Yale ICF Working Paper No. 05-24
July, 2005

Pricing Electronic Mail To Solve the Problem of Spam

Robert E. Kraut, Carnegie Mellon University – David A Tepper School of Business and School of Computer Science

Shyam Sunder, Yale School of Management

Rahul Telang, Carnegie Mellon University – H. John Heinz III School of Public Policy and Management

James Morris, Carnegie Mellon University – School of Computer Science

Pricing Electronic Mail To Solve the Problem of Spam

Robert E. Kraut[1] **Shyam Sunder**[2] **Rahul Telang**[1] **James Morris**[1]

[1]Carnegie Mellon University

5000 Forbes Avenue

Pittsburgh, PA 15213

{Robert.Kraut|Rahul.Telang|James.Morris}@cmu.edu

+1 412 268-7694

[2]Yale University

PO Box 208200

135 Prospect Street

New Haven CT 06520-8200

Shyam.Sunder@Yale.edu

+1 203 432 6160

Abstract

Junk email or spam is rapidly choking off email as a reliable and efficient means of communication over the Internet. While the demand for human attention increases rapidly with the volume of information and communication, the 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 inform recipients about the value of a message they read it. This paper presents an economic model and the results of two laboratory experiments to explore the consequences of a pricing system for electronic mail. Charging postage for email causes senders to be more selective and to send fewer messages. However, the recipients did not use the postage paid by senders as a signal of message importance. These results suggest that markets for attention have potential for addressing the problem of spam, but that their design needs more work.

Keywords: Economics of spam, junk, pricing e-mail, market for attention

Journal of Economic Literature Codes: D40, D61, L50, L86, L96, O33

Pricing Electronic Mail To Solve the Problem of Spam

As Herbert Simon noted over two decades ago, with the increasing capability of computers to generate, store, and transmit information, "a wealth of information" can lead to a "poverty of attention" (Simon, 1982). At about the same time, Poole et al. (1984) empirically demonstrated that the supply of information has been growing faster than our ability to consume it. Over the past hundred years, the volume of words and images has grown exponentially, much faster than the number of recipients and the time they can devote to processing the information. The growth in information supplies means that a higher proportion of the information produced remains unread by many who could benefit from it. This 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–unsolicited, bulk electronic mail, known as spam.

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

Economics lies at the core of this spam glut. Computer technology has cut the cost of delivering messages by orders of magnitude. Relative to the fixed cost of hardware and software, the marginal cost of delivering an electronic mail message is negligible. The marginal cost of sending a marketing message to 1 million recipients by electronic mail is less than $2,000, while the same solicitation sent by conventional, bulk-rate postal mail would cost over $190,000 in postage, not counting paper and printing costs (Hansell, 2003). In the face of these low costs, it is economically rational for individual commercial emailers to distribute their messages as widely as possible. Experts estimate that commercial electronic mail is profitable if even one recipient in 100,000 makes a purchase (Hansell, 2003).

The present study explores whether an economic solution—pricing—can solve this problem caused by the economics of communication. A postage approach was one of the alternatives proposed at recent Senate hearings in May 2003 to deal with the problems of unsolicited commercial email (see Gross, 2003). In the introduction, we briefly review two other approaches for dealing with spam—filtering and regulation—and then summarize the basic idea of a pricing model. In Section 2, we develop an economic model of pricing. It demonstrates that, in principle, 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 3, we describe two laboratory experiments that test some of the core predictions of the economic model. The experiments show that per-message pricing does indeed improve targeting of messages. The final section of the paper considers how to translate the idealized conditions assumed by the economic model and the laboratory experiments into real policies.

Filters

Email filters and rules allow recipients to flag messages with a priority (Cranor, and LaMacchia, 1998). Filters are not effective for several reasons. First, they do not deter sending behavior unwanted by recipients. Commercial mailers can and do continually change or disguise the source of their messages in order to get through the filters. In a recent 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 hours (Krim, 2003). Second, effective filter rules are difficult to program and maintain. 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., 1989). Even modern, research-based Bayesian filters (Sahami, Dumais, Heckerman, & Eric Horvitz, 1998) are only 92 to 95 percent accurate with static email, and fail to account well for the continual evolution of strategies on the part of the senders of spam. Spammers disguise their identities or adopt new ones faster than we can devise more discriminating filters. Third, designing and deploying effective filters consume scarce recipient resources, getting past the filters consumes sender resources, and the messages themselves consume network and storage resources. Fourth, the recipients program filters without regard for the legitimate interests of the senders.

Regulation

Government regulations to place limits on commercial electronic mail exist in both Europe and the US. The European Union's Directive 2002/58/EC prohibits false identities on commercial email and requires that recipients must explicitly elect to receive commercial email before it can be sent to them. In the United States, at least 25 states currently have some form of commercial email legislation (see http://www.spamlaws.com/state/index.html for a current list). For example, in 2002, Utah legislation requires unsolicited commercial electronic mail to include the sender's name and physical address, along with opt-out instructions and accurate routing information. It requires these messages to contain a subject-line label indicating that the message contains advertising.

At the federal level, the Federal Trade Commission (2002) proposed a national "do not send registry" to restrict unsolicited telemarketing. More than a dozen spam-related bills have been introduced in Congress over the last two years. The Senate Commerce Committee unanimously approved the most promising bill in 2002, called the Controlling the Assault of Non-Solicited Pornography and Marketing Act (an awkward title chosen to yield the acronym Can Spam). It would require commercial emailers to truthfully identify themselves and valid "opt-out" options on all unsolicited commercial email. The FTC could impose civil fines up to $10 per unlawful message.

While it is possible for the regulatory approach to work, the highly distributed design of the Internet makes such regulation difficult to enforce. Even though the European Union has had more stringent email legislation for a longer time than the United States, recent field studies find that the patterns of junk email messages in the

U.S. and Europe remain similar (Jamal, Maier and Sunder, 2003a,b). One of the problems is that a large proportion of commercial electronic mail originates outside of the recipients' national borders. Even if the regulatory approach were to work within a single country, the global scope of the Internet renders enforcement of national laws across boundaries ineffective in the absence of an international enforcement regime.

Moreover, interest groups and direct marketers are lobbying to tone down the strongest aspects of anti-spam legislation. The groups argue that many of the bills would unfairly restrict email marketing and put electronic commerce at a disadvantage. According to American Direct Marketing Association (DMA), direct marketing is a $7.1 billion market in the US, and consumers save close to $1.5 billion due to such efforts (Grimes, 2003).

Pricing

Pricing is a third approach to dealing with the problem of spam. The pricing of email is an example of using a market mechanism to allocate scarce resources—human attention in this case. It is based on a fundamental principle of economics that market-determined prices can help allocate scarce resources in a Pareto-efficient manner. A Pareto-efficient allocation of resources among people is such that there does not exist an alternative scheme of allocating resources that makes anyone better off without also making one or more people 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 many other markets, efficiently allocate the resources through decentralized decision-making. It allows people to decide when to send or receive messages, based on the value of the messages they wish to send and receive and the value of their attention, without revealing private information.

A pricing system for email would conform to the spirit of the Internet in that it (1) would not depend on a big brother regulating the message traffic, (2) would not block anyone from sending a message to anyone else, (3) would allow electronic message pricing to evolve naturally in the market place, and (4) could be institutionalized so that it imposes no net cost for the users if the money paid by the senders was redistributed to the recipients of their messages.

Pricing electronic mail based on volume or importance is not a new idea. Researchers in academia and industry have explored how it might work (e.g., Dwork and Naor, 1993; Zandt, 2001; Dwork, Goldberg, & Naor, In press; Gross, 2003; Malone, Grant, Turbak, Brobst & Cohen, 1987). This is analogous to the practice of the U.S. Postal Service in charging different rates for express, registered, first class, second class, and bulk mail services. 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.

Zandt (2001) developed an economic model to explain why differential pricing of electronic mail should be helpful to both senders and recipients. 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. The 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. We extend Zandt's model 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 have information to differentiate among recipients in terms of their potential responsiveness to the message.

Zandt's and our formal models focus on the benefits of sender selectivity. Prices can also signal the sender's assessment of the message to the recipient and help the latter decide which of the competing messages deserve attention. The price that the sender paid is the major reason that recipients open their express mail before their bulk mail. Credibility of the signal should increase with the cost to the sender. A "high-priority" label on a message works only if the recipient has reason to believe, from prior experience or its cost to the sender, that such labels are not used indiscriminately.

The next section presents a model of the economics of email pricing. Section 3 derives hypotheses from the model and describes two experiments to test these hypotheses in the laboratory. Section 4 is a discussion, drawing implications and extensions from the model and experiments.

An Economic Model

Imagine a world with a large number of advertisers or senders who sell products that appeal to particular classes of consumers or recipients. If an advertising message reaches an interested recipient and the recipient responds to it, both the sender and recipient benefit. For example, if a nursery advertises a sale on perennials and its message reaches a gardening enthusiast who wants to find and buy a plant on sale, both the nursery and gardener benefit. However, both would lose if the nursery failed to send out advertisements, if the advertisement failed to reach the interested gardener, or if the gardener failed to process the message because it was buried in a flood of other messages.

To make this situation more general, suppose there are m senders and R recipients of messages[1]. Each sender is interested in only one of the recipients, although the sender does not know which one, and can send messages to any recipient. Each recipient may receive messages from any sender; the recipient is interested in messages originating from only one sender but she does not know which one. There is a potential overlap of benefits for senders and receivers. If the sender sends a message, which matches with the receiver's interest, then both realize some benefit.

The recipient's capacity to process (i.e., read, comprehend, and respond to) messages during the relevant time interval is limited and exogenously given as f; she ignores any messages received beyond this processing capacity. If the m senders send one message to each of the r recipients, each receiver will be receiving exactly m messages. Out of these m messages received, the probability of any given message being processed is $min(1, f/m)$.

Free Email (Baseline Results)

When email is free, the senders should send their messages to every possible recipient, as long as the gross payoff to the sender and the probability of the messages being processed by the recipients are positive.

---

[1] For ease of exposition, we use masculine pronouns for senders and feminine for recipients.

Of the m messages she receives, the recipient can process no more than f of them. If the payoff to the recipient from processing and responding to an interesting message is $Payoff_{\text{recipient}}$, her expected benefit is:

$$ExpectedBenefit_{recipient} = Payoff_{recipient} \times \min(1, f/m) \quad (1)$$

Since the sender has sent the message to every possible recipient, the person who would be interested in his message receives the message with certainty, and processes and responds to it with probability $min(1, f/m)$. Similarly, if the sender's payoff from getting a response to a message is $Payoff_{\text{sender}}$, his expected benefit from sending one message to each recipient is:

$$ExpectedBenefit_{sender} = Payoff_{sender} \times \min(1, f/m) \quad (2)$$

In the following sections we examine how different pricing policies may change senders' and recipients' behavior and their payoffs.

When No Targeting Is Possible

First, consider the case where the senders have no information to distinguish among the recipients, and recipients have no information to distinguish among the senders without opening their messages. Therefore, while the receivers are heterogeneous and would prefer to get only the "desired" message, senders cannot distinguish among them. Therefore, the nursery mentioned in our previous example does not know who is and is not interested in gardening; it may send the message to recipients not interested in gardening.

Flat-rate Pricing [2]

Under flat-rate pricing, senders pay a fixed amount ($Cost_{fixed}$) to send a message to an unlimited number of recipients. Suppose that each sender sends messages to *r* (less than or equal to *R*) recipients. Given the message processing capacity of the recipients (*f*), there is probability *r/m* that the message will reach an interested recipient, and probability *min(1,f/r)* that this recipient will have the ability to process the message. Clearly, the joint probability that a message will be responded to is simply *f/m*. More generally, an interested recipient will process a message when the probability is *min(r/m, f/m)*. If a recipient responds to the message, she gets some payoff but must pay postage of $Cost_{fixed}$ for responding to each sender of interest[3]. If $Payoff_{\text{recipient}}$ is less than $Cost_{fixed}$ the recipient would not respond, and get a net payoff of zero; otherwise, she would respond and get a net payoff of $Payoff_{\text{recipient}}$ - $Cost_{fixed}$. We will continue to assume that $Payoff_{\text{recipient}} > Cost_{fixed}$ in our subsequent discussion; otherwise there will be no market for emails. Thus, at the time a message is received, the recipient expects a net payoff from the message given by:

$$ExpectedBenefit_{recipient} = \min(r/m, f/m)(Payoff_{recipient} - Cost_{\text{fixed}}) \quad (3)$$

[2] Non-linear pricing strategies in different contexts have been studied extensively in economic literature (Wilson 1993)

[3] One could argue that receivers need not pay for sending messages if the main goal of pricing is to restrict the commercial bulk messages sent. Adding this generality simply shifts the profit function, but does not affect the nature of our results. More generally, an informationally decentralized postage regime cannot depend on the central availability of information which is inherently priv

Now, the sender has probability $min(r/m, f/m)$ of getting a response to one of the r messages he sends out and obtaining a gross payoff of $Payoff_{\text{sender}}$. The sender's expected payoff from sending r messages is

$$ExpectedBenefit_{sender} = Payoff_{sender} \bullet \min(r/m, f/m) - Cost_{\text{fixed}} \qquad (4)$$

The sender will send a message only if this net payoff is non-negative. The expected benefit maximizing strategy for the sender requires him to find the maximum probability that a message will be processed by an interested recipient, $min\ (r/m, f/m)$. This is achieved by setting r to its maximum possible value R and yielding probability $min(1, f/m)$ of getting a response. Since the marginal cost of sending an extra message is zero, it is always optimal for a sender to send messages to all receivers. Thus the sender will send messages to all receivers if $Payoff_{\text{sender}}.min(1,f/m) - Cost_{\text{fixed}}$ is non-negative. This strategy will yield an expected benefit for the sender $Payoff_{\text{sender}}.min(1,f/m) - Cost_{\text{fixed}}$ when messages are sent, and zero otherwise. The same optimal strategy will yield an expected benefit of $min(1,f/m).(Payoff_{\text{recipient}} - Cost_{\text{fixed}})$ for the recipients senders choose to send messages to, and zero otherwise.

Note the following decision rules:

If $Payoff_{\text{sender}}.\ min(1,f/m) - Cost_{\text{fixed}} >= 0$, senders send one message to every recipient and fraction $min(1,f/m)$ of these messages yield responses, giving senders a payoff of $Payoff_{\text{sender}}\ .\ min(1,f/m) - Cost_{\text{fixed}}$, and recipients a payoff of $min(1,f/m).(Payoff_{\text{recipient}} - Cost_{\text{fixed}})$.

If $Payoff_{\text{sender}}\ .\ min(1,f/m) - Cost_{\text{fixed}} < 0$, then all senders cannot send m messages and hope to be profitable. In this case, some senders will drop out of the market but the rest will continue to send the messages to all *m* receivers, because the marginal cost is zero.

Interestingly, even if the number of messages drop, the receiver's benefits are unlikely to improve. Note that that if the recipients do respond, they would always respond to the same fraction of messages *min(1,f/m)*). Therefore, as the number of messages decline, the probability of getting the message from the right sender (*r/m*) will decrease, canceling out the increase in probability of processing (*f/s*).

As the cost of communication $Cost_{\text{fixed}}$ rises, the range of parameters over which the recipients find it worth responding to the messages shrinks, reducing the processing of messages. Importantly, this fraction of messages recipients process is not influenced by the cost of communication. Therefore, while a higher cost of communication can change the number of messages received, it will not affect the receiver's benefits. Similarly, the overall benefits to senders will not change when these costs increase. Therefore, when communication does take place, the cost of communication reduces the expected profits of both the senders and the recipients.

Usage-Based Pricing

Now consider the case when senders have to pay for each message sent to each recipient. We assume that there is a communication cost of $c_{per\ message} > 0$ per message, and senders pay $c \times r$ to send a message to *r* different recipients. The benefit functions (3) and (4) still apply; for calculating profits we simply subtract cost $c_{permessage} \times r$ instead of $Cost_{\text{fixed}}$.

Therefore, at the time a message is received, the recipient expects a net payoff from the message given by:

$$ExpectedBenefit_{recipient} = \min(r/m, f/m)(Payoff_{recipient} - c_{permessage}) \qquad (5)$$

Similarly, with probability *min(r/m, f/m)*, the sender gets a response to one of the *r* messages he sends out and obtains a gross payoff of $Payoff_{\text{sender}}$. The sender's net payoff from sending *r* messages is $Payoff_{\text{sender}}.min(r/m,f/m) - rc_{\text{per message}}$.

If it was profitable to send messages to all *m* receivers, i.e., $Payoff_{\text{sender}}.min(1,f/m) - mc_{\text{per message}} > 0$, then the sender would send messages to every one of the *m* recipients. When the cost of communication rises, eventually it will be unprofitable to send m messages to a receiver. In this case, the number of messages sent to that receiver reduces from *m*. Therefore, senders start rationing the messages, or one or more of the senders stop sending messages. So the decision rule is:

If $Payoff_{\text{sender}}.\ min(1,f/m)- mc >= 0,$ senders send one message to every recipient and fraction *min(1,f/m)* of these messages yield responses, giving senders a payoff of $Payoff_{\text{sender}} .\ min(1,f/m) - mc,$ and recipient a payoff of $min(1,f/m).(\ Payoff_{\text{recipient}} - c)$.

If $Payoff_{\text{sender}} .\ min(1,f/m) - mc < 0,$ then all senders cannot send *m* messages and hope to be profitable. In this case, either some senders drop out of the market or all senders randomly send fewer messages ($r < m$).

Since receivers will respond to the fraction *min(*1,*f/m)* of the messages and even usage-based pricing will not affect this response, the results derived from flat-rate pricing will continue to hold when applied to usage-based pricing. Again, since the fraction (*min(*1,*f/m*), is unchanged, the receiver and sender benefits remain unchanged as well.

It should be noted that depending on $c_{per\ message}$ and C_{fixed}, the number of senders sending messages and hence the total volume of communication will be different in both cases, but the overall benefits (without costs) to all senders and receivers will be unchanged. More precisely, the nature of the equilibrium remains same.

To summarize, overall benefits in flat-rate and usage-based pricing for all senders and recipients are the same. Note that the realized profits may be different, depending on the cost of communication. Moreover, when senders have no differentiating information about the recipients, the total benefits for senders and recipients are not affected by the change in prices under either pricing scheme. However, under both pricing schemes the total communication declines with an increase in price.

Targeting Recipients

Now consider the case where recipients are heterogeneous and the senders can distinguish among them in terms of their likely value. Hence, senders can partially rank-order recipients in terms of the relevance of a message to them and thus the likely benefit each will receive from processing the message. With usage-based pricing, the sender will target messages, sending them only to those recipients for whom the benefit of sending the message is more than its cost. The key in the case of targeting is the response function of the receiver. When a receiver receives r messages, she receives them from the top r most relevant senders. Therefore, the joint probability of one of the r messages being responded to is now $Relevant/m.min(1, f/r)$. Here $Relevant/m$ is the overall probability that one of the r senders is the "desired" sender. $Min(1, f/r)$ is simply the probability that the message will be noticed and processed by the recipient. More generally, the response function is $min(Relevant/m, R/m\ f/r)$. The important thing to note is that $Relevant > r$ because of partial rank ordering.

The receiver benefit function can be written as

$$ExpectedBenefit_{recipient} = \min(Relevant/m, \frac{Relevant}{m}\frac{f}{r})(Payoff_{recipient} - c_{permessage}) \quad (6)$$

If a sender sends a message to the top r receivers then its benefit function is

$$ExpectedBenefit_{sender} = Payoff_{sender} \bullet \min(Relevant/m, \frac{Relevant}{m}\frac{f}{r}) - rc_{permessage} \quad (7)$$

Therefore, if sending a message to all m receivers was profitable— i.e. $Payoff_{sender}.\ min(1,f/m) - mc_{per\ message} > 0$— then the sender sends a message to everyone and there is no need for targeting. Note that in this case $Relevant = r = m$ and hence the receiver response function is simply $min(1,f/m)$. However, when the cost of communication is such that sending all m messages is not profitable, then the lowest value sender drops out first. Therefore, now if only r messages are profitable, then the receiver response function as in (6) is $\min(\frac{Relevant}{m}, \frac{Relevant}{m}\frac{f}{r})$. Note that because of targeting $R/r > 1$, and the receiver benefit increases with higher costs of communication. It is immediately clear from (7) that the sender's benefits increase as well. The decision rule is

If $Payoff_{sender}.\ min(1,f/m) - mc_{per\ message} >= 0$, senders send one message to every recipient and a fraction $min(1,f/m)$ of these messages yield responses, giving senders a payoff of $Payoff_{sender}\ .\ min(1,f/m) - m\ c_{per\ message}$, and recipients a payoff of $min(1,f/m).(Payoff_{recipient} - c_{per\ message})$.

If $Payoff_{sender}\ .\ min(1,f/m) - m\ c_{per\ message} < 0$, then all senders cannot send m messages and hope to be profitable. In this case, the senders send r messages to their top receivers and realize the payoff of $Payoff_{\text{sender}} \min(\frac{Relevant}{m}, \frac{Relevant}{m}\frac{f}{r}) c_{per\ message}\ .r$. The

receivers realize benefits of $\min(\frac{Relevant}{m}, \frac{Relevant}{m}\frac{f}{r})\ (Payoff_{\text{recipient}} - c_{per\ message})$.

However, with flat-rate pricing, whenever a sender sends a message, he sends it to everyone, even to its lowest valued receiver. Therefore, targeting in the flat-rate case is not likely to increase benefits. In summary, it is the sender's ability to target recipients, combined with the usage-based pricing scheme, which leads to higher benefits to senders and recipients.

It should be noted, however, that while usage-based pricing reduces the amount of communication and information overload, charging too high a price could lead to sub-optimal benefits for both senders and recipients. Once the number of messages received is equal to *f*, raising the cost will not lead to more benefits. Too high a cost can reduce the number of messages below *f*, hence reducing the benefit for both sender and recipient. On the other hand, if the cost is too low, it may not reduce the number of messages enough to prevent information overload.

Signaling

In many situations, a sender can signal the importance of his message by paying more. For example, people generally open their express mail before opening their bulk mail because the sender spent more to send the message.

The goal of signaling is to raise the probability that a recipient will notice a message. Suppose a recipient is receiving r-1 messages. The r^{th} sender has two options to choose from. Either he can send a standard message at price *c* or he can send a priority message at price *c*'. We also assume that r^{th} sender has some targeting information available about the recipient, and believes that there is a probability R that this receiver is interested in his message[4]. In the case of a standard message, the recipient response function is $R.min(1, f/r)$. R is simply the probability that the r^{th} sender is the "desired" sender and $min(1, f/r)$ is simply the probability of the message being noticed and processed by the receiver. Therefore, the receiver's payoff for this one message is

$$ExpectedBenefit_{recipient} = R \bullet \min(1, f / r)(Payoff_{recipient} - c)$$

(8)

Similarly, sender's payoff for sending this one message is

$$ExpectedBenefit_{sender} = Payoff_{sender} \bullet R \bullet \min(1, \frac{f}{r}) - c$$

(9)

Now consider a case where the sender can send priority messages. Priority messages can increase the probability of message being noticed to $\min(1, \frac{f(1+S)}{r})$ where S is the signaling power of the message. Following the same analysis, the expected benefit of the sender sending a priority message is

$$ExpectedBenefit_{sender} = Payoff_{sender} \bullet R \bullet \min(1, \frac{f(1+S)}{r}) - c'$$

(10)

The sender will send the priority message only if (10) is greater than (9). Or,

---

[4] If there was no targeting information available then this probability R is simply $1/m$.

$$Payoff_{sender} \bullet R \bullet \min(1, \frac{f(1+S)}{r}) - c' > Payoff_{sender} \bullet R \bullet \min(1, \frac{f}{r}) - c$$

$$Payoff_{sender} \bullet R \bullet \left(\min(1, \frac{f(1+S)}{r}) - \min(1, \frac{f}{r}) \right) > c' - c$$

(11)

If there is no information overload, i.e. $f > r$, then the sender will never use priority messages because the standard message itself is noticed by the receiver and there is no need to signal. When $f < r$, the expression (11) simplifies to

$$Payoff_{sender} \bullet R \bullet \left(\frac{f(1+S)}{r}) - \frac{f}{r}) \right) > c' - c$$

Or,

$$Payoff_{sender} \bullet R \bullet \frac{f\,S}{r} > c' - c$$

Notice that the sender will send priority messages if:

The difference in cost of priority message and standard message is low.

The signaling power *S* of priority message is high.

The sender has targeting information available about the receiver (higher *R*).

Design of Experiments

We conducted two laboratory experiments to test hypotheses derived from the economic model, by examining the effects of different postage regimes on message-sending behavior, message-receiving behavior, and the overall social welfare of the participants. As in most experiments, our results only show what the effects of an intervention—in this case, per message pricing and targeting—*can be* under controlled circumstances. Producing the same effects in the real world is a further 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 communication as high-priority or standard-priority. They performed 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 condition. Under the second, they paid per recipient and paid more for sending high-priority than for standard-priority mail. We provide the details below.

The experiments simulated a situation with the following conditions:

(A) Players had a foreground task—completing crossword puzzles—whose performance was personally valuable, but whose value varied with personal endowment and goals. This foreground task determined the opportunity costs of the time players spent processing email.

(B) Players had an incentive to send electronic mail to other players, which may or may not have been of value to the recipients.

(C) 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 communication tasks, and among the messages they received.

(D) Players differed in terms of the types of messages they found valuable.

Experiment 1

Subjects and task.

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

Subjects could send email messages to one or more players to request help or to provide help. To simplify data analysis, the players were required to characterize the content of their messages and were allowed to include only a single query or answer per message. Upon receiving a message purporting to contain help, the recipient indicated whether he or she would use the answer. If so, the sender was awarded $.15 per letter in the answer.

Although exchanging messages was potentially rewarding, it also competed for the time a subject could work on the puzzle, which was an alternative way to earn money. To emphasize the opportunity cost of communicating, we imposed a time delay when players read their mail. After clicking on a message, the player had to wait five seconds before the text of the message was displayed. During this time they could not work on their puzzles.

The subjects could send messages as either standard-priority or high-priority. Standard-priority messages were delivered in approximately 20 seconds and appeared in a player's inbox in standard text. High-priority messages were delivered immediately and appeared in the recipient's inbox in bold text.

To increase the burden of processing messages, a server sent out email approximately every seven seconds. The server-generated messages appeared to come from another player. The server-generated messages were all sent via standard-priority.

None of the messages contained a subject field. As a result, they revealed few cues to their value before being opened. To determine whether a message was potentially valuable, the recipients had to open it, relying on its priority level, or on previous experience with the sender. The 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 subjects 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*. We calculate the postage spent to send a high-priority message sent to three people to illustrate the experimental conditions. In the fixed-rate condition, this message would cost $.02 ($.02 x 1 message), but it would cost $.12 ($.04 x 3 addressees) in the variable-rate condition.

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

The customized email system.

We built an experimental environment, including a customized email system, for exchanging messages and keeping track of costs and earnings. Four windows appeared on the subjects' screens (See Figure 1). One of these (top left) was a customized version of Excel™ containing a crossword puzzle and its clues. Two windows were used for email, one for sending messages (bottom middle) and one for viewing messages (right). The fourth window (bottom left) was used for calculating the player's cash earnings. The values in the earnings window were updated every thirty seconds to reflect rewards from helping other players and postage charges. At the end of each round, a server computer graded the crossword puzzles and updated players' earnings.

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

Because messages were nested within players within rounds of the experiment, and because players were nested within the experimental session, we used hierarchical linear models to account for the non-independence 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.

Measures.

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

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:

The hypotheses derived from the economic model are divided into three categories. H1 refers to sending behaviors, H2 to attentional behaviors, and H3 to overall performance of the system. All of the hypotheses compare behavior under expensive, variable-rate, recipient-based pricing compared to inexpensive, fixed-rate, message-based pricing.

Sending behavior:

H1a: Under variable-rate pricing compared to fixed-rate pricing, the senders send fewer messages.

H1b: Under variable-rate pricing compared to fixed-rate pricing, the senders send messages with fewer recipients per message.

H1c: Under variable-rate pricing compared to fixed-rate pricing, the senders use the high-priority designation less often.

Attentional behavior:

H2a: Under variable-rate pricing compared to fixed-rate pricing, the recipients read a larger fraction of messages.

H2b: Under variable-rate pricing compared to fixed-rate pricing, the recipients will read a larger fraction of high-priority messages compared to low-priority messages.

H2c: Under variable-rate pricing compared to fixed-rate pricing, the recipients reply to a larger fraction of messages.

H2d: Under variable-rate pricing compared to fixed-rate pricing, the recipients accept a larger fraction of help offered.

Overall performance:

H3a: Under variable-rate pricing compared to fixed-rate pricing, the players earn more from completing puzzles completion.

H3b: Under variable-rate pricing compared to fixed-rate pricing, the players earn more from offering help.

Results and Discussion

The subjects used high-priority frequently, in 77 percent of messages. Across all conditions, they spent $.91 per round on postage. Forty-five percent of their messages were questions and 52 percent were answers. Sixty-four percent of their messages were read. Fifteen percent (primarily questions) received responses.

Effects of postage regime on message-sending.

The postage regime influenced message-sending in economically sensible ways. Table 1 shows these effects. In the variable-rate postage condition (paying per recipient and for high-priority messages), the players rationed their messages. They sent a third fewer messages per round ($p < 0.13$), addressed each to fewer recipients ($p < 0.01$) and used the high-priority delivery option less often ($p < 0.01$). In the fixed-price case, when the cost of a high-priority message was the same as that of a standard message, the users predominantly sent high-priority messages. The results confirm hypotheses H1a, H1b, and H1c.

Effects of message-sending behavior on attention

The way the players sent messages had powerful consequences for the attention that the addressees paid to them. In the variable-rate condition, the players read a higher percentage of the messages ($p < 0.01$) and also replied to more of them ($p < 0.06$), confirming hypotheses H2a and H2c. They could pay more attention to the fewer messages received under the variable-rate postage. The players read a higher proportion of messages when fewer originals were sent ($r = -0.14$, $p < 0.02$) and when each original was sent to fewer recipients ($r = -0.31$, $p < 0.001$). The recipients did not accept more help in the variable-rating condition ($p = 0.60$), providing no support for hypothesis H2d.

The players read high-priority messages more often than standard-priority ones. They read 75 percent of the high-priority messages versus only 12 percent of the standard-priority messages ($r = 0.72$, $p < 0.0001$). These large differences in reading rates occurred partly because recipients could be sure that the high-priority messages were sent by real players and not the server.

We had expected high-priority messages would be read most when this designation was costly (i.e., in the variable-rate postage condition). The message-reading frequencies in Figure 3 are consistent with this prediction, although the postage regime X priority interaction did not achieve statistical significance ($F(1,1315)=2.22$, $p< 0.13$). Thus, hypothesis H2b was only weakly supported.

Effects of postage regime on economic outcomes

We expected that these effects of the variable-rate postage regime—reducing the volume of communication and increasing readership—would help players more efficiently allocate their time between their foreground task (the crossword puzzles) and communication. We expected them to earn more money under the variable postage condition, at least before subtracting postage costs.

This prediction was disconfirmed. The subjects in the variable-rate postage condition netted significantly less than those in the fixed-rate postage condition. They earned insignificantly less for completing puzzles, but spent substantially more for sending messages (see Table 1).

While the result is somewhat puzzling, our economic model provides an insight. The users did not have targeting information in this experiment. Therefore, when the cost of communication went up, they rationed their messages as expected. But when they reduced the number of messages, their probability of reaching a "desired receiver" was also reduced. This essentially cancelled out the gains of lower information overload. We use the number from Table 1 to calculate the sender's profits in two different pricing regimes.

The response rate in the fixed-rate regime was 12 percent. A sender sends 19.2 unique messages. Therefore, on average, only $0.12x19.2 = 2.3$ messages are replied to. If the sender makes *p* dollars for each message replied to then from (1) the benefit to the sender is:

$$benefit_{sender} = p_{sender} x2.3$$
$$= 2.3\, p_{sender}$$

The response rate in the variable-rate regime was 17 percent. A sender sends 12.5 unique messages. Therefore only $0.17x12.5 = 2.12$ messages are replied to. Therefore, the benefit for the sender is

$$benefit_{sender} = 2.12 p_{sender}$$

Under these parameters, the benefits do not change substantially. We performed similar calculations for the receivers and reached the same conclusions. Clearly, simply enforcing the pricing regime (either fixed or variable) without providing targeting information is unlikely to change social welfare.

Experiment 2

Experiment 2 was designed to simulate a condition where the senders of messages could target the recipients. In Experiment 1, each message had the potential to benefit the senders and the recipients equally. People who asked questions benefited if they received good answers, and players who sent answers received rewards if their answers were used. We introduced variation in targeting by assigning players domains of expertise and testing to see whether the players were more likely to address their questions to the experts under costly communication. We contrasted the potential to target questions with the rules associated with a new type of message—advertisements. The senders were rewarded when the recipients read advertisements, but the recipients were penalized because opening advertisement messages left them less time to do their crossword puzzles. The senders could not target their advertisements, because they benefited equally regardless of the recipient who opened it.

In summary, Experiment 2 was designed to test the hypothesis that pricing communication would change the senders' behavior most when the recipients were heterogeneous. We did this by creating two classes of messages. The recipients were heterogeneous with regard to *questions*, because the senders benefited more when they were sent to a domain expert. However, the recipients were homogeneous with regard to *advertisements*; the senders benefited equally, regardless of which recipients read their advertisements.

Methods

The procedures for Experiment 2 were very similar to Experiment 1, with the following differences.

Advertising

All subjects could send questions and answers, as in Experiment 1. In addition, two subjects in each experimental session could send advertising messages, earning $.25 per recipient who opened an advertisement. Opening an advertisement consumed five seconds of the recipient's time.

Expertise

As in Experiment 1, the answers to the puzzles were distributed among the players. In Experiment 2, the answers were distributed so that each player was an "expert" in one domain. The expert was 80 percent likely to have an answer versus 50 percent for a random player. All the players were given a table showing the distribution of expertise across players.

Testable Hypotheses:

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

Sending behaviors

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

H1d: 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.

H1e: 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. Two additional hypotheses are applicable:

H2e: Recipients are more likely to read advertising messages with usage-based than with message-based pricing.

H2f: Recipients are less likely to open and read advertising messages than non-advertising messages.

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

Results and Discussion

In Experiment 2, the subjects had targeting information available when asking questions but not when sending advertising messages. The subjects used high-priority for 56 percent of the messages. Across all conditions, they spent $.64 per round on postage. Thirty-three percent were questions, 45 percent were answers, and 23 percent were advertisements. Recipients read 55 percent of the messages and replied to 45 percent of the questions.

Effects of postage regime on message sending

Costly communication influenced the participants' sending behavior in economically sensible ways. Under the variable-rate postage, the participants sent fewer messages, addressed each message to fewer participants, and used high-priority less often. In addition, as predicted, the subjects in the variable-rate postage condition were more likely to target their queries to the experts (see Table 2).

The players sent most of their messages to multiple recipients, and did so more for advertisements (93 percent) than for questions (74 percent) or answers (75 percent). To test the prediction of our model that variable-rate pricing along with targeting would reduce the number of recipients per message, we performed the following test. In this experiment, the players could send advertisements, questions, or answers. The players had no targeting information when they sent advertisement and answers, but they did for questions. We used the interaction of message type with the postage regime to predict the number of message sent. Consistent with H1d, variable-rate postage inhibited sending to multiple recipients most when the senders could differentiate among the recipients. We plot this in Figure 3, showing the percent of potential recipients sent to for the three message types 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 and answers. Since no targeting information was available for advertisements and answers, the postage regime did not affect them significantly.

We also tested the interaction of high-priority messaging with targeting. We argued in our model that when targeting information is available, users tend to send high-priority messages as opposed to standard messages in the variable-postage regime. Targeting information was available only for questions. About 42 percent of the questions were sent using high-priority messaging in the variable regime, as compared to only 15 percent of the advertisements. But about 41 percent of the answers were also sent using high-priority, even though no targeting information was available.

Effects of postage regime on attention

While the recipients read and replied to slightly more of their messages in the variable-rate postage condition, neither effect approached statistical significance. The players were more likely to read high-priority messages (59 percent read high-priority messages read versus 41 percent who read standard-priority messages), highlighting the importance of signaling. Contrary to expectations (H2b), variable-rate postage did not enhance the signaling power of the high-priority option ($p > 0.50$). The response rate was critically dependant on whether the right receiver got the message. Since the queries could be targeted, we expected that queries with a high-priority signal would have a better chance of being replied to. In the variable-postage regime, 30 percent of the high-priority queries were replied to as opposed to 22 percent of the standard messages.

The participants were 18 percent less likely to read advertising messages than those containing questions or answers ($p < 0.001$). The players could not know the type of message before incurring the cost of opening it; the senders' identity was their only cue to content. The recipients 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 b were disconfirmed. The subjects in variable-rate postage sessions earned reliably less than those in fixed-rate sessions. They earned less from completing the puzzles and paid more for postage (see Table 2).

Although our economic model argues that the total welfare will increase under variable-rate pricing when targeting is possible, we do 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 under variable-rate pricing. Second, probably the high-priority messages were too costly. As we noted above, a high-priority query was likely to be replied to 30 percent of the time, leading to certain benefits to the sender, while a standard query was likely to be replied to 22 percent of the time. If we simply assume that each message replied to generated p dollars to the sender, then for the high-priority messaging to work, this had to be true:

$$0.3\,p - 0.04 > 0.22\,p - 0.02$$

Or, $0.08\,p > 0.02$

But p was at most \$0.15 in our case. Therefore, the senders could not realize the benefit of the high-priority messages even in the variable-rate regime.

Discussion

To summarize, this paper developed an economic model of the influence of pricing and targeting electronic mail on the behavior of email senders and recipients and of the benefit they gained. The major conclusions are that under the current regime, with practically free electronic 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, sub-optimal for both advertisers and their targets. Because the recipients can read, understand, and respond to only a fraction of the mail they receive, and because they can't adequately distinguish useful from worthless mail without processing it to at least some degree, they 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 the 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 email does not change this situation and does not increase the benefits to either the senders or the recipients. Under a fixed fee, some advertisers stop sending mail altogether, but those who find it profitable to send mail should still send it to all possible recipients. One may intuitively believe that reducing the volume of communication per se would lead to better payoffs for the recipients. Indeed, many proposals to curb spam are based on this premise. However, our model points out that this is not the case. The consumers' loss from the failure to receive relevant messages cancels out the benefits from wasting less time on irrelevant ones. The benefits to the senders and the recipients increase only when the senders can distinguish the interested from the uninterested consumers, and when they incur a cost per recipient for sending messages. Under these conditions, the senders are motivated to send only to interested recipients, increasing their own and the recipients' benefits.

The paper also reports the results from two experiments testing hypotheses based on the economic model. The experiments examined the consequences of making electronic communication more costly, by charging per recipient, and imposing a surcharge for high-priority messages. In both experiments, usage-based postage had consequences that follow from the model. Under conditions of variable-rate pricing, senders sent fewer messages, sent each message to fewer recipients, and used the high-priority option less often. Experiment 2 tested the key 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 per-recipient cost for communication. This prediction was confirmed. The senders were more likely to target their messages to relevant recipients under variable-rate postage, when they had information to differentiate among the recipients.

The evidence is strong that the message-sending behavior influenced the frequency of recipients reading and responding behavior. The players read a higher proportion of messages when they had fewer to read, and when they were directed to fewer addressees. In Experiment 1 (but not in Experiment 2) players read and responded to a higher proportion of messages under variable-rate pricing. In both experiments they read and responded to more messages labeled high-priority than standard-priority. However, in neither experiment did the cost of high-priority messages change their signaling power.

These changes in the economics of communication made it more efficient, reducing volume and increasing relevance. We had expected that the usage-based pricing would benefit both the senders and the recipients if the senders could target messages. Our experiments did in fact allow for targeting, and the senders targeted more under usage-based pricing. However, the benefits to senders and recipients declined because the senders sent too few messages. We suspect that the overall social welfare did not increase under usage-based pricing because we chose parameters that made communication too expensive 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 the sender and the recipient benefit from relevant communication. Our model points out that both the sender's 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 upon the diagnosticity of the cues senders use to address their messages. Empirical studies that vary the diagnosticity of the information as well as targeting and pricing 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. Our economic model indicates that with the design of an appropriate pricing scheme, both the senders and the recipients could benefit from the pricing of electronic mail. One hurdle in implementing such a vision is convincing people to accept having to pay for what is currently a free service. The transition from a free to a paid service has been accomplished in other domains, including shifting from flat-rate pricing 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.

Similar shifts are already occurring in the domain of electronic mail. In 2002, the Daum Corporation, the largest Internet portal in Korea, created an Online Stamp Service, to charge bulk emailers a fee to send messages to its subscribers (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 email (approximately .08 cents, US) charged to those who send the most mail. The mailers receive a rebate, however, proportional to the fraction of subscribers who rate their bulk email as "informative." The subscribers receive points redeemable for gifts for each email they rate.

The Online Stamp service became operational April 1, 2002. Figure 4 shows the daily volume of bulk email sent to Daum subscribers before and after the initiation of the

Online Stamp service. Bulk email traffic in the three months following the introduction of the service was 46 percent of the volume in the prior three months. Company officials report that the quality of commercial email also changed following the introduction of the Online Stamp service, to become more informative. The company collected no systematic data, however, about whether the introduction of the Online Stamp system changed the likelihood of emailers targeting particular subscribers, or the likelihood of subscribers reading commercial email or responding to it. As a result, there is no evidence about whether the reduction in bulk email changed the benefits for either the senders or the recipients.

In the pricing scheme used in the experiments, the 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 the recipients of messages rather than to a bank either on delivery or when read. This latter pricing scheme would not only reduce the volume of messages and improve targeting, but might also induce the recipients to read messages that the senders consider important. In this sense, the 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 addition, the information about the postage attached to the message could be made available to the recipient before she decides whether to open the message.

Although our models and experiments have treated postage for electronic 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. Our analysis and empirical results apply to these cashless pricing mechanisms.

The pricing scheme we describe in this paper 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 (Horvitz et al, 2003) and Hudson (Hudson et al, 2003) have developed algorithms to assess an individual's interruptability. Postage could vary with a recipient's interruptability.

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.

ACKNOWLEDGMENTS

The Intel Corporation supported this project. 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.

References

1. Back, A. (2002, August 1). Hardcash: A denial of service counter-measure. Retrieved June 1, 2003, from the World Wide Web: http://www.cypherspace.org/~adam/hashcash/hashcash.pdf

2. Boyle, C. D. B., & Clarke, M. R. B. (1985). An intelligent mail filter. Proceedings of the HCI'85 Conference on People and Computers: Designing the Interface (331-341).

3. Bryk, A. & Raudenbush, S. (1992). Hierarchical linear models: Applications and data analysis methods. Newbury Park, CA: Sage.

4. Cranor, L. F. and LaMacchia, B. A. (1998). Spam! Communications of the ACM. 41(8), 74-83.

5. Dwork, C. and 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, (pp139-147), NY: Springer-Verlag.

6. Dwork, C., A. Goldberg, & Naor, M. (In press). On memory-bound functions for fighting spam. Proceedings of CRYPTO 2003.

7. Federal Trade Commission. (2002). Telemarketing Sales Rule, 16 CFR Part 310. Washington, DC: U.S. Printing Office.

8. Grimes, Anne. (2003). "Digits: Spam pays," The Wall Street Journal, May 22, 2003, page B3.

9. Gross, Grant. (2003). Spam hearing: E-mail tax, international treaty proposed. InfoWorld, May 21, 2003. http://www.infoworld.com/article/03/05/21/HNspamtax_1.html.

10.Hansell, Saul. (2003). Internet is losing ground in battle against spam. The New York Times. Apr 22, 2003, p 1.

11.Hayek, Friedrich A. (1945). The uses of knowledge in society. American Economic Review, 35(4), 519-530.

12.Horvitz, E., Jacobs, A., & Hovel, D. (1999). Attention-sensitive alerting. Proceedings of UAI '99, Conference on Uncertainty and Artificial Intelligence, (pp 305-313). Stockholm, Sweden.

13.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.

14.Hudson, S. E., Fogarty, J., Atkeson, C. G., Avrahami, D., Forlizzi, J., Kiesler, S., Lee, J. C., & Yang, J. (2003). Predicting human interruptibility with sensors: A Wizard of Oz feasibility study. Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI 2003). New York: ACM Press.

15.Jamal, Karim, Michael Maier, & Sunder, S. (2003b). Effects of standards and enforcement on disclosure and business practices: A comparative study of e-commerce privacy in US and UK. Yale University Working Paper.

16.Jamal. K., Maier, M. & Sunder S. (2003). Privacy in e-commerce: Development of reporting standards, disclosure and assurance services in an unregulated market. (2003). Journal of Accounting Research 41 (2), 285-309.

17.Kraut, R. E. & Attewell, P. (1997). Media use in a global corporation: Electronic mail and organizational knowledge. In S. Kiesler (Ed.) Culture of the Internet. (pp. 323-342). Mahwah, NJ: Erlbaum.

18.Krim, Jonathan. 2003. "A spammer speaks out: In Hill testimony, bulk e-mailer says Internet providers use same tactics," Washington Post, May 22, 2003, p. A01.

19.Mackay, Wendy E. (1988). Diversity in the use of electronic mail: A preliminary inquiry. ACM Transactions on Office Information Systems,6(4), 380-397

20.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(2), 115-131.

21.Malone, T.W., Grant, K.R., Turbak, F.A., Brobst, S.A. and Cohen, M.D., (1987). Intelligent information-sharing systems. Communications of the ACM, 30(5), 390-402.

22.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, Massachusetts: MIT Press.

23.Perlow, L. A. (1999). The time famine: Toward a sociology of work time. Administrative Science Quarterly, 44(1), 57-81.

24.Pollock, Stephen (1988). A rule-based message filtering system. ACM Transactions on Office Information Systems, 6(3), 232-254.

25.Pool, I., Inose, H., Takasaki, N. & Hurwitz, R. (1984). Communication flows: A census in the United States and Japan. New York: North-Holland.

26. Sahami, M., Dumais, S., Heckerman, D., & Horvitz, E. (1998, July). On Bayesian spam filtering: A Bayesian approach to filtering junk email. Paper presented at the AAAI Workshop on Learning for Text Categorization, Madison, WI.

27. Simon, H. A. (1982). Designing organizations for an information-rich world: Models of Bounded Rationality (pp. 171-185). Cambridge, MA: MIT Press.

28. Telang, Rahul. (2002). Information overload and pricing of communication services. Carnegie Mellon University Working Paper.

29. United States Government Printing Office (1998). United States Code, Title 39 POSTAL SERVICE. Prohibition of pandering advertisements

30. Webb, Cynthia L. 2003. "'Digital dreck' times 180 million," Washington Post, May 22, 2003. http://www.washingtonpost.com/wp-dyn/articles/A25239-2003May22.html.

31. Wilson, R.B. (1993). Nonlinear pricing. New York: Oxford University Press.

32. Zandt, van T. (2001, May 18). Information overload in a network of targeted communication. Downloaded from http://zandtwerk.insead.fr, September 14, 2001.

Table 1: Experiment 1. Effects of postage regime

Dependent variable	Postage Regime		p
	Fixed	Variable	
Sending behavior			
Unique messages sent [b]	19.2	12.5	*0.13*
Recipients per message [a]	3.0	2.4	*0.01*
percent high-priority message [a]	90percent	68percent	*0.01*
Attention behavior			
percent messages read [a]	57percent	77percent	*0.01*
percent messages replied to [a]	12percent	17percent	*0.06*
percent help accepted [a]	10percent	17percent	*0.60*
Performance			
Total earnings [b]	$2.64	$1.71	*0.05*
Puzzle completion earnings [b]	$2.26	$2.08	*0.45*
Reward for help [b]	$.79	$0.85	*0.47*
Paid for postage [b]	$.40	$1.21	*0.10*

(N=1474 messages nested within 5 rounds per player,

nested within 55 players nested within 11 experimental sessions)

[a] Based on a hierarchical linear model, with the message as the unit of analysis.
[b] Based on a hierarchical linear model, with the player nested within round as the unit of analysis

Table 2: Experiment 2. Effects of postage regime

Dependent variable	Postage Regime		p
	Fixed	Variable	
Sending behavior			
Unique messages sent [b]	13.6	9.8	*0.01*
Recipients per message [a]	3.6	3.0	*0.02*
Percent High-priority message [a]	81percent	41percent	*0.001*
Percent Queries matching addressee's expertise [a]	32percent	45percent	*0.01*
Attention behavior			
Percent Messages read [a]	42percent	45percent	*0.99*
Percent Messages replied to [a]	04percent	05percent	*0.60*
Performance			
Total Earnings [b]	$2.07	1.47	*0.01*
Puzzle Completion Earnings [b]	1.61	1.39	*0.02*
Reward for help [b]	.76	.66	*0.47*
Reward for advertising [b]	.06	.23	*0.19*
Paid for postage [b]	.36	.81	*0.001*

(N=7609 messages nested within 5 rounds per player,
nested within 120 players nested within 24 experimental sessions)

[a] Message nested within player as the unit of analysis.
[b] Player nested within round as the unit of analysis

Figure 1: Crossword puzzle controls

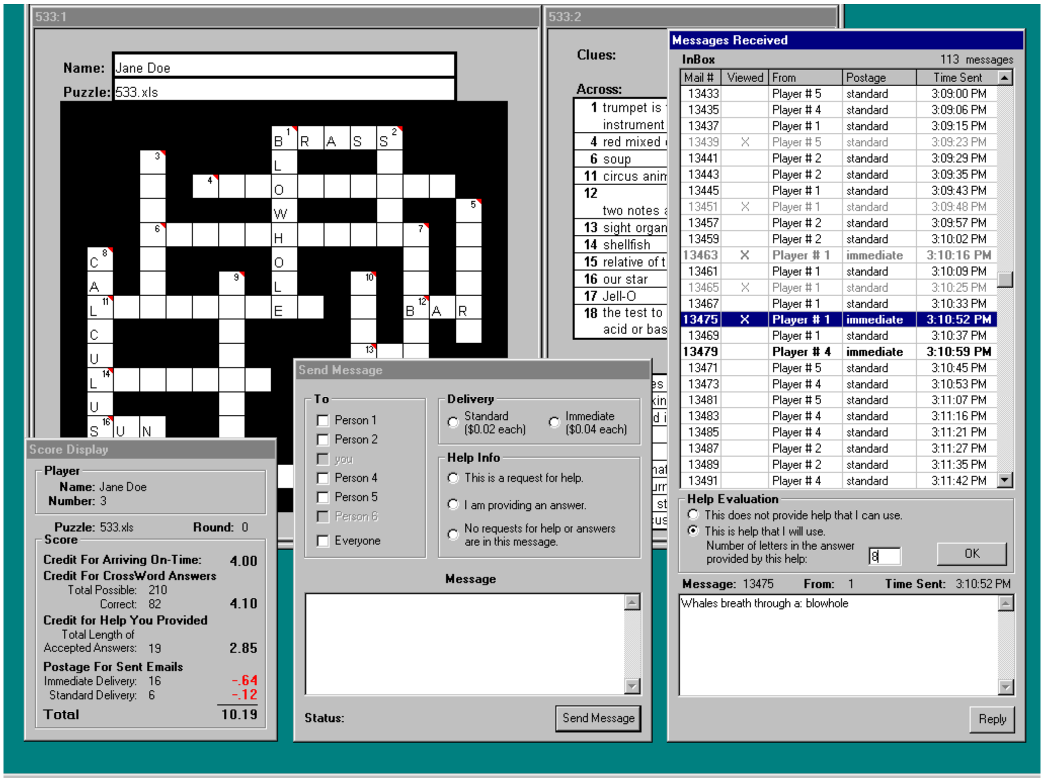

Figure 2: Interaction of postage regime and priority of a message with the likelihood of being read

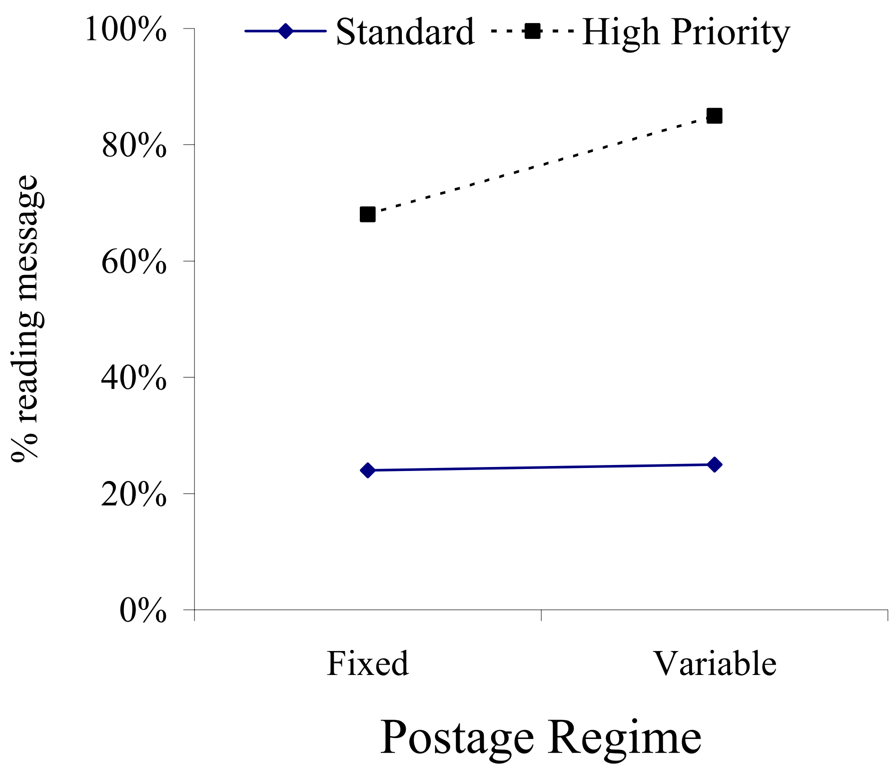

Figure 3. Percentage of all recipients addressed by message type and postage condition

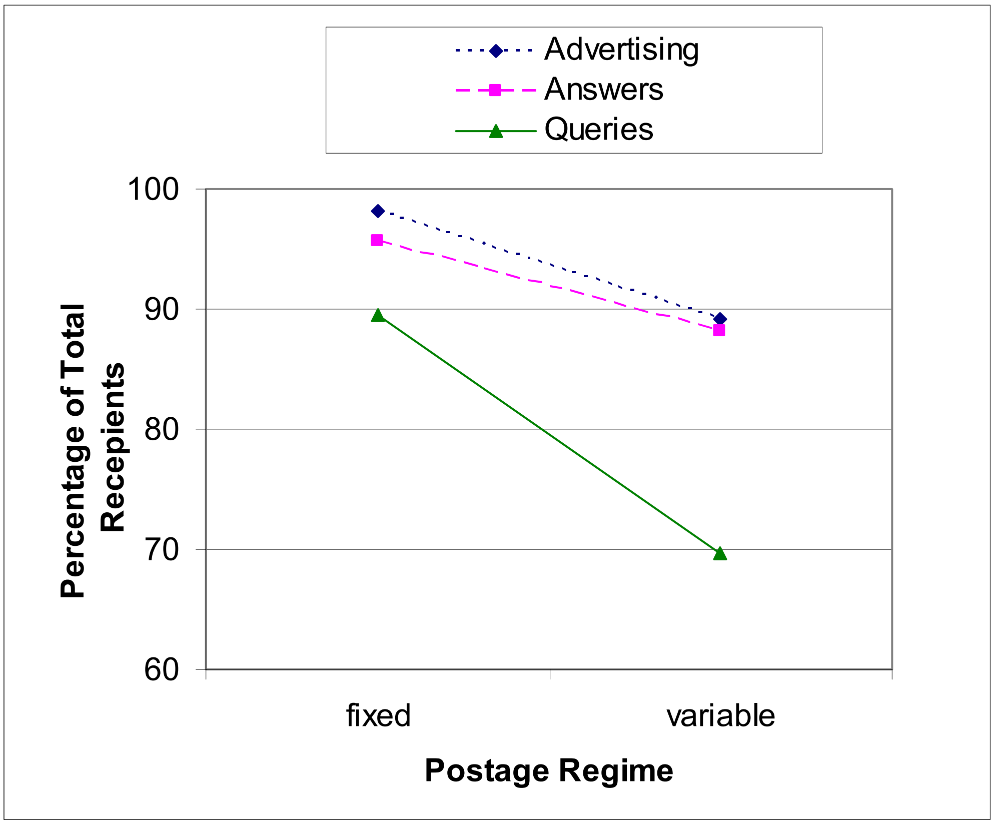

Figure 4. Daily email volume before and after the initiation of Daum's Online Stamp service.