would be especially relevant to computer chat, instant messaging, telephone calls, and other synchronous communication, where traditionally recipients accept communication when it arrives.

Regardless of the exact pricing mechanism, more research is needed to identify appropriate cost functions so that they reduce the volume of communication and increase the targeting of messages without reducing communication to harmful levels.

NOTES

Background. This article is based on research originally reported in Kraut, Sunder, Morris, Telang, Filer, and Cronin (2002).

Acknowledgments. Matt Cronin helped plan this research, and Darin Filer wrote the crossword puzzle software used in the experiments. Catherine Chan, Jessica Morgan, Yuquin Ren, and Allan Roth helped to run the experiments. Nicole Lazard provided statistical advice. Yonnie Kim and the Daum Corporation generously provided raw data about Daum's Online Stamp system.

Support. This project was supported by the Intel Corporation and NSF Grant IIS-0325351.

Authors' Present Addresses. Robert E. Kraut, Human Computer Interaction Institute, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213. E-mail: robert.kraut@cmu.edu. James Morris, Carnegie Mellon University West Coast Campus, Office 214, Building 23 (MS-23-11), Moffett Field, CA 94035. E-mail: james.morris@cmu.edu. Rahul Telang, Heinz School of Public Policy and Management, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213. E-mail: rtelang@andrew.cmu.edu. Shyam Sunder, School of Management, Yale University, PO Box 208200, 135 Prospect Street, New Haven, CT 06520-8200. E-mail: shyam.sunder@yale.edu.

HCI Editorial Record. First manuscript received June 3, 2003. Revision received May 24, 2004. Accepted by Paul Moody. Final manuscript received October 25, 2004.
— Editor

REFERENCES

- Back, A. (2002). *Hardcash: A denial of service counter-measure*. Retrieved June 1, 2003, from www.cyperspace.org/~adam/hashcash/hashcash.pdf
- Bälter, O., & Sidner, C. L. (2002). Bifrost Inbox Organizer: Giving users control over the inbox. *Proceedings of the ACM Nordic Conference on Human-Computer Interaction*. New York: ACM.
- Bryk, A., & Raudenbush, S. (1992). *Hierarchical linear models: Applications and data analysis methods*. Newbury Park, CA: Sage.
- Cranor, L. F., & LaMacchia, B. A. (1998). Spam! *Communications of the ACM*, 41(8), 74–83.
- Dwork, C., & Naor, M. (1993). Pricing via processing for combating junk mail. In E. Brickell (Ed.), *Advances in Cryptology-Crypto '92: Lecture notes in computer science 740* (pp. 139–147). New York: Springer-Verlag.
- Dwork, C. A., Goldberg, X., & Naor, M. (2003). On memory-bound functions for fighting spam. *Proceedings of the 23rd Annual International Cryptology Conference (CRYPTO 2003)*. Santa Barbara, CA: Springer-Verlag.
- Fahlman, S. E. (2002). Selling interrupt rights: A way to control unwanted e-mail and telephone calls. *IBM Systems Journal*, 41, 759–766.
- GoodmailSystems. (2004). *Frequently asked questions*. Retrieved May 22, 2004, from <http://www.goodmailsystems.com/faq.html>
- Grimes, A. (2003, May 22). Digits: Spam pays. *The Wall Street Journal*, pp. B3.
- Gross, G. (2003, May 21). *Spam hearing: E-mail tax, international treaty proposed*. *InfoWorld*. Retrieved May 22, 2004, from www.infoworld.com/article/03/05/21/HNspamtax_1.html
- Hansell, S. (2003, April 22). Internet is losing ground in battle against spam. *The New York Times*, p 1.
- Hayek, F. A. (1945). The uses of knowledge in society. *American Economic Review*, 35, 519–530.
- Horvitz, E., Kadie, C. M., Paek, T., & Hovel, D. (2003). Models of attention in computing and communication: From principles to applications. *Communications of the ACM*, 46(3), 52–59.
- Hudson, S. E., Fogarty, J., Atkeson, C. G., Avrahami, D., Forlizzi, J., Kiesler, S., et al. (2003). Predicting human interruptibility with sensors: A Wizard of Oz feasibility study. *Proceedings of the CHI 2003 Conference on Human Factors in Computing Systems*. New York: ACM.
- Jamal, K., Maier, M., & Sunder, S. (2003). Privacy in e-commerce: Development of reporting standards, disclosure and assurance services in an unregulated market. *Journal of Accounting Research*, 41, 285–309.

- Jamal, K., Maier, M., & Sunder, S. (2005). Enforced standards versus evaluation by general acceptance: A comparative study of e-commerce privacy disclosure and practice in the US and the UK. *Journal of Accounting Research*, 43(1) 73–96.
- Kraut, R. E., Sunder, S., Morris, J., Telang, R., Filer, D., & Cronin, M. (2002). Markets for attention: Will postage for email help? *Proceedings of the 2002 Conference on Computer Supported Cooperative Work*. New York: ACM.
- Krim, J. (2003, May 22). A spammer speaks out: In Hill testimony, bulk e-mailer says Internet providers use same tactics. *The Washington Post*, p. A01.
- Loder, T., Van Alstyne, M., & Walsh, R. (2004). An economic response to unsolicited communication. Working paper, April 7, 2005.
- Malone, T. W., Grant, K. R., Lai, K., Rao, R., & Rosenblitt, D. (1987). Semi-structured messages are surprisingly useful for computer-supported coordination. *ACM Transactions on Office Information Systems*, 5, 115–131.
- Malone, T. W., Grant, K. R., Turbak, F. A., Brobst, S. A., & Cohen, M. D., (1987). Intelligent information-sharing systems. *Communications of the ACM*, 30(5), 390–402.
- Mayer, M. (1977). The telephone and the uses of time. In I. de Sola Pool (Ed.), *The social impact of the telephone* (pp. 225–245). Cambridge, MA: MIT Press.
- Poole, I., Inose, H., Takasaki, N., & Hurwitz, R. (1984). *Communication flows: A census in the United States and Japan*. New York: North-Holland.
- Rainie, L., & Fallows, D. (2004, March). *The impact of CAN-SPAM legislation* (Pew Internet and American Life Project). Retrieved May 22, 2004, from <http://www.pewinternet.org/reports/toc.asp?Report=116>
- Sahami, M., Dumais, S., Heckerman, D., & Horvitz, E. (1998). *A Bayesian approach to filtering junk e-mail*. Retrieved November 29, 2004, from <http://research.microsoft.com/~horvitz/junkfilter.htm>
- Simon, H. A. (1982). *Designing organizations for an information-rich world: Models of bounded rationality* (pp. 171–185). Cambridge, MA: MIT Press.
- Spence, M. (1973). Job market signaling. *Quarterly Journal of Economics*, 87, 355–374.
- Webb, C. L. (2003, May 22). ‘Digital dreck’ times 180 million. *The Washington Post*. Retrieved May 25, 2003, from <http://www.washingtonpost.com/wp-dyn/articles/A25239-2003May22.html>
- Zandt, van T. (2001, May 18). *Information overload in a network of targeted communication*. Retrieved November 29, 2004, from <http://zandtwerk.insead.fr>