

EPISTEMIC COWARDICE AND OTHER PEOPLE'S TIME

On the Cost of Deciding Late

Fredrik Olsson

When an organisation makes a decision, the decision can be traced afterwards. There is a date, someone who made it, and a result that followed. If the result is bad, there is something to go back over and someone to ask. A delay leaves no such trace. Nobody records the months a question spends waiting. And when a decision turns out well, no one goes back to ask why it was made in June when it could have been made in March.

Consider a company deciding whether to renew a contract with a supplier. By March it has most of the information it needs, but the decision is not made until June. Nothing new arrives in those three months, and nothing new is expected. The question comes up at the right meeting more than once and is left open each time. By June the supplier has raised its price by four per cent, on a contract worth a few million a year. Had the company renewed in March, it would have kept the old price. The increase shows up in next year's accounts as a supplier price rise. The entry is correct, and it gives no hint of what caused it.

Delays have many causes, and most of them are reasonable. Some information may still be on its way, or the question may be stuck behind other work. This essay is about a narrower case. The person responsible could settle the question and does not, for reasons that have more to do with their own position than with the facts.

Defensive Decisions

Researchers have measured how often people at work protect themselves at the organisation's expense. A survey of 950 managers in public administration asked about defensive decisions. These are cases where a manager thought one option was best for the organisation, but deliberately chose a weaker one that protected them personally. On average, two and a half of each manager's ten most important decisions in the past year were of this kind, and four in five managers reported at least one. The study traced the problem to organisations that respond to mistakes by looking

for someone to blame, rather than by trying to learn from them.¹ Doctors show the same pattern, which in medicine is called defensive practice. Among doctors in specialties where lawsuits are common, 92 per cent said they ordered extra tests, procedures or referrals to protect themselves, and 42 per cent said they avoided risky procedures or patients.²

Almost all of this is defensive *action*: someone chose the weaker option, or ordered the extra test. Each of these is an act. It has a name attached and a cost that can be counted, which is why a survey can find it. Holding a decision open leaves nothing of the kind, because nothing is decided. The surveys do not measure it directly. The case for it rests on how the incentives work, not on a figure.

Why Waiting Is Safe

When the outcome is uncertain, making a decision is a gamble for the person who makes it. If it goes well, they get some credit. If it goes badly, they get the blame. People judge a decision by how it turned out, even when they knew the odds it was taken against.³ They also judge harm caused by doing something more harshly than the same harm caused by doing nothing.⁴

Saying no counts as a decision here too. A no has a date and a name, and it is judged by how it turns out. Waiting is different. While a decision sits unmade, it has a real cost, but most of the cost comes later and falls on other people. Some of it does come back to the person who waited, as a reputation for being slow or a project that runs late. But it comes back slowly and spread thin, never as a single event with their name on it, because waiting leaves nothing to judge. People are held to a record of what they decided and how it turned out, and that record counts very little when they wait. Call it the scoreboard. On the scoreboard, deciding is a gamble and waiting is close to free.

The shape of the two costs matters as much as their size. The blame for a bad decision

¹Artinger, Artinger, and Gigerenzer, "C. Y. A.: frequency and causes of defensive decisions in public administration," *Business Research* 12 (2019): 9–25. The study documents self-protection over the option judged best, not delay in particular; the application to waiting is the author's.

²Studdert et al., "Defensive Medicine Among High-Risk Specialist Physicians in a Volatile Malpractice Environment," *JAMA* 293 (2005): 2609–17. A survey of 824 physicians in six high-risk specialties in Pennsylvania; evidence that the mechanism exists and can be large, not a base rate for medicine or anywhere else.

³Baron and Hershey, "Outcome Bias in Decision Evaluation," *Journal of Personality and Social Psychology* 54 (1988): 569–79. Vignette experiments: evidence for how decisions are judged, not field evidence.

⁴Spranca, Minsk, and Baron, "Omission and Commission in Judgment and Choice," *Journal of Experimental Social Psychology* 27 (1991): 76–105; Ritov and Baron, "Reluctance to Vaccinate: Omission Bias and Ambiguity," *Journal of Behavioral Decision Making* 3 (1990): 263–77.

is uncertain: it may never come, but if it does, it is large, it has a date, and it is personal. The cost of waiting is small, slow and fairly certain. Someone who is averse to risk will take a small, certain cost over a gamble with a large downside, even when the gamble is no worse on average. A no can be the smaller gamble, but it is still a gamble. Only waiting is not. So even where waiting is not free, it is safe, and it pulls hardest on the people most averse to risk.

This is the author's assumption, not a measured finding: that only a small part of the cost of a delay reaches the person who chose it, and later and more thinly than the blame for a bad decision would. Readers can test it against decisions they have watched. If it is wrong, and people who delay already pay as much as a bad decision would cost them, then no remedy is needed, and the waiting has some other explanation.

What Epistemic Cowardice Is

No decision is ever made with certainty. What a decision needs is enough evidence to go on, and how much is enough depends on how hard the decision would be to undo. A decision that can easily be reversed needs less than one that cannot. So "we are not sure yet" is true of every decision ever made, and on its own it is never a reason to wait. Waiting is worth its cost only while new information is on its way, can arrive in time, and could change the decision.⁵ After that point, waiting longer gains nothing.

Epistemic cowardice is keeping a decision open after that point, when the person has already made up their mind, because deciding either way would put them at risk and waiting feels safe. It is not indecision. The person knows what they would do, and does not do it. Nor is it a lack of authority, where nobody in the room has the power to decide. Here someone does. What separates it from prudence is what the person knows. Prudence is waiting while information that could change the decision is still on its way. Epistemic cowardice is waiting when the person knows that nothing more is coming, in order to protect themselves. And if the evidence is thin and nothing more is coming, waiting still gains nothing. The decision has to be made on what is known, and saying no is a decision too. If the decision can be undone later, making it now costs little.

It is called epistemic because what the person gives up is their own judgment. They still have a view about what should be done, but they will not say it where saying

⁵In decision analysis this is the value of information: Howard, "Information Value Theory," *IEEE Transactions on Systems Science and Cybernetics* 2 (1966): 22–26. Information is worth acquiring only insofar as some answer it could give would change the decision.

it would cost them. The reason they give for waiting is about the evidence. The real reason is not. The two are hard to tell apart, because both are expressed in the language of care. From the outside, caution about the evidence looks the same whether the person is protecting the outcome or their own record.

Two things follow. First, “epistemic coward” is not a label to pin on anyone else. Anyone who calls a slow colleague a coward is claiming to know a motive they cannot see. The term works better as a question people can ask themselves: am I holding this open for the sake of the decision, or for my own sake? Second, having a name for it does not make it go away. The person is reading the scoreboard correctly, and a name does not change the score. On this account, epistemic cowardice is what ordinary self-interest produces when waiting is cheap and safe and deciding is a gamble, rather than a sign of weak character.

Who Pays

Some of the cost is money. It appears in the accounts under other names, never under a heading that says “decided too late”. While the weeks pass, an opportunity closes, a price goes up, a delivery slot is lost, and nobody adds it up.⁶

The rest of the cost is not money, and it falls on particular people. A team cannot start until the matter is settled. Someone has to tell a customer that nothing has happened yet. Finished work sits unused, waiting for a decision to put it to use. These people pay with their own time and effort. They did not choose the delay, and they cannot end it. What they can do is absorb it: they finish the work late instead of dropping it, and they rebuild the schedule around the wait. This kind of work leaves no record, because its whole effect is to make the disruption look normal. It is also why so little of the cost gets back to the person who delayed. The people who pay it hide it by coping.

There is also a cost that spreads. Colleagues can see that a delay costs the person who chose it little. When someone waits and is rarely held to account for it, the people around them learn that deciding is risky and waiting is safe. So the next person with a decision to make waits too. Each delay makes the next one more likely, because everyone is reading the same scoreboard.

⁶Reinertsen, *The Principles of Product Development Flow* (Redondo Beach: Celeritas, 2009), argues that this cost of delay is both real and routinely left unquantified. The claim about who it lands on is not his.

Making Delay Cost More

The scoreboard has two sides, and both can be changed, though not at the same speed. The blame for a bad decision depends partly on the organisation. The survey of managers in public administration, described earlier, traced defensive decisions to organisations that respond to mistakes by looking for someone to blame. Changing that means changing a culture. That is slow, and never complete, because people judge a decision by its outcome even when they knew the odds. The other side, that waiting costs so little, is only an arrangement, and a rule can change it by next week. The remedy is to make the cost of a delay land on the person who chose it, so that a delay has an owner, a deadline and a consequence, just as a decision does. Three devices do this.

The first device is to ask two questions out loud. Would more time change the decision? And what is the wait costing? As long as nobody asks, the wish to wait cannot be examined. Once someone asks, “we are not sure enough” has to be backed by two claims that everyone present can check. One is that useful information can still arrive in time. The other is that waiting for it is worth the cost. If the person holding the decision can answer both, the delay is an honest one. This is the one test for epistemic cowardice that does not require anyone to guess at motives. The answers are about the evidence, and nobody is accused of anything. When the answers do not come, what often comes instead is another reason to wait, and then another when that one has been dealt with. This does not take bad faith. It is what the scoreboard produces as long as deciding is judged and waiting is not.

The second device is a name and a date. A decision that belongs to a committee belongs to no one in particular. A matter that is waiting for more information can wait forever if nobody is responsible for the waiting. Naming one person as the owner, and setting a date, makes the open decision that person’s, due by that date. If it is late, the delay is theirs, and it is on the record. It also becomes possible to ask what other work it was put behind.

The third device is a default: an action, agreed in advance, that happens automatically if nobody has decided by the date. A board might rule that if next year’s budget is not approved by the end of the quarter, this year’s figures will simply carry over. That is rarely what anyone wants, so the carry-over puts a price on the delay. Once a default is in place, leaving the matter open is the same as choosing the default, and a choice is judged in a way that silence is not.

Each device can be gamed, as any measure can be once people are judged by it.⁷ A

⁷The formulation is Marilyn Strathern’s, in “Improving Ratings: Audit in the British University

deadline can be met with an empty decision made only to hit the date. A default is harder to game, because it takes effect whether anyone acts or not. But someone still sets it. If the person holding the matter would rather have the default than make the decision, they can pick a convenient date and simply let it pass.

Only the first device needs no authority. Anyone in the room can ask the two questions, including of someone more senior, and about a decision that is not theirs. The questions do not attack the person holding the decision, because neither of them is about that person. One is about time and the other about cost. The other two devices are standing rules, set in advance by whoever is in charge of how the work is done. Suppose that person answers for the organisation's overall results rather than for any single delay. Then they see the effects of delay only as a total, and cannot tell which delay caused which cost. That is why the remedy is a rule set in advance, not a judgment made case by case.

The three devices are not equally strong. The two questions open a delay to scrutiny, and a name and a date put a person's name on it. Both carry some cost for whoever delays, but less than the blame for a bad decision. The default goes further. It turns waiting into a choice with a result, and that result is judged like any other decision. Waiting becomes a gamble too, and loses the safety that made it attractive. None of the devices makes anyone less self-interested. They work by changing what waiting costs. For a risk-averse person, self-interest favours deciding only once waiting costs more than the average blame for deciding, because to such a person a gamble weighs more than its average.

Copyright © 2026 Fredrik Olsson. All rights reserved.

System,” *European Review* 5 (1997): 305–21, generalising Charles Goodhart’s 1975 observation that a statistical regularity collapses once it is used for control.