

EDITION 2026.1-ANALYST

The work worth automating is the work whose mistakes travel

We scored every billing and operations workflow a home health or hospice agency runs. About a third is worth automating now, and the pattern points at cash posting, claim quality and denial recovery first.

Moative Research · 2026-09-22 · [Read the data](#)





THE WORK WORTH AUTOMATING IS not the easy work. It is the work whose mistakes travel. When a referral is keyed in wrong, the bad data rides the episode from intake to the final claim, and a clerk's typo surfaces months later as a denial nobody can explain. That is the pattern our ranking found, workflow after workflow: the jobs that pay back automation fastest are the ones where one small error spreads into other people's jobs.

We scored every workflow a home health or hospice agency runs, from referral intake to the acquisition playbook, on what it is worth and how

tractable it is. About a third of everything we looked at cleared the bar on both at once. And the winners are the duller jobs in the building: checking a claim before it goes out the door, posting the checks that come back, confirming the patient is covered before the first visit.

The most useful thing the ranking demolished is the excuse everyone reaches for. How hard a workflow is to build scored exactly the same for the work worth starting tomorrow and the work we would leave alone. Difficulty is not why this work goes undone. It goes undone because nobody has connected the mistake to the money it costs.

The ranking says what is worth doing. A second list, the dependencies, says what has to work first. They disagree, and the disagreement is the plan: start with the cash file, not the biggest problem.



Every one of the 169 jobs an agency's billing and operations run on, placed by what automating it is worth against how tractable it is. Hover or tab to any dot.

The chart follows the argument. As you read, the jobs each section is about light up here.

The winner is the dullest job in the building

Claim editing and cash posting beat every clever idea in the study, and they beat it every single time we re-ran the numbers.

First place, in every re-run

We re-ran the whole ranking five thousand times, nudging every scoring weight up and down by a quarter. The claim-editing workflow came first in every single run. Priority score 80.4 of 100.

The clocks are the law

A late home health notice of admission cuts the thirty-day period payment by one thirtieth per late day. A late hospice notice of election makes the missed days the agency's own liability. Sources: 42 CFR 484.205 and 418.24,^{1,2}.

The workflow that came out on top is claim editing, the job of checking assessments, codes, notices and deadlines before a claim leaves the agency. It is not glamorous. It wins because the clocks attached to it are statutory, not managerial: a home health notice of admission must be filed within five calendar days of the start of care³, and every day it is late cuts the whole period's payment by one thirtieth¹. A hospice has the same five-day trap on its notice of election, and days missed there are the agency's own liability, unbillable to the family². One missed deadline is not one clerical slip. It is a payment fraction gone, on every affected episode.

Right behind it sits payment posting, the work of matching what a payer actually sent against what the claim asked for. A home health or hospice agency lives or dies on knowing which episodes are fully paid, which are short, and which are silently accruing underpayments. When posting is

done by hand for smaller payers, the agency is flying blind on its own cash, and everything downstream, from denial follow-up to period revenue checks, inherits the blindness.

The boring pattern is the point. The jobs that topped the study are ones where the mistake has a price tag written into regulation, the volume is daily, and the fix is mostly reading a document against a rule. That is exactly the work a model is good at and a person is slow at. Every cleverer idea in the study scored lower than these two, and no amount of re-scoring moved them off the podium.

The jobs that pay back automation fastest are the ones where a mistake has a price tag written into regulation.

Difficulty is not why this work goes undone

The thing that most reliably separates the worth doing from the rest is how far a mistake travels, not how hard the build is.

The widest gap we found

Work worth building is work whose errors spread into other people's jobs. That gap was larger than for any other input we tested: scored 4.6 against 2.9 out of 5.

Difficulty told us nothing

How hard a job is to build scored the same for the work worth starting and the work we would leave alone: 2.5 against 2.5 out of 5.

Crowded is not a warning

Existing AI competition scored slightly higher among workflows worth doing than among the rest, 3.0 against 3.5 out of 5. Busy tooling marks a real problem.

To find what actually separates the work worth automating from the work worth leaving, we split the workflows into two groups: the ones that scored above the line on both value and tractability, and the ones below it on both. Then we compared the groups input by input. Two results matter. The first is the big gap: the single widest difference between the groups was how far an error spreads. Work worth doing is work whose mistakes land in other people's queues, other departments, other months. The second is the flat line: how hard a workflow is to implement scored identically in both groups. The work nobody has automated is not the hard work. It is the work nobody has noticed.

This kills the most common planning mistake an agency owner makes: picking the project by how easy it looks. Ease told us nothing about value. Two other assumptions died with it. Software platform variation also scored the same in both groups, so waiting for a system upgrade is not a reason to wait. And AI competition actually ran slightly higher among the workflows worth doing, meaning crowded tooling is a sign the problem is real, not a reason to stay out.

What does separate the groups, in order: error spread, then revenue lost when the work fails, then how often the work happens, then whether the customer will actually pay for the fix. Those four inputs are the whole story. If a workflow's mistakes stay in one person's drawer, leave it alone regardless of how painful it feels.

The work nobody has automated is not the hard work. It is the work nobody has noticed.

Start with the cash file, not the biggest problem

*The most depended-on workflow in the study is a middling performer on paper.
Do it first anyway.*

The quiet foundation

Electronic remittance ingestion ranks thirty-ninth of one hundred and sixty-nine on value, but a chain of other workflows, from underpayment recovery to period revenue integrity, cannot run properly until it works.

Value and dependency disagree

Of the workflows other workflows are built on, roughly three quarters rank well below where their dependents rank. The order to build in is not the order of prizes.

Ranking workflows by value gives you one list. Asking which workflows other workflows are built on gives you a different one. The second list is the spine, and the most useful thing in this study is where the two disagree.

The clearest case: electronic remittance ingestion, the job of pulling in the payment files payers send and posting them line by line. On raw value it sits in the middle of the pack, notable for nothing. On dependencies it towers over everything: a long chain of other jobs, from underpayment recovery to denial routing to period revenue integrity, cannot work properly until this one does.

An operator reading only the value ranking would start somewhere flashier and quietly starve their own project. Every downstream fix a vendor promises, from spotting underpayments to aging the receivables correctly, assumes the cash file is in the machine and the payments are posted line by line. Post it by hand for the small payers, as most agencies still do, and the downstream fixes are decorations.

The same inversion shows up in eligibility. The batch eligibility check, the job of asking payers whether the patient in front of you is covered, is another middle-of-the-pack performer on value and another foundation on

the dependency list. Most of the start-of-care work, from visit verification to authorizations, presumes the coverage answer is in the episode record already. Neither job will win you a demo. Both will make every later job cheaper to build.

Every downstream fix a vendor promises assumes the cash file is already in the machine. Post it by hand and the downstream fixes are decorations.

Take the documentation fights once the plumbing works

The highest-value work left on the table is hard for honest reasons: the data is messy and the rules bend to the patient in front of you.

Money provably on the table

Customer willingness to pay scored the maximum, five out of five, for eligibility documentation readiness, denial rebuttals, audit responses and appeal letters alike.

Held back by the same two things

Data maturity and implementation complexity each scored three out of five for nearly every hard, high-value workflow in the study. The knowledge is codifiable; the records are scattered.

Some of the highest-value workflows in the study are ones we would not start first. They are the documentation fights: proving a patient is homebound before the claim, drafting the rebuttal when a payer says the visit was not medically necessary, assembling the package a review contractor demands inside its deadline. Each one carries top marks on revenue at stake, how often the work happens, and how far mistakes travel. Each one is held back by the same two things: the agency's records are scattered across systems, and the knowledge lives in the heads of a few senior people.

Consider what breaks today. The clinical manager writes denial rebuttals when there is time, and many age out of the appeal window unwritten. Appeal letters take a senior biller and a clinical manager the better part of an hour, together. Audit responses are assembled by hand from four systems over two weeks, and a late response is an automatic loss. These are not exotic problems. They are the agency's best people doing assembly work while the deadline runs.

The reason to take these fights is that the money is provably there and the buyers know it: customer willingness to pay scored at the top for every one of them. The reason to take them second is sequencing. The rebuttal needs the record in one place and the denial routed correctly first. Build the routing, the eligibility spine and the cash file, and these fights stop being heroic and become a Tuesday.

One caution belongs here. Nobody should point a model at hospice eligibility judgment itself. Terminal prognosis bends to the patient in front of you, and a system that cannot see the patient will be confidently wrong. The winnable version is narrower and better: check the documentation against the criteria and the reviewer's checklist before the claim goes out, and hand the gaps back to the clinician while there is still time to fix them.

Build the routing, the eligibility spine and the cash file, and the documentation fights stop being heroic and become a Tuesday.

Say no to the shiny stuff, and mean it

The bottom of the study is full of things an agency owner would love to delegate.

The scores say the tool earns nothing there yet.

Painful, and still not worth it

The lowest-scoring workflows, from interface error handling to equipment tracking across branches, scored near the bottom on value, tractability, or both. Rare, low-stakes work earns nothing from automation.

Real value, wrong form

The acquisition integration playbook carries a platform-variation score of five out of five, the highest in the study, and almost nothing else. Valuable work in the wrong form for a model.

A ranking that cannot say where the tool does not help is a brochure. The bottom of our study is instructive because it is full of real pain: interface errors between the agency system and everything else, onboarding and competency records, equipment idle at one branch while another waits for kits, the acquisition playbook that turns every purchase into a project. None of it scored worth doing now, and the reasons differ in a way that teaches the whole method.

The enterprise back office fails on value, not difficulty. Training records, reputation management, partnership programs: the mistakes stay local, the frequency is low, and the labor saved does not buy the build. The acquisition playbook fails on tractability: it is genuinely valuable, but every acquisition is different, the data is a mess by definition, and the workflow is rare enough that no learning accumulates. Same verdict, opposite mechanism.

That is the discipline worth importing to every decision. Ask how far a mistake travels, how much revenue rides on it, how often it happens, and whether the customer will pay. When all four answers are small or rare, AI is the wrong tool no matter how loud the pain. When three are large and one is stuck on scattered data, that is the next project after the plumbing works.

When a mistake stays in one drawer and the work happens twice a year, AI is the wrong tool no matter how loud the pain.

The order to do it in

The ranking says what is worth doing. This says what has to work first, and they are not the same list.

1 **Wire the remittance files into the system and post payments line by line**

It is a middling performer on value and the single most depended-on workflow in the study. Underpayment recovery, receivables aging, denial follow-up and period revenue checks all presume it is done. Doing it first makes every later project cheaper, and doing it late quietly starves them.

Ranks thirty-ninth of one hundred and sixty-nine; more downstream work depends on it than on anything else.

2 **Automate the eligibility check and write the answer into the episode**

The batch coverage check is the other quiet foundation. Most start-of-care work, from visit verification to authorizations, assumes the coverage answer is already in the record. The data is standard, the rules are published, and the exceptions are nameable, which makes it the most tractable kind of automation there is.

Data maturity and knowledge codifiability both scored the maximum, five out of five.

3 Route every denial against a payer-rule set before anyone reads it

Today one or two senior billers decide the path for every denial in the agency, which is why many die unread in a queue. Routing is where the appeal, rebill, correct or write-off decision gets made, and the appeal window is short. Route the denial, show the coverage rule that fired, and send the low-confidence cases to a person.

One of the highest-value workflows in the study; a handful of other denial jobs cannot run until routing works.

4 Let the model draft the rebuttals and appeal letters, with citations

Once denials are routed, the drafting work is assembly from the record: coverage criteria, the assessment, the notes, the decline evidence. A senior biller and a clinical manager currently spend the better part of an hour per letter, and unwritten letters age out of the window. The model drafts, the humans approve and sign, and the agency stops losing appeals to the calendar.

Revenue at stake and human skill required both scored the maximum, five out of five.

5 Then take the documentation fights, starting with eligibility documentation readiness

Proving homebound status, skilled need and terminal prognosis before the claim is the single biggest prize the study left off the immediate list: top marks on revenue at stake, frequency and error spread, held back only by scattered records. With routing, eligibility data and the cash file in place, the assembly stops being heroic. Keep the judgment with the clinician; automate the checklist.

Opportunity score of 84.2 of 100 against a readiness score of 57.1, the widest value-versus-readiness gap worth closing.

6 **Leave the enterprise back office alone for now**

Training records, reputation management, equipment tracking, partnership programs and acquisition playbooks scored at the bottom for reasons the middle sections explain: mistakes stay local, work is rare, or every instance is different. Revisit when the revenue-cycle spine is running, not before.

Every one of the six lowest-ranked workflows carries the same watch verdict.

Where AI earns least

Worth automating and cheap to automate are different claims. These fail the first test, and saying so is what makes the rest of this worth reading.

System interfaces and integration monitoring

Interface errors between the agency system and its payers are real pain, but the work is rare, the fixes are bespoke to each pair of systems, and a mistake

here does not travel into other people's jobs. The build costs more than the labor it saves.

Onboarding, competency and training records

The records matter at audit time and not before, the frequency is low, and the errors stay in one drawer. Nobody will pay for a fix to a problem that surfaces once a year, whatever the audit anxiety feels like.

Acquisition integration

Genuinely valuable and almost perfectly untractable. Every acquisition is a different mess of records, enrollments and open episodes, the thirty-six-month transfer rule binds, and the work happens too rarely for any learning to accumulate. Project management with a checklist, not a model.

Equipment and fleet tracking across branches

Idle equipment is visible waste, but the data to see it barely exists, the savings are operational rather than revenue, and the mistake travels nowhere. A spreadsheet with an owner beats a model here.

Partnership and channel programs

Health system and value-based partnerships are strategy work done a handful of times a year by senior people. The bottleneck is trust and negotiation, neither of which a model speeds up. Track the pipeline; do not automate the relationship.

What this cannot tell you

This edition was scored by analysts. No operator panel sat in the room, so a working billing lead might reorder the middle of the ranking, though the top and bottom are unlikely to move.

Readiness scores reflect how tractable the work looked on paper, including our own assumptions about what an agency's data looks like. A specific agency with unusually clean records will find some watch-quadrant work easier than we judged, and some easy-looking work harder.

The dependency spine is inferred from how the workflows logically connect, not measured in any real agency. Where we say a job opens the way for others, that is the structure of the work, not a field observation.

Scores are a snapshot of one edition. Payer rules, review programs and payment rates shift every year, and the value of a workflow moves with them.

SOURCES

1. [eCFR \(Office of the Federal Register and GPO\), 42 CFR 484.205\(j\)\(3\)\(ii\) \(trigger at \(j\)\(3\) intro; exceptions at \(j\)\(4\)\)](#), checked 2026-09-21
2. [eCFR \(Office of the Federal Register\), 42 CFR 418.24\(e\)\(1\)](#), checked 2026-09-21
3. [Office of the Federal Register \(eCFR\), 42 CFR 484.205\(j\)\(1\)](#), checked 2026-09-21

LICENCE

Published by Moative Research under [CC BY 4.0](#).

Reuse it, including commercially, with attribution.

HOW TO CHECK IT

Every regulatory fact links to the primary source it was checked against. The scoring method, its inputs and its limits are on the [method page](#), including the disclosure of who paid for this and what was removed from the scores because of it.

CORRECTIONS

Write to research@moative.com and it will be fixed with the change noted.