One Bot, One Job

A chatbot answers when you ask, whereas an agent does a job. You already manage human agents, like an MA who rooms patients, a biller who works denials, a front desk that confirms tomorrow's appointments. Each has one job, a thing to produce, and a rule for when to get the doctor. An AI agent is the same shape and everything here is free or nearly free and can quickly and easily be set up today.

WHERE AI CAN HELP

The Monday review pile. Online reviews of your doctors and staff pile up all week on Google and Yelp. Nobody reads them until one goes tragically viral, and by then the staff member it names has heard about it from a patient.

The thank-you that never gets written. You meant to thank the PCP who sent you three patients last month. It is now the month after.

The person who does everything. Your best front desk person also answers phones, chases billing, and orders supplies. Each job gets done well enough, none of them great, and when she is out, all of it stops. A bot works the same way. One bot, one job is how you get one you can trust and stop thinking about.

THE THREE PARTS OF ANY AGENT

1. The job. One sentence, including where it looks. "Every Monday at 7am, summarize last week's Google reviews."

2. What it produces. One paragraph in your inbox, a row in a spreadsheet, a draft reply.

3. When it stops and asks. "If a review names a staff member or mentions safety, do not summarize, flag it to me." Write this part first. If you cannot say when it should stop and come get you, the job is not ready to hand off.

PASTE THIS IN, THEN FILL IN YOUR INFO

PASTE INTO CHATGPT OR CLAUDE

"I want to hand off one job to AI. **The job:** [what it does, when, and where it looks, for example, every Monday read my Google reviews]. **What I want back:** [for example, one paragraph in my inbox]. **Stop and ask me if:** [for example, a review names a staff member]. Write me step-by-step instructions I can reuse every week, the way I would train a new employee on day one. Then tell me if I should set this up as a scheduled task in Claude or ChatGPT, or just paste it in myself each week."

No money, no patient data, nothing sent to patients. Agents draft, humans send.

FIVE AGENTS TO BUILD TODAY

1 **The review reader.**
Every Monday, summarize last week's Google and Yelp reviews. Flag any that name staff or mention safety. Agent output is one paragraph in your inbox.
15 minutes

2 **The referral thank-you drafter.**
Every Friday, give your agent the list of referring doctors for the week. It will draft one warm, specific note per name. It gives you drafts to approve and send.
15 minutes

3 **The inbox triage bot.**
Every morning, it reads one email folder you pick and sorts them into three short lists: needs you today, can wait, and ignore.
10 minutes

4 **The meeting prep bot.**
The night before any calendar event tagged "meeting," pull the last three emails with that person and write a five-line brief. It is in your inbox by 7pm.
15 minutes

5 **The website watcher.**
Once a week, it checks two or three web pages you care about, like your state medical board's news page or the American College of Rheumatology's guidelines page, and emails you a short note in plain English on what changed. If nothing changed, the note says so.
10 minutes

This handout is general information, not legal or compliance advice. No patient information enters any tool without a signed BAA.

DEEPER DIVE

Why one job works, and how to keep a bot honest

Know why a single-job agent is safe to leave alone, and you will know which job to hand off next without anyone handing you a list.

1. Why one job ONE PLACE TO FAIL

An agent with one job has one place to fail, so you can see the failure. An agent with five jobs fails somewhere in the middle and you find out three weeks later. This is the same reason a practice does not hire one person to be the biller, the MA, and the front desk. Not because the person is not smart. Because no one can check their work. The vote in the room, one job or two, is the whole skill.

2. Why "stops and asks" matters most THE ESCALATION RULE

The third part is what makes an agent safe to leave alone. A new hire with no rule for when to get the doctor will either interrupt constantly or, worse, decide on their own. Write the rule before the job. A review that names a staff member, a vendor email asking for money, a message that mentions a patient by name, each of those is a stop. The agent's output on a stop is a flag, never an action.

3. Agents draft, humans send THE SEND BUTTON

An agent that sends is an agent that can send the wrong thing to the wrong patient at 3am. An agent that drafts costs you a click and never costs you a complaint. Keep the send button human for anything that leaves the practice. The same goes for money and for anything that touches a chart. Nothing on this page needs patient information, and nothing on it should ever get any.

4. Every agent has a boss TWO WEEKS OF READING

You read what it produces for the first two weeks before you trust it. Keep a note of every miss. Most misses are a vague first part or a missing third part. Fix the parts, not the bot. After four weeks of right answers in a row, move it from a prompt you run by hand to a task on a clock.

aiforprivatepractice.com/rppa

5. Where agents actually run today TWO HOMES

A saved prompt you run by hand each week. Not really an agent, but the same three parts, and where most people should start until the job is stable. **A scheduled task in Claude or ChatGPT.** Runs on a clock, reads what you connect, produces the output, and stops. Paid plan required, with the by-hand version as the fallback. Start with the first, and graduate to the second once the output has been right four weeks in a row.

```
Every Monday at 7am, read last week's reviews
from [SOURCE]. Write one paragraph for me.
Name the themes, count the stars, and quote
the single most useful line. If any review
names a staff member or mentions safety, do
not summarize it. Put it at the top under the
heading NEEDS YOU and stop.
```

What an agent is not

It is not a second brain, a digital twin, or an employee replacement. It is a very fast, very literal new hire who reads only what you point at and does exactly what the three parts say. That literalness is the feature. It is also why the three parts matter. Vague instructions produce confident nonsense on a schedule.

When it goes wrong

If the bot sent something it should have drafted, you skipped rule three. Pull its send permission, keep the draft permission, and apologize like a human. **If the output is right but useless,** the first part is too broad. Narrow the job until the output changes what you do on Monday. **If it stopped and asked too often,** good, that is the safe failure. Loosen the third part one clause at a time, never all at once.

Sources worth your time

The Claude and ChatGPT help pages on scheduled tasks · Your own job descriptions, since every one of them already has these three parts · The Self-Defense handout from this same workshop for the practice AI policy that governs all of this.