The Speed of Waiting

Fast delivery, a quick reply and a finished task measure different things. An evidence-led essay on email, waiting and human work.

2026.11.16 04:16 DENNIS HEDEGREEN ANALYSE V1.0
HTTPS://HEDEGREENRESEARCH.COM/ARTICLES/THE-SPEED-OF-WAITING/

When a fast message still leaves a task unfinished.

Imagine trying to arrange a meeting. A date has been suggested, a time accepted and an invitation promised. The invitation never appears. The next day comes a spontaneous question: is there time to talk now? A reply follows within a quarter of an hour. And then there is more waiting.

Nobody needs a horse, a railway or a postman. The messages cross the distance almost instantly. The meeting still has not happened.

This is an illustrative situation, not a measured case. It exposes the question hiding inside the word *fast*. What, exactly, has become faster?

The letter was not always slow

London had unusually dense local postal services long before email. The London Museum describes William Dockwra's 1680 Penny Post, with several daily deliveries in busy districts and a four-hour delivery guarantee. The national reforms introduced in 1840 later made postage cheaper and simpler. These were different services in different periods.

A transcription of Charles Dickens Jr.'s 1879 London guide specifies twelve daily deliveries in the Eastern Central district and eleven within the town limits of the other districts. Suburban services were less frequent. This is a contemporary description available through a later transcription, not a measurement of achieved delivery times.

Twelve deliveries is a startling number. It also belongs to a particular place: a dense capital with postal infrastructure that cannot stand in for the whole nineteenth century.

A delivery schedule tells us when a letter could reach someone. It does not tell us when that person read it, whether they answered or how long the matter took to settle. We need to keep that distinction when we turn to email.

A real speed advantage

A 1997 doctoral study by Kimberly Post Good selected 528 university employees with email accounts and assigned them to four survey treatments. Among completed surveys returned through the original channel, mean return time was 3.05 days for email and 14.68 days for post. Six email questionnaires returned by post

were excluded from that timing analysis. These are elapsed return times among responders, not minutes spent answering.

The response-rate comparison was more complicated. After follow-up, personalised postal surveys returned 88 of 131 questionnaires (67.2%), against 111 of 217 (51.2%) across three email subgroups. The email groups differed in personalisation and prenotification; reminder schedules also differed between channels. This was not a comparison in which everything except the medium stayed the same.

The electronic method plainly returned completed questionnaires sooner in this setting. That matters: digital transport can remove days from a process. But a faster return among responders and a higher chance of getting an answer are different outcomes. Neither establishes how quickly an ordinary conversation reaches a decision.

Forty-seven minutes — but whose forty-seven minutes?

Kooti and colleagues' 2015 study examined more than two million opted-in Yahoo Mail users. Its wider dataset contained sixteen billion emails; the selected conversations between pairs contained about 187 million. Pairs had to have exchanged at least five replies in each direction. These were already reciprocating contacts, not a sample of every request someone might send.

Among observed replies in those conversations, the median delay was 47 minutes and more than ninety percent arrived within a day. Reply time ran from the incoming message, or the first in a consecutive series, to the subsequent reply. The paper also found a lower fraction of messages answered as incoming load rose, while sent replies could become faster. These are observational patterns within its dataset.

A message never answered cannot appear as a completed reply in that median. That is a boundary of the statistic, not a mistake by the researchers, who separately examined response fractions.

A busy person could answer several easy questions quickly and leave an important request unresolved. That example would not prove that their whole workflow had become slower. It shows why reply speed alone cannot settle the question. We need to know which requests obtained a useful answer, and what happened afterwards.

Four clocks

The first clock is **transport time**: from sending to arrival. Digital messages often cross that distance almost instantly. A physical letter depends on route, service, period and distance.

The second is **reply time**: from the request to its first answer. An automatic receipt or a polite acknowledgement may stop this clock without advancing the task. A useful answer needs its own definition and timestamp.

The third is **resolution time**: from opening a practical matter to reaching its outcome. A meeting has been booked, an invoice paid or a technical problem solved. A reply and a resolution can occur together, but they do not have to.

The fourth is **human work time**: the person-minutes spent reading, composing, checking, forwarding and chasing, across everyone involved. Three calendar days can contain ten minutes of work. An hour-long exchange can consume much more than an hour of labour when several people participate.

Waiting has another boundary. If someone can leave the matter alone and do something else, the whole delay is not active labour. If they must stay available, keep checking or remain unable to leave, that restriction should be recorded separately. Counting every hour until an answer as an hour of work would create another misleading saving.

In We Saved Eight Hours. Did We? (https://hedegreenresearch.com/articles/we-saved-eight-hours-did-we/), I asked whether two processes still provide comparable output and service before subtracting their human-time requirements. Email makes that problem ordinary. Faster transport does not, by itself, establish how much human work a completed task requires. Only after establishing a real reduction can Saved Hours (https://hedegreenresearch.com/articles/the-saved-hours-doctrine/) ask who receives the benefit.

The cost of expecting an answer

Giurge and Bohns reported eight preregistered studies involving 4,004 participants in 2021. They examined non-urgent work emails outside normal working hours. Recipients overestimated how quickly senders expected replies, with consequences for perceived stress. Explicitly clarifying that a quick answer was unnecessary reduced that expectation gap in their experiments.

This does not establish that every recipient feels compelled to reply immediately. It shows why perceived urgency and intended urgency need not match. A system can deliver the message successfully while leaving people uncertain about what the arrival demands of them.

An immediate reply is not always the right objective. The ability to answer at a chosen time can be one of email's benefits. A useful communication system should let people distinguish a genuine deadline from someone else's convenient moment to send.

A faster inbox can hold more work

Microsoft's June 2025 WorkLab report describes 117 incoming emails daily and 153 Teams messages per weekday for the average worker in its analysed population. Its Microsoft 365 telemetry excludes education and EU tenants and ends on 15 February 2025. These are company-reported workplace figures with those boundaries, not a population estimate for all workers.

Volume alone cannot tell us whether the communication is useful. More messages might open opportunities and relationships that would have been impossible when contact cost more.

There is nevertheless a human cost worth measuring. In a 2012 field intervention, thirteen workers experienced a baseline period followed by five workdays without new email. Window logs showed fewer switches and longer stretches with one window active, a proxy for task focus. The physiological stress

comparison had usable paired data for only seven participants; it was not a stress result from thirteen complete records. This was a small, short study in one organisation, with other ways of contacting colleagues still available.

A 2016 observational study of forty workers, averaging twelve study days, associated longer email duration with lower perceived productivity and higher measured stress. The authors explicitly cautioned against causal interpretation. They did not find evidence that batching email reduced stress; that is not proof that batching can never help.

Neither study supplies a historical balance sheet for email. They give us narrower reasons to examine the effort and pressure of handling messages alongside the speed of delivery.

Measure the task

The evidence reviewed here supports real benefits from digital communication, including faster survey returns in one comparison. It also shows why a delivery schedule, a return-time average and a conditional reply median cannot be placed on one historical graph and called the speed of human coordination.

I have not established a comparable cross-era estimate of end-to-end task resolution. That is a limit of this review, not a claim that no relevant research exists, and not permission to conclude that email made people slower.

A useful next study could start with a narrow task: arranging an appointment, for example. Define the beginning, the useful reply and the completed booking before collecting data. Record reminders, other channels, everyone else's work and whether the outcome was useful. Keep requests that never receive a reply in the cohort.

Report how many tasks finish within a fixed window as well as how long completed tasks take. Cases still open at the observation cutoff have been observed waiting up to that point; their eventual completion time is unknown. Explicit refusals, withdrawals and abandoned requests need their own outcomes rather than being treated automatically as merely delayed successes.

Any future historical comparison would also have to confront who could use each channel, which records survived, and whether the tasks were genuinely comparable.

Perhaps the best communication system is one that helps people reach useful outcomes with fewer demands on their attention while preserving the right to answer later. That is an objective to test, not a result these studies have already delivered.

An instant message is not an instant agreement.

The message can arrive. The answer can still wait.

— Dennis Hedegreen, trying to see the structure

Sources and scope

1. **London Museum — The Penny Post sparked a letter-writing revolution.** Read the museum's account (https://www.londonmuseum.org.uk/collections/london-stories/penny-post-sparked-letter-writing-revolution/). Institutional historical summary; London-specific.
2. **Charles Dickens Jr. — Dickens's Dictionary of London (1879), "Postal Regulations."** Read the Victorian London transcription (https://www.victorianlondon.org/dickens/dickens-que.htm). The original printed edition was not independently inspected here. Town Deliveries specifies twelve daily deliveries in Eastern Central and eleven in the other town districts; these are schedules, not reply times.
3. **Kimberly Post Good — A Study of Factors Affecting Responses in Electronic Mail Surveys (1997).** University record and abstract (https://scholarworks.wmich.edu/dissertations/1635/); mirrored dissertation scan (https://scispace.com/pdf/a-study-of-factors-affecting-responses-in-electronic-mail-4pw6zpd9pd.pdf). Methods and Tables 11–13, printed pp. 28–33 and 47–52. Direct university PDF access was unavailable during this review. Timing is conditional on returned questionnaires; treatments and reminder schedules differed.
4. **Kooti et al. — Evolution of Conversations in the Age of Email Overload (2015).** Read the primary paper (https://arxiv.org/html/1504.00704), §§2, 3.1 and 4.5. Selected reciprocating Yahoo contacts; completed-reply timing is not the response probability for all received messages.
5. **Giurge and Bohns — You Don't Need to Answer Right Away! (2021).** Read the accepted manuscript at London Business School (https://lbsresearch.london.edu/id/eprint/1928/1/The%20email%20response%20bias.pdf). Eight preregistered studies, N=4,004; principally non-urgent workplace messages outside normal work hours, with boundary-condition studies.
6. **Mark, Voida and Cardello — "A Pace Not Dictated by Electrons": An Empirical Study of Work Without Email (2012).** Read the author-hosted paper (https://ics.uci.edu/~gmark/Home_page/Publications_files/CHI%202012.pdf), methods and Tables 2–3. Thirteen workers overall, seven paired physiological records; one organisation and a short sequential intervention.
7. **Mark et al. — Email Duration, Batching and Self-interruption (2016).** Read the institutional publication page (https://www.microsoft.com/en-us/research/publication/email-duration-batching-and-self-interruption-patterns-of-email-use-on-productivity-and-stress/); author manuscript (https://www.microsoft.com/en-us/research/wp-content/uploads/2016/06/Email20Duration20Camera20Ready20submission3-1.pdf), methods, findings and "Notes on causality." Observational, forty workers and about twelve study days; perceived productivity, not a causal output estimate.
8. **Microsoft WorkLab — Breaking down the infinite workday (17 June 2025).** Read the report and methodology (https://www.microsoft.com/en-us/worklab/work-trend-index/breaking-down-infinite-workday). Corporate Microsoft 365 telemetry through 15 February 2025, excluding education and EU tenants; not representative of all workers.

RELATION MEMORY

NEARBY we-saved-eight-hours-did-we · Applies comparable-output and human-time boundary questions to communication.

NEARBY the-saved-hours-doctrine · Measurement must precede claims about who receives a real time saving.

CANONICAL URL: HTTPS://HEDEGREENRESEARCH.COM/ARTICLES/THE-SPEED-OF-WAITING/