

The Honest Inbox

Contents

1. What Email Is For (And What It Is Not)
 - What it is bad at
 - What it is good at
 2. The Cost of a Quick Question
 - The discipline of writing the question
 3. On Replying Slowly
 - What slow replies cost
 - What to tell people
 4. The Letter You Did Not Send
 - The drawer is the point
 - A practical version
 5. Notes on the Subject Line
 - A pattern that works
 - The cost of un-subject lines
 6. Why Inbox Zero Is A Tax
 - The alternative
 - The hidden cost
- Afterword
License

The Honest Inbox

Six essays on email, attention, and async communication

— T. Voss

Contents

1. [What Email Is For \(And What It Is Not\)](#)
 2. [The Cost of a Quick Question](#)
 3. [On Replying Slowly](#)
 4. [The Letter You Did Not Send](#)
 5. [Notes on the Subject Line](#)
 6. [Why Inbox Zero Is A Tax](#)
-

1. What Email Is For (And What It Is Not)

Email was designed as a slow, asynchronous medium for messages that did not require an immediate response. That is not what it is anymore. That is what it was supposed to be.

The drift happened gradually. The first email message I ever received, in the early 1990s, sat in my inbox for two days before I noticed. Nobody minded. The default expectation was that email was checked occasionally, in the way you check a mailbox at the end of the road. You wrote a letter, you put it in the mailbox, and you waited until the recipient happened by their mailbox and noticed it. The medium was patient by design.

Then we got mobile phones with email clients on them, and the patience evaporated. The mailbox started to follow us. The notification badge appeared. The expectation that “you would have seen it by now” tightened from two days to two hours to twenty minutes. And once an expectation hardens, it does not soften without deliberate effort from everyone involved.

The result is that we now use email for two purposes that do not actually fit it: urgent communication, and ongoing conversation. Email is bad at both.

What it is bad at

Email is bad at urgent communication because it has no mechanism for guaranteeing the recipient has seen it. The “read receipt” check is unreliable and socially awkward to enable. If something is genuinely urgent, you need a channel that announces itself: a phone call, an SMS, a face-to-face request. Email is the wrong channel for “I need this in the next hour.”

Email is bad at ongoing conversation because each message stands alone, with no shared context that grows over time the way a real conversation does. Threading helps a little. But anyone who has tried to onboard a new team member into a fifteen-message email thread will tell you that the thread is more confusing than helpful. Real conversations have shared memory. Email threads are reconstruction projects.

What it is good at

Email is good at delivering self-contained messages that the recipient can read when they have time and respond to thoughtfully. Examples that fit:

- “Here is the document I promised. Take a look when you can.”
- “I wanted to introduce you to my colleague Maya. Bios for both of you below.”
- “Following up on our conversation last week: I have decided to take you up on the offer. Details below.”
- “I am writing to thank you for the generous afternoon yesterday.”

These messages do not require an instant response. They benefit from being written carefully and read carefully. They are the kind of communication email was built for, and they are still the kind it does best.

The discipline is to use email for what email is good at, and to use other channels for everything else.

2. The Cost of a Quick Question

The phrase “quick question” is one of the most expensive utterances in modern work, and the cost is almost entirely borne by the recipient.

When someone sends you a quick question by email, what they are actually doing is offloading their cognitive cost onto you. They have a thing they want to know. Rather than spend the five minutes to find the answer themselves (search the wiki, read the document, ask a search engine, try the thing), they spend the thirty seconds to send you a message. The thirty seconds feels free to them. The cost to you, however, is rarely thirty seconds.

The cost to you is:

- The interruption to whatever you were doing
- The time to read and parse the question
- The time to consider whether you actually know the answer
- The time to write the answer in a form that will be useful
- The time to recover from the interruption and get back to what you were doing

Recent research on interruption costs in knowledge work puts the recovery time at roughly twenty minutes per interruption. If you were doing deep work, that is twenty minutes of compounded loss, plus the time spent on the interruption itself.

So when someone sends you a “quick question,” they have paid thirty seconds and you have paid, on average, twenty-five to thirty minutes. That is a 50-to-1 trade. Multiply by a few dozen quick questions per week and you have an explanation for why the responsive team members feel exhausted while the askers feel productive.

The discipline of writing the question

The honest fix is to put the cost back on the asker. The discipline I would suggest, both for asking and for receiving, is this: write the question in a form that costs the asker fifteen

minutes, not thirty seconds. That means:

- Stating the context clearly
- Stating what you have already tried
- Stating what specifically you are stuck on
- Stating what good answer would look like

A question written this way is much easier to answer, because the answerer does not have to reconstruct the problem from scratch. It is also much less likely to be sent in the first place, because the discipline of writing it often surfaces the answer on its own.

The pattern in functioning engineering teams is sometimes called “rubber duck debugging” applied to questions. The mere act of writing out the question in full often makes the answer obvious. The question never gets sent. The interruption never happens. Everyone wins.

I would estimate that fewer than 10% of the “quick questions” I receive could not have been answered by the asker spending five more minutes on the problem. The other 90% are essentially attention theft, charged at a 50-to-1 rate. Make the askers do the work.

3. On Replying Slowly

I have spent the last three years deliberately replying to email slowly. Most replies go out within 24 to 48 hours of receipt. Some take a week. A few take a month. A small number never get replied to at all, by design.

The first thing I noticed is that almost nothing breaks.

The second thing I noticed is that my replies got better. When I respond within minutes, I am responding to the literal words on the screen, often without having read them carefully. When I respond a day later, I have had time to think. I have noticed what the sender was actually asking, which is often different from what they wrote. I have considered the second-order effects of my answer. I have, in the cases that warrant it, slept on it.

The replies that result are shorter, more useful, and less likely to require follow-up. The total volume of email I generate is meaningfully lower because each reply tends to be the right reply, not a first draft I will have to correct in the next message.

What slow replies cost

The cost of slow replies is real and worth naming. Some people will be frustrated. Some opportunities will pass because the offerer needed a faster response. Some relationships will cool because the other person experiences your slowness as indifference.

I have made peace with this. The cost of fast replies is also real: it is paid in the quality of every other thing I do during the day, because every fast reply interrupts something. The slow inbox is not free, but it is, on the long view, the cheaper of the two.

The signal I look for, to decide whether the slow inbox is hurting me, is not “are people frustrated” (some always are) but “am I missing things that matter.” In three years, I can count on one hand the things that mattered which the slow inbox cost me. In the same period, the slow inbox has bought me back something like a thousand hours of contiguous attention that I would otherwise have given up to email.

What to tell people

If you choose to reply slowly, the kind thing to do is tell people. I have, for years, kept a one-line note in my email signature that says, in effect: “I check email twice a day. If you need something faster, please text me.” This sets expectations cleanly. The people who actually need fast access have it. The rest get my best attention on a slower schedule, which is, almost always, more attention than they got from my old fast self.

This is the only way the slow inbox works at scale. You cannot just be slow and let people