

A penny for your thoughts: e-mail and the under-valuation of expert time¹

[This comment is an extension of Daniel R. Shiman's 'When E-Mail Becomes Junk Mail: The Welfare Implications of the Advancement of Communications Technology', published in the *Review of Industrial Organization*, vol. 11 (1996), pp. 35-48.]

Abstract: The cost incurred by the sender of an e-mail does not reflect the costs to the recipient, leading to a larger number of messages being sent than is optimal for the general welfare. As a solution, we suggest a per-message e-mail tax on the sender similar to that proposed by Shiman (1996), with the addition of a recipient-determined 'contact list'. The recipient derives utility from messages sent by those on this list, and they are not to be taxed when sending messages to the recipient. We recommend that the tax revenue be divided between a regulating body and the e-mail recipient's Internet service provider.

Chris Willmore
The University of Warwick

December 16th, 1999

I. Introduction

E-mail has drastically reduced the cost and increased the speed of communication. The Internet's searchability makes finding people simpler than ever before, so that people can easily contact the authors of a book or article they've read, and other minor celebrities.

And they do. Experts in any given field find themselves swept away by a tsunami of fan mail, intricate queries, commentary on their work, requests for free advice or opinions and other clutter. Daniel Shirman (Shirman, 1996) recognised that there is a cost involved in reading a message, even if no answer is given. For the affected experts, the time spent in scanning their mail can quickly add up, particularly since public relations and professional considerations often make filtering the messages or deleting them unread unfeasible. If an answer is given, then there is the additional opportunity cost of time that could have been spent more productively.

¹ The author would like to thank A. Oswald and L. Willmore for their helpful comments on an earlier version of this note.

The problem lies in that the senders of the letters seldom place any value on the expert's time. The present paper analyses the situation and suggests a solution for this under-valuation of expert hours. A reporter looking for quotes for her article is used as our model e-mailer, but this is merely for concreteness of illustration. She could well have been a doctoral student asking for a critique of a chapter of her thesis, or a well-meaning Netizen sending a combination fan letter and commentary on the expert's latest work.

II. The Model

We assume the existence of a reporter who can contact experts, each with a probability x_i of answering before the deadline. The expected gains to the reporter from contacting N experts is²

$$G_r = u(a_e) - sN$$

Here, a_e is the expected number of answers, given by

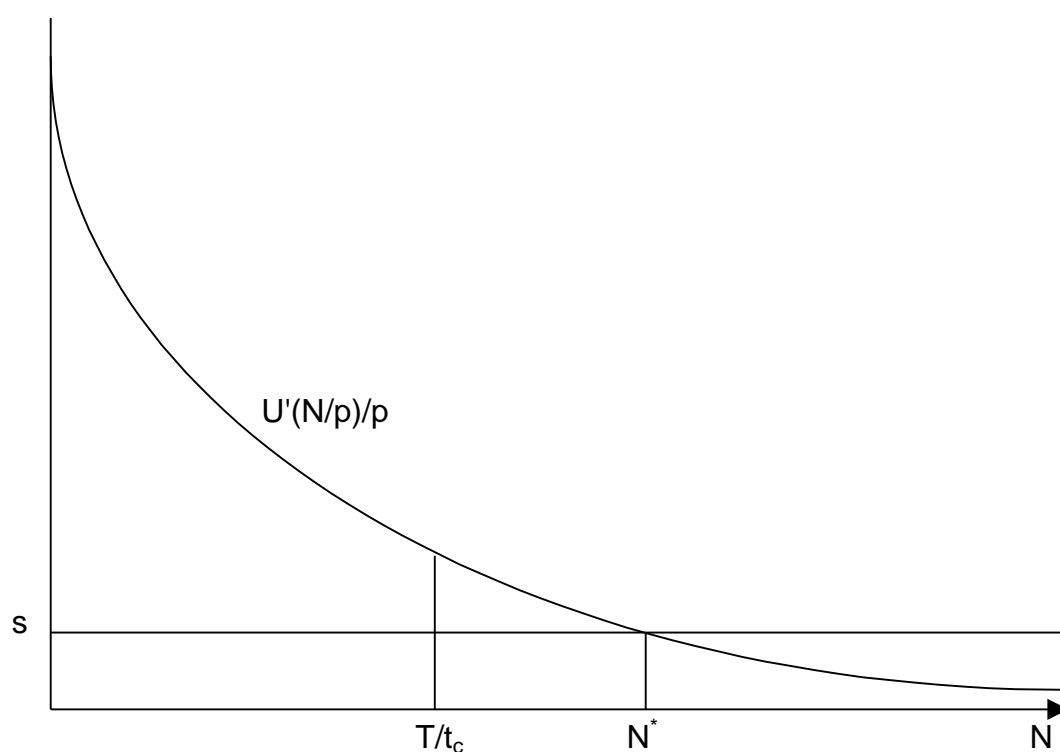
$$a_e = \sum_{i=1}^N x_i$$

Let us assume that $x_j \geq x_{j+1}$ for all j , so that increasing N means digging deeper into the expert pile. The first chosen are the most promising. The cost of sending an e-mail is s , and u is the utility gained from the expected answers. The utility term is increasing in a_e and shows sharply diminishing returns: the first expert opinion received is essential; the seventh scarcely counts as icing on the cake.

² This gain function was arrived at by questioning three professional journalists on the costs and benefits of expert opinions. Though simple, it matches reality as described by the three reporters.

For purposes of illustration, let us suppose that all $x_i = x = 1/p$ for all i , where p is some positive integer. Then $a_e = N/p$, and the reporter will contact experts until the costs of sending messages is equal to the benefit from the expected replies: that is, s is equal to $u'(N/p)/p$. The cost of sending an e-mail is typically low enough so that time constraints take precedence over cost considerations. If it takes time t_c to contact each expert, then we must have $Nt_c < T$, where T is the time left before the deadline.

From the above we may construct the following diagram:

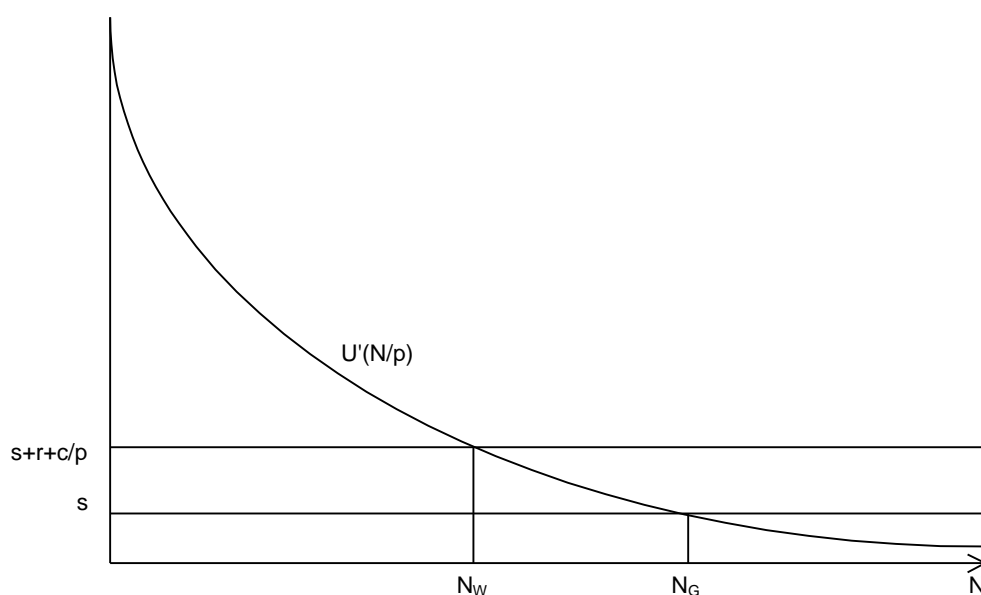


Recall that $u(a_e)$ is the utility from a_e expected answers. In the absence of time constraints, to maximise gains the reporter will set marginal utility equal to the marginal cost of sending a message, and the point chosen will be N^* , the

intersection of u' and s . In making her choice, the reporter doesn't take into consideration the value of the expert's time³. The total welfare function is

$$W = u(a_e) - (s + r)N - c a_e$$

where r is the cost of scanning and reading a letter. We follow Shiman in assuming that experts also receive personal messages (from friends, family, colleagues etc.) yielding them a total utility U inclusive of reading costs, and that $U > r$ so that they will read all messages that they receive. That is, all experts incur a cost r from an e-mail being sent to them, regardless of whether they answer or not. The average opportunity cost of answering a letter is given by c , and weighted by a_e , the number of expected answers⁴. If we once more assume that the probability of answering is $1/p$ for all experts, and omit time constraints, we obtain the following diagram:



³ Except, possibly, in the calculation of x for a given expert. The reporter may assume that a particularly well-known (and presumably particularly busy) expert will have a lower probability of giving her an answer.

N_W is the optimal number of experts consulted according to the general welfare equation, and N_G the reporter's optimum. Note that $N_G > N_W$. Because the costs of sending e-mail are minuscule, a reporter will be willing to contact experts for even a negligible gain in utility. Meanwhile, the work and leisure foregone by the expert in answering the message represent a much higher opportunity cost.

III. Solutions

We need a way to make the reporter's costs reflect at least a portion of the opportunity costs of the experts she consults⁵. The strongly diminishing returns of the utility from answers, $u(a_e)$, implies that even a small increase in sending costs will greatly reduce the reporter's optimal N . Taxing e-mail at the rate of τ /message would accomplish this. Setting τ to the level proposed in Shiman's paper has the added advantage of taking advertisers into account. A tax would discourage advertisers and make reporters pickier when contacting experts, but it would also decrease the number of personal, utility-generating messages sent. This difficulty can be bypassed by allowing each e-mail user to maintain a 'contact list'⁶ of known addresses that she is willing to receive messages from. The people on this list would have τ waived when sending e-mail to that user. The author has observed that there are those (especially newcomers to the Internet) that enjoy receiving ads and random messages in their mail. For these, an option could be added allowing users to set themselves as τ -free recipients for all messages.

⁴ We have assumed that all experts incur the same opportunity costs for clarity of exposition. It is a straightforward, if messy, matter to attribute a specific wage rate and answering time to each expert.

⁵ I will not deal with another plausible solution: the experts in question can charge for their answers. This has its public relations down side, and opens up an entirely different can of worms. There are a few web sites devoted to maintaining a stable of experts-for-hire, but I am unfamiliar with their success. The issue merits investigation.

⁶ Such lists are standard features of instant-message clients such as ICQ or AOL Instant Messenger (AIM). Yahoo!'s web mail includes a 'block address' option, from which one may infer the existence of a maintained list of addresses that are denied e-mail access to the account holder. Thus, the necessary technology is already in place.

Who gets τ ? The Internet is notoriously international - there is no single government to which one must pay taxes. The obvious (but not necessarily correct) choice is to split the tax revenue between the W3 Consortium (W3C), currently chaired by the man who invented the Internet, and the e-mail account holders' Internet service providers (ISPs). The W3C already defines HTML standards, and is better-suited than most other existing bodies to exert a regulatory function. Meanwhile, ISPs could offer the contact lists as a feature and take τ as an extra source of income. This would work as a subsidy, paying for the costs of maintaining the lists and allowing them to lower (perhaps waive?) their base connection rates and on-site advertising intensity. Some leeway should probably be allowed in setting τ , so that there can be rate competition as in other telecommunications industries. (Note that telephone transmissions also have a negligible marginal cost.) To ensure that there is no cheating – that is, to make sure that no one charges a zero τ - billing should be from the message recipient's ISP to the sender's own provider.

IV. A simple numerical example

Let's take a look at what the magnitude of the tax would be, given the assumptions we've been making all along. Working in dollars, let us assume that the cost of sending an e-mail, s , is \$0.01/message, that it takes an expert 2 minutes to scan a message and 20 minutes to reply to the average query, and that the going wage for experts is \$25/hour. Once more, we assume a uniform probability of answering $1/p$, and in this case we take $p=2$ – a 50/50 chance of getting an answer. For ease of calculation we'll also assume that all the taxes paid by the reporter on messages sent return to the expert in some form, perhaps through lower connection rates from her Internet service provider. Plugging in numbers, our welfare function becomes

$$W = u\left(\frac{N}{2}\right) - \left(.01 + \frac{25}{30}\right)N - \frac{25}{3}\frac{N}{2}$$

Differentiating,

$$\frac{\partial}{\partial N} W = .500 D(u)\left(\frac{1}{2}N\right) - 5.01$$

If we include a tax τ /message in the costs of sending a message to the reporter and plug in our numbers, her gain function becomes

$$G_{\tau} = u\left(\frac{N}{2}\right) - (.01 + \tau)N$$

Differentiating,

$$\frac{\partial}{\partial N} G_{\tau} = \frac{1}{2} D(u)\left(\frac{1}{2}N\right) - .01 - \tau$$

A quick inspection of the expression for the two derivatives shows that the optimum tax level is \$5/message⁷. This seems a bit steep, representing as it does about a quarter of the usual monthly Internet connection fee in the US. Though good at fending off nosy reporters, it would be punitive for amateur senders who wish to, say, get in touch with a friend whose contact list they aren't on yet. Could a non-linear tax bring the reporter's preferred N up to par with the general optimum and be less harsh on first-time senders? After all, we wish to discourage *excessive* e-mail, and not suppress it altogether.

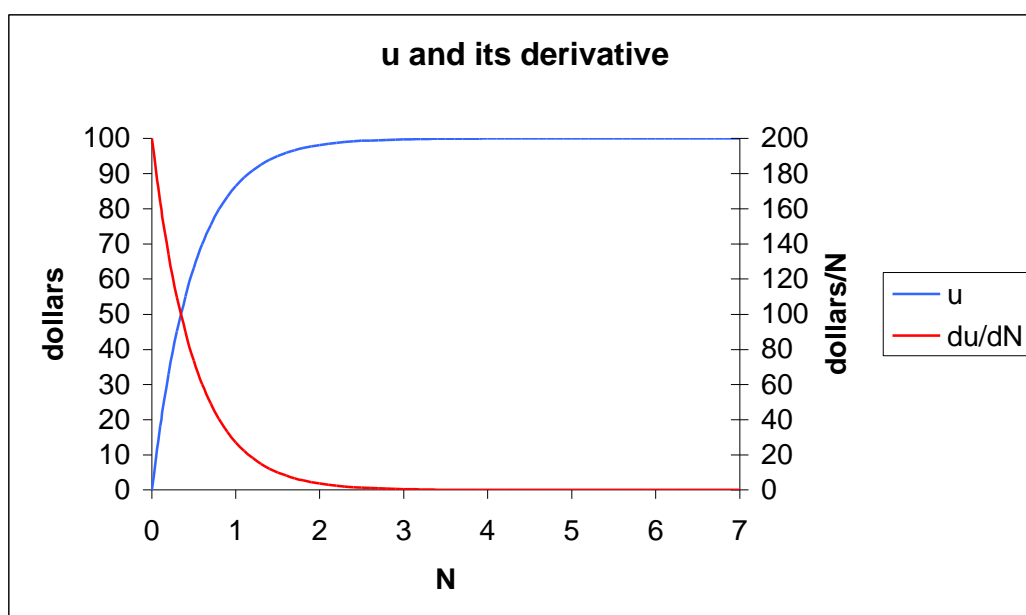
⁷ To find the maxima, these derivatives would be set to zero and solved for N , and the optimum N obtained from W would be compared to the $N(\tau)$ obtained from G_{τ} to solve for τ . Clearly, setting τ equal to the difference in costs between the general welfare and reporter's gain functions is a solution for this linear tax case.

In order to find out, we'll need to provide an explicit mathematical expression for u . Working on previous assumptions, we require a function that provides most of the utility at $N=1$, and for which the jump from $N=6$ to $N=7$ gives only a negligible gain. Such a function⁸ and its derivative are shown below:

$$u = \text{umax} \left(1 - e^{\left(-4 \frac{N}{p} \right)} \right)$$

Setting $\text{umax}=\$100$, $p=2$,

$$u = 100 - 100 e^{(-2 N)}$$



The graph of du/dN is consistent with those sketched earlier. Plugging in the explicit function for u into our numerical general welfare function, we obtain

⁸ It should be possible to provide an empirical basis for this function by tallying the density distribution of experts cited by journalists. The maximum utility, in dollars, is presumably the amount that the reporter is paid for the article.

$$W = 100. - 100. e^{(-2. N)} - 5.01 N$$

Differentiating,

$$\frac{\partial}{\partial N} W = 200. e^{(-2. N)} - 5.01$$

To find the maxima of W , we set the left-hand-side to 0 and solve for N , yielding $N_{\text{opt}} = 1.843440726$. In our case, a fractional N is nonsensical, and we round this off to $N_{\text{opt}} = 2$.

Now that we have the general optimum, we wish to gauge the effects of varying degrees of nonlinearity on the first-message tax. Let us take the nonlinear tax to be given by

$$T = \tau N^x$$

where τ is the tax on the first message, and x is some rational number greater than or equal to 1.

The reporter's numerical gain function is now

$$G_T = u\left(\frac{N}{p}\right) - .01 N - \tau N^x$$

$$G_T = 100 - 100 e^{(-2 N)} - .01 N - \tau N^x$$

Differentiating,

$$\frac{\partial}{\partial N} G_T = 200 e^{(-2 N)} - .01 - \frac{\tau N^x x}{N}$$

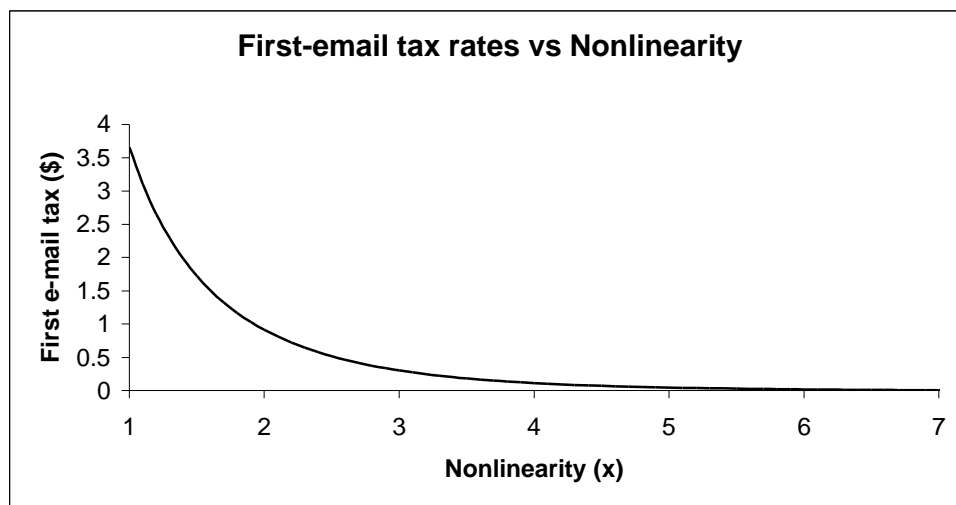
And if we substitute $N=N_{opt}$,

$$\frac{\partial}{\partial N} G_T = 200 e^{(-4)} - .01 - \frac{1}{2} \tau 2^x x$$

To find the τ that will maximise the reporter's welfare, we set the left-hand-side to zero. This allows us to find an explicit expression for τ in terms of x :

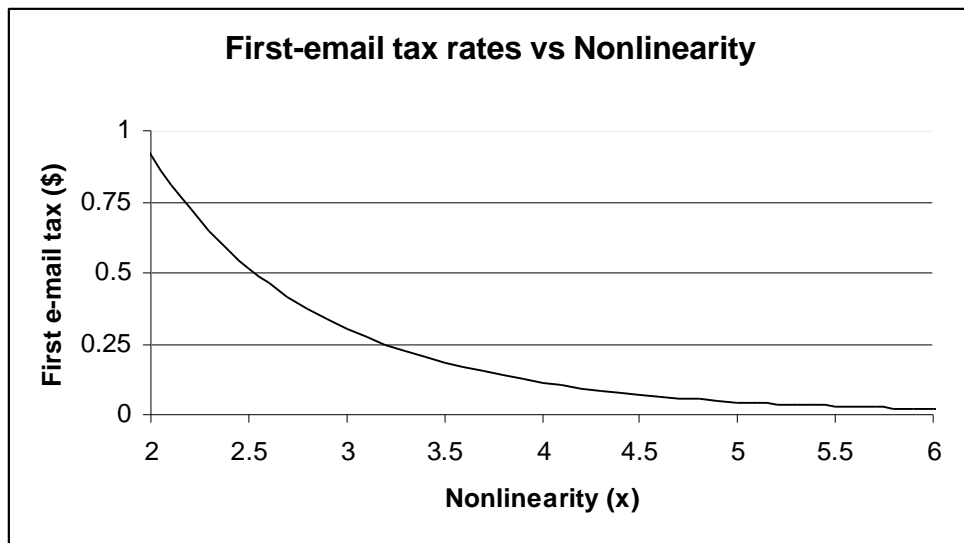
$$\tau = 7.31 \frac{1}{2 \cdot x}$$

The relationship is graphed below:



The discrepancy between the five dollars calculated earlier as the proper τ for $x=1$ and the value shown in the chart is entirely due to the rounding up of N_{opt} to $N_{opt}=2$.

The region where the initial tax is in less than one dollar and more than one cent is of particular interest, since it represents what is probably a reasonable tax rate:



So far, so good, then. It appears all one has to do is choose a value for τ that doesn't prove too onerous for the average utility-providing user (that is, one who will eventually appear on the recipient's contact list), and pick off the appropriate level of non-linearity from the graph above. One can even frame the choice mathematically by including a community of senders and recipients of all sorts in the general welfare function, and formally derive the utility-maximising levels of τ and x .

A tax, however, is only ideal insofar as it can be collected, and the implementation of a non-linear tax structure presents at least two major difficulties. The first has to do with the definition of N . We've treated it as the number of messages sent, but always with the implicit caveat of *for this round of bugging the experts*. If N were the number of messages ever sent by the user, period, and the tax rose as the fifth power of messages sent, the charge per e-mail would soon become prohibitive. What should the time in between 'resets' of N be, then? A day? A month? An hour? A reasonable solution would be to allow the per-email tax to rise until some societally acceptable maximum before resetting it to the minimum. With this the analysis becomes a lot less straightforward, and even in our simplest case the optimum values of τ and x are not easily solved for. Then there's the problem of technological tax evasion.

There would be a strong incentive to cheat by changing the clock of the sender's computer. In theory, this problem could be solved by forcing all e-mails to pass through a node of a network of routers with synchronized system times, but in practice policing the system could prove well-nigh impossible.

The second major difficulty is even less amenable to resolution. How can one track how many e-mails each *person* sends? With a non-linear tax, a reporter might well choose to swap accounts or even computers in between e-mails to avoid the accumulation of N on any given address. The problem is akin to that of a real-world post office trying to charge non-linearly for delivery of letters. In a short time, it would find itself faced with a large increase in its usual letter-sending population, as evidenced by the names on the 'sender' portion of the envelope. Identity is too indefinite and malleable a thing in cyberspace to allow for a precise head count – with current technology. There is always the chance that the requisite infrastructure will be developed and implemented. It is with this possibility in mind that the above discussion is retained.

For the time being, the best choice is to use a linear tax and either use the discreteness of N to good advantage, or to compromise between the needs of society at large and the needs of the reporter. Let's take the first case. In our simple (linear) example, the N closest to the derived societal optimum is $N_{opt} = 2$. To make the reporter choose $N = 2$, all we must do is set the tax *just* high enough so that the reporter will not choose $N = 3$. It is a matter of straightforward algebra to find that this happens at about $\tau = \$0.50/\text{message}$, a considerable reduction from the previously calculated $\$3.60$. It might be that 50 cents a message is still considered punitive, given that some utility-providing message senders will have to pay the cost until they are added to their acquaintances' contact lists. If so, a trade-off must be made between a lower τ and a higher N . Appropriate assumptions about utility can be made and inserted into the general welfare function to find the optimum balance.

References:

Shiman, D.R., 1996, *When E-Mail Becomes Junk Mail: The Welfare Implications of the Advancement of Communications Technology*, Review of Industrial Organization, vol. 11, pp. 35-48