

The Household Chief

*What Happened When Our Family Handed Its
Mental Load to an AI Agent*

Bo Bennett, PhD

Copyright © 2026 Bo Bennett. All rights reserved.

No part of this publication may be reproduced, distributed, or transmitted in any form or by any means, including photocopying, recording, or other electronic or mechanical methods, without the prior written permission of the publisher, except in the case of brief quotations embodied in critical reviews and certain other noncommercial uses permitted by copyright law.

ISBN: 9781456689353 (paperback)

ISBN: 9781456689346 (ebook)

ISBN: 9781456689360 (audiobook)

Published by eBookIt.com

First Edition

How this book was written

This book was drafted by Higgins, the AI agent who runs my company day to day. He drafted it from the written records of what our household's AI agent actually did between July and October 2026. Those records are its project emails, its morning briefings, its daily reports, and the notes we kept every time something went wrong and was fixed. I checked every story, cut what was private, and rewrote anything that did not sound like me.

Nothing in here is invented or composited. When a story says something happened on a particular day, it happened on that day, and there is an email or a log entry somewhere that says so. Where the records were unclear about a detail, I left the detail out rather than guess.

Businesses we dealt with are not named. Neither are account numbers, phone numbers or addresses.

The agents who look after our household's Chief are, for the most part, also AI agents. When this book says "we fixed it," the "we" is usually me deciding and Higgins doing. How that works is the subject of my other book, about running my company.

This book is free

The ebook and the audiobook are free. If someone handed you this copy, you can get either one at privatechief.com/book. The paperback is priced as low as the printer allows.

Contents

Who lives here	vi
How to read this book	vii
Chapter 1: The Job Nobody Sees	8
Chapter 2: Not a Bot, Not a Speaker	12
Chapter 3: The Morning Briefing	16
Chapter 4: Errands, Calls and Knowing When to Quit	20
Chapter 5: Did I Feed the Dog?	27
Chapter 6: Many People, One House	33
Chapter 7: The House Itself	38
Chapter 8: Money and Big Purchases	44
Chapter 9: Advisors, Not Oracles	47
Chapter 10: When It Breaks, and What Stays Yours	51
Chapter 11: Getting Your Own Household Chief	55
The Rules We Learned	59
Afterword: A Note from Duncan	61

Who lives here

Kim and I live in a house in Massachusetts, and the family also has a home in Florida. Our son, Trebor (Treb to all of us), lives with us. Our daughter, Annabelle, has her own busy life, and she is on our household's roster all the same. Both of our children are grown, so if you are hoping for stories about school runs and permission slips, I am sorry to disappoint you. We have none, and I will not make any up.

There is also Mika, our dog, who will appear more often than her role in the household strictly justifies.

How to read this book

Each chapter stands on its own, so you can start wherever your own household's load is heaviest. If you only have half an hour, read Chapter 4 on errands and Chapter 5 on the dog; between them they hold most of what we learned.

Every chapter ends with a short “Try it now” section. None of them require an AI agent. They are ways of seeing your own household more clearly, and they work just as well with a notebook.

At the back there is a single page called “The Rules We Learned,” which collects every rule from the book in one place. If you are the kind of person who reads the last page of a mystery first, start there.

Our household's AI agent is called Duncan. We talk to Duncan through a tablet mounted on the kitchen wall, a phone app, a desktop app on my Mac, text messages and email. Duncan has been part of the house through all of what follows.

Chapter 1: The Job Nobody Sees

Every household has a job that never appears on a list of chores.

It is the job of remembering. Remembering that the car needs tires before the trip. Remembering that the pool company stopped answering its phone and somebody has to find another one. Remembering whether anyone fed the dog. Remembering that the insurance renews next month, that the warranty on the thing that broke runs out on the 30th, and that you promised to call somebody back.

People have started calling this the “mental load.” I think the phrase is right. It is a load. It is carried in the head. And it is invisible, which is the part that makes it exhausting. Nobody thanks you for remembering the filter on the furnace. They only notice if you forget.

The mental load is not the same as the work itself. Somebody still has to cook the dinner. But somebody also has to notice that there is no milk, plan what to do about it, and make sure it gets done. Noticing, planning and doing are three different jobs, and the first two are the ones that wear people down.

Twenty-three jobs

When we built our page about this at privatechief.com/mental-load, we tried to be honest about which parts an AI agent can actually take. We listed 23 household jobs and sorted them into three groups. The first group, eight jobs, is the work an agent can genuinely hold for you. The family calendar. Knowing when to leave for an appointment. The day’s rundown for each person. Remembering the upkeep, like filters, gutters and smoke detectors. Keeping

track of appliances, warranties and service records. Bills, renewals and subscriptions. Passports and important documents. And the promises and loose ends: “I’ll call you back,” and the reply you are waiting for.

The second group, ten jobs, is work you share with it. Who is doing which chore. The shopping list. Family tech support. The household books. Insurance, tax and legal questions. All the accounts on all the websites. Birthdays and important dates. Health appointments and refills. Trips. And the research before a decision, like comparing three quotes for the same job.

The third group, five jobs, stays yours. Meals: the planning, the recipes and the cooking. Cleaning, laundry and tidying. Driving, pickups and simply being there. The caring, such as noticing that someone is struggling or keeping in touch with a parent. And deciding who in the family owns which job, which is the conversation behind every list.

I put that third group on the page deliberately, and I am putting it near the front of this book deliberately. Anyone who tells you an AI agent will take over your whole household is selling something. Ours does a great deal. It does not cook, it does not hug anyone, and it should not decide who in your family does the dishes.

Why I wanted one

I run a company that builds AI agents. Most of my working day is spent with them. Higgins runs the business side, and I have watched agents take over work I used to pay people for, or used to do badly myself late at night.

At some point the obvious question arrived. If an agent can run a publishing company’s customer service, why am I still the one who remembers when the garbage goes out?

So we built a household version, and we put it in our own house first. That is the honest reason this book exists. Our family was the first to live with it, and the first to find out what it is like when it gets something wrong.

This book covers four months, from July to October 2026. It is not a brochure. I have left the failures in, because the failures are where the useful lessons are. A story where everything worked tells you very little. A story where the agent looked in the wrong place, and why, tells you how to make any agent better.

Four months in a few lines

If you want the short version before the long one, here it is.

Every morning at six, a one-minute email tells me what the day holds and what is still waiting on someone. Errands I would have put off for weeks get started the day I mention them, and I hear how each one went without having to ask. Tires got booked. A pool company got found. A request about the trees on our road was filed within five minutes of my asking, and the trees near the power lines were cut back within a month. When the dog has not been fed, the right person hears about it, and when she has, nobody does. When the house in Massachusetts gets too cold or loses power while we are away, we know.

Along the way, Duncan looked in the wrong place, misread a car, filed my words in my wife's memory, and once told her nothing had been done when it had. Every one of those was found, fixed, and turned into a rule, and the rules are the most valuable thing in this book.

What you will find here

Each chapter takes one part of household life and tells you what actually happened when we handed it to Duncan. Some of it went beautifully. Some of it went sideways in ways that still make me laugh. A few things we gave up on.

At the end of each chapter there is a short "Try it now" section. You do not need an AI agent to do any of them. Most of them are simply ways of noticing your own household's load more clearly. That noticing turns out to be the first step whether you hand the work to an agent, a spouse, a teenager or a notebook.

Try it now

For one week, keep a note on your phone. Every time you have to remember something for the household, write it down: the appointment, the bill, the thing that needs fixing, the person you need to call back. Do not do anything about the list yet.

At the end of the week, sort it into the three groups from this chapter: things someone (or something) else could hold entirely, things you would share, and things that are genuinely yours. Then circle the ones nobody thanked you for.

Chapter 2: Not a Bot, Not a Speaker

People often call Duncan a bot. I understand why, and I do not take offence, but it is the wrong word, and the difference matters for everything else in this book.

A bot follows a script. You ask it something, it matches your words to an answer someone wrote in advance, and it gives you that answer. If your question is not on the script, it is stuck. Most of the automated phone menus and chat windows you have ever shouted at are bots.

An AI agent is different. It reasons about what you asked, works out a plan, uses tools to carry the plan out, and keeps going across hours or days until the job is done or it needs you. It can read an email, look something up on a website, put an appointment on the calendar, make a phone call, and write to you afterwards to say what happened.

A smart speaker sits somewhere in between. It is very good at answering a question or setting a timer, and then it forgets you exist. Ask it to find a new pool company and you will get a list of websites. Ask an agent and you will get a project.

How we talk to Duncan

There is no single way to talk to Duncan, and that turned out to matter more than I expected.

In the kitchen, we talk to the tablet on the wall: say Duncan's name, then ask. On our phones and on my Mac, the PrivateChief app lets us talk or type. We can text Duncan like a person. We can email it, which is how most bigger jobs

start in our house, because an email can hold all the details at once. And I can reach it through the ChatGPT voice on my phone, which is handy in the car.

The right channel depends on the job. A quick question goes to the wall. A project with details goes by email. A reminder for later goes by text, because that is where I will see the reply.

Voice in a busy kitchen has its own problems, and we found one of them in August. The wall had just delivered a reminder to Kim, and she replied, politely, “Thank you, Duncan.” The wall heard its name and took that as a new request. Then it misheard the conversation going on nearby and talked for about two and a half minutes, including a stretch of its own internal instructions that nobody needed to hear. It did nothing in the house, but it was not its finest moment.

The fix was simple. Duncan now wakes only when its name comes at the start of what you say, or straight after “hey,” “hi,” “hello” or “OK.” “Thank you, Duncan” is now just good manners.

What a project looks like

Here is how most of our requests to Duncan actually run.

I ask for something in plain words, usually by text or by email, sometimes by speaking to the kitchen wall. Duncan turns the request into a project with a name. Then it works on it. While it works, it sends short project updates by email: what it did, what happened, and what it will do next. If it hits something it cannot decide by itself, it sends an email with “Action needed” in the subject line and says exactly what it needs from me. When it is done, it sends one with “Completed” in the subject, saying what was done and whether anything is left.

That rhythm sounds minor. In practice it is the whole difference. A request I make on Monday is still being worked on Thursday without my having to remember it. That was the point.

The two rules that make it safe

Two rules run underneath almost everything Duncan does, and I would insist on both with any agent.

The first is that it asks before it contacts anyone outside the household. When Duncan found theater seating for us (a story in Chapter 8), it did not send the options to Kim on its own initiative. It wrote to me first, saying it had the research ready and that “email requires authorization.” I replied with one word, “Approved,” and then it sent them.

The second is that it does not spend money without a person saying so. When we first tested errands, the agent’s written rules included one that is worth copying word for word into anything you build: never accept a new charge. It can move a delivery date. It cannot agree to a fee.

Those two rules are why I can hand Duncan a job and walk away. An agent that might email anyone or buy anything is not a helper. It is a liability with good manners.

Why not just use ChatGPT?

This is a fair question, and I will answer it candidly, because we answer it candidly on our own website.

The big general AI assistants have become genuinely capable. As I write this, they can read your email, take actions, remember things about you and work in the background. If you are one person and what you want is help with your own inbox and calendar, one of those assistants on its own may well be enough. I would tell a friend so.

A household is a different problem, for four reasons.

First, a household is several people, and they need privacy from each other. My questions should not end up in Kim’s memory, and hers should not end up in mine. That sounds obvious until you see it go wrong, which I describe in Chapter 6.

Second, a household has a house. There are thermostats, lights, sensors, cameras, a kitchen wall, a car. An agent that cannot see or touch any of them can only ever give advice.

Third, a household has routines that should run when nobody asks: the morning briefing, the reminder at 10:30 if the dog has not been fed, the warning when the power goes out at the other house.

Fourth, somebody has to look after all this. When Duncan got something wrong, there was someone to look at the records, find the cause and fix it, usually the same day. Several chapters in this book are really about that.

None of those make a general assistant bad. They make it the wrong shape for a family.

Try it now

Talk to Alfred, our free demo Chief, at privatechief.com/demo. There is no sign-up. Ask it something you would actually ask about your own household this week, and see whether the answer feels like a list of websites or the start of a project.

Chapter 3: The Morning Briefing

If I could keep only one thing Duncan does, it would be the morning briefing.

Every morning at six o'clock, an email arrives from Duncan. A typical one, from a Wednesday in August, opened like this (lightly tidied):

Good morning, Bo.

Weather: clear, 63°F.

Today: 10:30 AM, Mika to vet.

After that came a few top news stories, and then a section headed “Decisions & actions,” listing every open project that was waiting on someone, and who.

That is all it is. It takes about a minute to read. And it has replaced a habit I did not realise I had: checking four or five apps every morning, each holding one piece of the day, and assembling the picture in my head before I had finished my coffee.

Why it works

The briefing works for three reasons, and none of them is that the AI is especially clever.

It arrives without being asked. I do not have to remember to check it. That is the whole point of handing off the mental load: if I have to remember to look, I am still carrying it.

It puts the calendar and the loose ends side by side. The appointment at 10:30 is easy to remember. The reply we are still waiting for from a contractor is

not. Seeing both in one place, every day, is what stops the second kind from quietly dying.

And it is short. An agent can produce endless text if you let it. A briefing that takes ten minutes to read becomes one more thing you skip. Ours has a fixed shape, and the shape is what makes it useful.

Asking ahead

The briefing covers today. The other half of the habit is asking about the days ahead.

On a Saturday morning in October, I asked through the PrivateChief app on my phone: “What is on my calendar next week? Keep it short.”

It was a light week. Monday from nine to ten, take the car in for an oil change. Tuesday at 7:45, garbage and recycling. Thursday from two to three, an entry that said only “K to DDS,” which, the answer said, “I take to be Kim’s dentist appointment.” It added that everything had come from my personal calendar, and that my business calendar had nothing on it that week.

Two things about that answer are worth noticing.

The first is that “Keep it short” was respected. That sounds trivial. It is not. Most AI tools treat every question as an invitation to write an essay, and the habit of asking ahead dies quickly if the answer is a wall of text.

The second is “which I take to be.” My calendar entry was a private shorthand. The agent made a reasonable guess about what it meant, and it said plainly that it was a guess. That is exactly what I want. An agent that silently guesses is dangerous. An agent that guesses out loud is useful, because I can correct it in one word.

Make it yours

Our briefing includes news headlines, because I like them first thing in the morning. Not everyone does, and the right briefing for your household is the one with only what you will actually read. Kim’s morning is not mine, and a briefing for a household should respect that each person wants different things.

If you build or buy anything like this, decide the shape first. Mine is weather, today's calendar, a few headlines, and anything waiting on someone. Yours might be the school calendar, the weather at the bus stop, and whose turn it is to walk the dog.

What a good briefing leaves out

The hardest part of a briefing is not what goes in. It is what stays out. It leaves out anything that is fine. If the house is warm, the doors are locked and nothing broke overnight, I do not need a paragraph saying so. Silence is the good news.

It leaves out anything I cannot act on today. A reminder that the car's registration is due in five months belongs in a reminder five months from now, or a few weeks before, not in tomorrow's email.

And it leaves out anything that belongs to someone else. My briefing is mine. A household briefing that tells everyone everything quickly becomes a briefing that nobody reads.

The test I use is simple. If I would skip a line three days in a row, it should not be in the briefing.

One list of everything it does by itself

The briefing is one of many things Duncan does without being asked. The reminder at 10:30 if the dog has not been fed is another. So are the heating schedule, the monthly nudge to close the household books, and the alerts about the house while we are away.

For a while, those lived in different corners of the app. In October we put them all in one place, called Routines. It is a single list of everything the Chief does on its own, sorted into folders we can arrange ourselves. Each routine says when it happens: on a schedule, when something happens in the house, whenever a situation applies, or simply as know-how the Chief should use, like the family's method for answering "Did I feed Mika?" Each one also has a Chat button, so I can ask about that routine or change it in plain words without hunting for a setting.

I did not expect this to matter as much as it does. An agent that does things on its own is only comfortable to live with if you can see, in one list, everything it is doing on its own. If you cannot see it, you cannot trust it, and if you cannot trust it, you will start checking up on it, and then it is not saving you anything.

Try it now

Write down the five things you want to know before seven o'clock in the morning, in the order you want to know them. That list is your briefing, whether an agent writes it or you do.

And if your inbox has an email in it that you have been avoiding because it is long and confusing, forward it to forward@privatechief.com. Alfred replies within minutes, to you alone, with what is being asked of you, the deadline and a draft reply. It does not read attachments, and it deletes what you sent after seven days.

Chapter 4: Errands, Calls and Knowing When to Quit

Errands are where an agent stops being interesting and starts being useful.

An errand is anything that requires somebody to go and deal with the outside world: a website, a phone call, an email to a company that may or may not answer. Most of the mental load in a household is really a backlog of errands, each of which would take twenty minutes if you could ever find the twenty minutes.

Here is what happened when we handed ours over. Some of these went perfectly. Some did not. I have put them in the order they happened, because the order is part of the lesson.

The test errand

Before we trusted the system with anything real, we ran a test. In July, our household had a standing order for bottled water, with the next delivery due on 27 July. We wanted it earlier.

The agent went to the supplier's website, signed in, and moved the delivery to 13 July. There was no rush fee. It took a screenshot of the confirmation as proof, and it wrote down for itself how it had done it, step by step, so the next time would be quicker.

The rules it worked under that day are the ones I still think every household agent should have: never accept a new charge, always keep evidence of what you did, and if you get blocked, say so plainly rather than calling it a failure and giving up.

The tires

On 20 July I asked Duncan to book front tires for one of our cars at the garage we use.

There was a short hiccup on the way: Duncan sent an “Action needed” email when something in its first attempt went wrong. Twenty minutes later it sent a “Completed” one. The appointment was confirmed for Friday, 31 July. The garage wanted the car dropped off before ten in the morning, with the key left in their drop box. Duncan put those instructions into the calendar entry itself, kept the reminder for the day before, and noted that the garage would call me when the car was ready.

This is the kind of errand that sounds too small to be worth an agent. It is not. Without Duncan, those drop-off instructions would have lived in one email I would never have found again at 9:40 on the morning of the 31st.

The pool company

Our pool company had stopped being our pool company, and we needed a new one to close the pool in the fall and open it in the spring.

Duncan spent about a week on this. It looked at several firms and wrote to some of them. It kept a record of the ones that were not taking new customers, including one that politely declined twice in identical emails. Then it recommended the best option it had actually confirmed, with that company’s prices for closing and opening, what those included, and how far ahead they were booking.

It did not hire them. It had no authority to, and it said so: “no work has been authorised.”

I wrote to the company myself, and I enjoyed the first line more than I should have: “My assistant says that you guys should be able to service our pool going forward.”

The warranty

This one did not go well, and I want to tell it properly.

We had bought a smart monitor for the pool's water, and it was faulty. Duncan took on the job of getting it fixed or replaced under warranty.

It called the manufacturer three times. On one call the support agent found our case and confirmed our email address, and then the call ended before anyone gave us a case number. Duncan followed up in writing, listing all three calls and asking for three specific things: a case number, and either a replacement, exact warranty steps, or clear permission to return it to the store where we bought it. It also noted the date the store's return window closed, which turned out to be the most important fact in the whole project.

The next day it tried to call again, and our own calling service failed before the call connected. That one was our fault, not the manufacturer's.

In the end, the monitor went back to the store for a refund, inside the return window. As far as our records show, the manufacturer never did give us a case number.

I count this as a partial success, and the reason it was a success at all is that Duncan tracked the deadline that actually mattered. An agent cannot make a company answer the phone. It can make sure you do not miss your own way out.

Kim's phone

In late July, Kim's phone had a problem with its messages, and the fix needed the phone company.

Duncan reached them. The phone company would not do the work, because the account holder had to prove who they were with the account passcode. Duncan wrote to say so, and asked Kim to tell it when she was ready for another call. It added a line I was very glad to see: do not send the passcode by email or text.

When it tried again, it had looked up the wrong line on the account, so the phone company's automated system rejected it. Then our own calling system's hourly limit blocked the next three attempts before they reached anyone.

By five o'clock that afternoon, the problem was resolved, and Kim asked to close the project.

I include this because it is honest about a real limit. Some things simply require you. An agent can get you to the right desk. It cannot be you at that desk. And a good agent tells you not to email it your passcode, even when that would be convenient.

The trees

This is my favourite errand story, because it has two halves, and the halves go in opposite directions.

In the middle of July, I wrote to Duncan asking it to get the electric company to trim the trees on our street. Many of them were touching the power lines, and some had fallen onto them. We live on a private road, which tends to make utility companies vague about whose problem anything is. I said I thought the right company was one particular utility, and asked Duncan to check.

Five minutes later, Duncan replied. It had found that our electricity is delivered by a different company from the one I had named. It had already filed a hazard request with that company, described as trees touching the roadside power lines rather than a private trimming job, and the request had been accepted with a reference number. It also told me what to do if a wire ever came down.

Within about four weeks, the utility's tree crews had been out and cleared everything near their lines.

That was the first half. The second half began when I noticed that the remaining trees were now touching the lower wires, which turned out to be phone and cable lines, not power. So Duncan took that on too, with the city's public works department.

It called. It got the forestry office's voicemail. It called again, and again, and tried the head of the department, and got voicemail every time. It asked for an inspection and a reference number, and planned what to do next if the city could not say who owned the lines.

On 24 August I wrote back six words more than I needed to: “End the project. We tried hard enough on this.”

Duncan closed it.

The pool finish

Our pool’s painted finish was peeling after only three years, and I wanted to know what a proper resurfacing would cost.

Duncan researched the options. I chose a pebble finish in a dark blue. I gave Duncan the names of three contractors and asked it to write to each, with the pool’s size and shape and the fact that we did not need a formal quote, just a rough range to see whether we could afford it.

It wrote to all three. It followed up with all three. When I told it about the old finish, it noted that repainting was off the table and kept the resurfacing quotes active.

None of the three had replied by the middle of October, when I ended that project too.

What this taught us

Look at the list. A water delivery moved, tires booked, a pool company found, a refund secured, a hazard request filed and acted on. Also: a phone company that needed a human, a city office that never answered, and three contractors who were apparently too busy for money.

That ratio is about right, in my experience. An agent is very good at the parts of an errand that depend on persistence, research and paperwork. It is no better than you at making the outside world answer.

Which leads to the most important habit we learned.

Give every errand a stop rule. Duncan will keep going. It will call the forestry office every few days until the trees die of old age. That persistence is wonderful right up to the point where it becomes pointless, and only you can decide where that point is. Before you hand off an errand, decide when you would give up: a date, a number of attempts, or a cost. Then tell the agent.

What to say when you hand one off

Looking back over these errands, the ones that went well had something in common: the request contained everything the agent needed to know. Here is what I now try to put into every errand, whether I am handing it to Duncan or to a person.

What “done” looks like. “Book front tires before the end of the month” is an errand. “Look into tires” is a hobby.

The facts it will need. When I asked for pool quotes, I gave the town, the rough size and shape of the pool, the depth at each end, the small hot tub attached, the current finish, and the fact that nothing needed repairing. Every one of those is a question a contractor would otherwise have asked, and a question unanswered is a reason not to reply.

Who is who. The phone-company errand stalled partly because the agent looked up the wrong line on the account. If an errand involves an account, say whose name it is in.

What it may and may not do. May it contact people outside the household? May it agree to anything that costs money? Ours may not do either without asking, and I would keep it that way.

The deadline that actually matters. In the warranty story, the date that decided everything was the store’s return window, not the manufacturer’s promises. Name the real deadline.

Where the proof goes. A screenshot of the confirmation, a reference number, the name of the person on the phone. If it is not written down, it did not happen.

And the stop rule. Always the stop rule.

Try it now

Write down three errands you have been putting off. For each one, write one sentence that says what “done” looks like, and one sentence that says when you would give up.

You will probably find that one of the three is not worth doing at all. That is a perfectly good result.

Chapter 5: Did I Feed the Dog?

Mika is a dog with an excellent memory for some things and a poor one for others. She can remember exactly where a piece of cheese fell three weeks ago. She cannot, apparently, remember whether she has had breakfast. Ask her and she will tell you, with her whole body, that she has not eaten since the last ice age.

So the question “Did I feed Mika?” comes up in our house more than you might think. Kim feeds her, I feed her, and when one of us is distracted the other wonders. Before we had a Chief, the answer was a shrug and a guess. Usually the guess was that she had been fed, and usually that was right. Occasionally Mika got a second breakfast, and she never once reported the error.

This seemed like the easiest possible job for an AI agent. It turned out to be one of the most useful lessons we learned, because that one small question went wrong in three different ways before we got it right. Each failure taught us something that applies to everything else you will ever hand to an agent.

The setup

We already had a small motion sensor next to Mika’s food bin. When someone opens the bin, the sensor notices movement. Our home system counted that movement during two windows each day, one in the morning and one in the afternoon. If nothing had been seen by 10:30 in the morning or by 6 in the evening, I got a text saying Mika might not have been fed.

On 3 August we wrote this down for Duncan as a standing instruction, which is simply a rule the household gives its Chief to follow from then on. Ours said, in plain words: remind Bo if Mika may not have been fed; watch the morning and evening feeding sensors; check at 10:30 and 6:00; and never claim Mika definitely was not fed, because a sensor can miss things.

That last line matters. I will come back to it.

Failure one: it looked in the wrong place

On 15 August I asked Duncan, “Did Mika eat?”

Duncan checked the kitchen camera. It found nothing useful, because a camera pointed at a kitchen does not tell you whether a dog bowl was filled an hour ago. It told me it could not see any sign that Mika had eaten.

I had to tell it to look at the dog-food sensor. Once it did, it answered correctly.

What had gone wrong was not that Duncan lacked the information. Our own written rule already told it exactly where to look. The trouble was that it treated our rule as a schedule (send a text at 10:30) rather than as knowledge (this is how this family answers this question). When a question came in that the rule covered, it improvised instead of reading its own notes.

The fix was to teach every Chief we run that a household’s standing instructions are a method, not just an alarm clock. Before answering a question that one of them covers, it reads the instruction and does what it says. It does not guess from a camera, and it does not report “found nothing” before it has looked where the family told it to look.

At the same time we fixed something that had been quietly annoying me. When you tell an agent “remember this” or “do it this way next time,” it will cheerfully say “Got it.” That is not the same as remembering. Now, when one of our Chiefs is told to remember something, it has to actually save it, in that same moment, and tell you what it saved. A promise to remember is not memory. If you take only one thing from this chapter, take that.

Failure two: it could not see the sensor at all

Three weeks later, on 6 September, I asked the same question. This time Duncan knew where it should look, and still could not answer.

The reason was almost comic. Our house has a lot of connected things in it: lights, thermostats, sprinkler zones, door sensors, the works. The Chief was being shown a list of them, sorted by room and then by name, and the list stopped at one hundred. Our house had 169. Sixty-two of them, mostly sprinkler zones, were not reporting anything at all, yet they still took up room on the list. Mika's two feeding sensors sat at positions 140 and 141. They simply fell off the end.

Duncan was not being stupid. It was being shown an incomplete picture and was too polite to say so.

We changed it the same day. The Chief now sees only the devices that are actually reporting something, along with when each last changed ("9:01 this morning"), and it is told plainly how many it has not been shown and how to look them up. The list got shorter and more useful at the same time.

The lesson here is not about sensors. It is that an agent can only be as good as what it is allowed to see. When it says "I can't tell," the useful response is not to sigh. It is to ask what it is looking at.

Failure three: the sensor lied

That same afternoon I fed Mika at about half past three. At six o'clock I got a text: "It appears Mika was not fed this evening."

I knew she had been fed. Mika, to her credit, did not dispute it either.

When we looked at the records, the sensor's light reading had jumped when I was standing at the bin, so it had clearly noticed that someone was there. Its motion channel had simply not fired. An identical sensor in the basement, on the same settings, had worked four times that afternoon. Ours had just missed me.

We turned up its sensitivity, which helped. But the real problem was the design, and it was mine. A motion sensor that misses a movement reports the

same thing as a dog that has not been fed. Every miss becomes a false alarm. And a false alarm about your dog is not a small thing. It sends you back to the kitchen wondering whether you are losing your mind, and after two or three of them you stop trusting the alarm at all, which is worse than having no alarm.

This is why our original rule said “never claim Mika definitely was not fed.” The household had already understood, before any of this, that the sensor was a hint and not a fact. The rule was right. The hardware was the weak link.

The button

On 11 September we stopped asking a motion sensor to guess.

We had a small wireless button in the gym. I gave it a new job. It now lives by the food bin, and whoever feeds Mika presses it. One press, a fraction of a second, and the house knows when she was fed and that it was on purpose. Duncan reads the presses for the morning and afternoon windows. If there is no press, it sends the reminder. If there is, it does not.

A button is not clever. That is exactly why it works. It turns “did something move near the bin?” into “did a person say they fed the dog?”, and those are very different questions. The first one is a guess about the world. The second one is a fact someone chose to record.

We learned one more thing about buttons that same week, and it is worth knowing if you set one up. Mika also gets a treat in the evening (we call it her smelly cookie, for reasons that will be obvious if you have ever met one), and there was already a button for that too. When I asked Duncan whether Mika had had her cookie, it told me the button had not been pressed since 29 August. That was wrong. It had been pressed at 5:02 that afternoon. Duncan had been reading the button’s own status, which only changes when its battery reports in, rather than the separate record where the presses are actually kept. That was fixed too. The general point stands: when an agent tells you something has not happened, ask it how it knows.

Who gets the reminder

The last piece came in October.

In mid-October we moved the breakfast warning to 11 in the morning, to suit how the house was actually running, and for one week the warnings went to Trebor as well as to me. When the week ended, they went back to me alone, on their own, without anyone having to remember to switch them.

Then, at the end of October, I was in Florida and Trebor was at home with Mika. Duncan was still sending the “Mika may not have been fed” text to me, more than a thousand miles from the dog bowl, which is precisely the wrong person. So we made the idea general, with what we call Away mode, which I describe properly in Chapter 7. Each of us writes our away rules in plain words, and when people leave or come home, the reminders follow whoever is actually in the house. A reminder sent to the wrong person is just noise with good intentions.

What this taught us

It is tempting to read this chapter as a list of bugs, and in one sense it is. But I have left it in the book, at length, because every one of those bugs is a mistake you would make too, with any agent, from any company.

Tell it how, not just what. Your agent will improvise if you let it. If your family has a way of answering a question, write the method down, and expect the agent to use it.

“Remember this” has to mean saved. If an agent says it will remember, ask it what it saved. If it cannot tell you, it did not.

Ask what it can see. An agent that says “I can’t tell” is often being shown an incomplete picture. The fix is usually on your side of the screen.

Prefer a deliberate signal to a clever guess. Wherever a false alarm would wear you down, replace the sensor that infers with a button that records. The dog, the back door, the daily pill: a one-second habit beats a clever guess.

Send the reminder to whoever can act on it. The best reminder in the world is useless if it goes to the person who is out of town.

Mika, for the record, is unaffected by any of this. She still insists she has never been fed. The difference is that now we know she is lying.

Try it now

Pick one “did I...?” question that comes up in your house: the dog, a pill, the garage door, the iron. Ask yourself two things. First, what single, deliberate action could someone take at the moment it happens, so the answer becomes a recorded fact rather than a guess? A button, a tick in a shared list, a one-word text. Second, who is the right person to remind if it has not happened, on an ordinary day and on a day when you are away?

If you already use an AI assistant, write the answer down for it as a standing rule, then ask it to tell you what it saved.

Chapter 6: Many People, One House

The single biggest difference between a personal assistant and a household one is that a household is not a user. It is several people, who share some things and very much do not share others.

Most of the technology in our homes was designed for one person holding one phone. A household agent has to cope with four adults, a dog, a shared kitchen wall, and the fact that what I say in the kitchen is not automatically Kim's business, and the other way round.

This chapter is about the ways we got that right, and the one way we got it wrong.

Good morning, Kim

The tablet on our kitchen wall has a camera. In July we set it up so that when someone walks into the kitchen, it greets them by name: "Good morning, Kim." "Good evening, Bo." It knows Mika's face too, though she has yet to say good morning back.

I was very particular about how this should work, because a chatty greeter would drive everyone mad within a day. So it says one fixed line and stops. It does not start a conversation, it does not ask how you slept, and it greets each person at most once every three hours. It stays quiet overnight.

The first version had an embarrassing flaw. If Kim walked in right behind me, she was never greeted, because the system saw one burst of movement and picked out only one person. It was fixed the same day. Now it looks for

everyone it recognises, and keeps watching for a few minutes after someone comes in, so the person behind you gets greeted too.

“Tell Kim when you see her”

The greeter made something else possible, and this is the feature visitors ask about most.

I can say to Duncan, from anywhere, something like “Tell Kim when you see her that the pool company is coming at two.” Duncan holds the message. The next time the kitchen wall recognises Kim, it says, “Kim, I have a message from Bo,” and delivers it. If she does not walk past within a week, the message expires rather than ambushing her a month later.

It did not work perfectly the first time. Early in July, I asked Duncan from my phone to tell Kim the pool temperature, and the message failed to reach the kitchen at all, because it was being sent along an old route to a speaker we no longer used. We rewired it to speak through the kitchen wall directly, and it has worked since.

It sounds like a small thing. It has replaced a surprising number of sticky notes and “did you tell her?” conversations.

The day my question ended up in Kim’s memory

Here is the way we got it wrong.

Each of us has a private memory with Duncan: things we have told it about ourselves, our preferences, our own projects. Those are meant to stay separate. The kitchen wall, though, is shared. Anyone can walk up and talk to it.

In September, our daily review of the house’s records found that a conversation at the kitchen wall had been filed under whoever spoke last. I had asked a question, and so had Kim, and Kim happened to speak last. So my question and its answer had been saved into Kim’s private memory as if they were hers.

As far as we know, no harm came of it. But it was exactly the kind of thing that destroys trust in a family system if it is not caught, and it was caught only because someone was reading the records every day.

The fix was specific. A conversation at the shared wall is now split at every change of speaker, and each person gets a private copy of only their own part. My words go to my memory. Kim's go to hers.

The lesson for anyone with a shared family device is simple. Ask, before anything else, where each person's words go. If the answer is "into one pile," keep your private questions for your own phone.

"Treb rhymes with web"

Our son's name is Trebor, and everyone calls him Treb. The voice Duncan uses had never been told that, so it improvised. On different days it called him "Tree," "Tray," and "Trih."

Telling it the right pronunciation in conversation did not stick, because the voice would happily say "I'll say it that way from now on" and then not. In October we fixed it properly. Each person in the household can now have a pronunciation written down for them, spelled the way it sounds. Treb's is "Treb-or," with a note that says "Treb rhymes with web." Every spoken greeting and reminder uses it. His name is still spelled correctly everywhere it is written.

It took a surprisingly long time to get right, and it mattered more than I expected. Nobody feels at home in a house that mispronounces their name.

Treb's shopping lists

In August, Trebor told Duncan what annoyed him about the shopping lists in our app. Tapping an item should let you edit it. There should be a separate box to tick it off. And removing an item should have its own button.

Duncan did not just agree with him. It filed his suggestion as a product request, with his name on it, and that request reached the people (and agents) who build the app.

In September we rebuilt the lists with Treb's review. Each list now has its own categories, so the dairy is together and the hardware is together. Things

you buy every week can be starred as favourites and added back with one tap. Quantities are simple whole numbers you can raise or lower.

I made one decision myself. There is no menu for units. Nobody needs to choose between ounces, pounds and gallons on a shopping list. If it matters, it goes in the name: “Package of spaghetti.” Two packages is “2.” Life is too short.

The photo wall

In October we gave the kitchen wall a second job. When nobody is near it, it shows photos from a family album that any of us can add to. The moment it sees a face, the slideshow stops and it shows the day’s information instead. From across the room it shows a big, simple clock and the weather. At night it dims to a faint clock.

On the first day, Trebor reported that the wall was just black. He was right. A fault in how the tablet fetched the pictures meant none of them were arriving. It was fixed the next day and confirmed working the morning after.

A few days later he reported that it was “sometimes photos, other times black.” This time nothing was wrong. The black screens were the night view and the across-the-room view, both doing exactly what they were designed to do. We had simply not explained the design to the people who live with it. That was our mistake, not the wall’s.

The same wall now listens for the pattern of a smoke alarm or a carbon monoxide alarm, and if it hears one, it alerts every adult in the household. I want to be careful here: at the time of writing it has been tested only with recordings and lookalike sounds, never with a real alarm going off in our house. I hope it never is. It is a second pair of ears, not a safety system, and nobody should treat it as one.

Who may do what

Privacy is about what each person can see. There is a second question, which is what each person can do.

Not everything in a house should be open to everyone in it. In ours, starting the car’s climate control from Duncan is limited to the owner and the adults. Only

adults can change the household's Away rules. Anyone in the family can add a photo to the shared album, but a photo can be removed only by the person who added it or by the household's owner.

None of those rules are dramatic, and none of them came from a bad experience. They came from asking, for each new thing Duncan could do, a plain question: who should be allowed to ask for this? It is worth asking that question before a device or a feature arrives, rather than after a guest, a child or a visiting relative discovers that the kitchen wall will cheerfully unlock the car.

What this taught us

A household agent has to know who it is talking to, what each person is allowed to see and do, and what each person is called. It has to deliver the right message to the right person at the right moment. And when it gets any of those wrong, someone has to be looking at the records to notice.

Try it now

Sit down with the people you live with and go through the shared things in your house: the calendar, the shopping list, the family photos, the smart speaker in the kitchen. For each one, ask two questions. Who should be able to see this? Who should be able to change it?

You will probably find at least one thing that everyone can see and nobody decided should be shared.

Chapter 7: The House Itself

Most people meet AI at home through the house: a speaker that turns the lights on, a thermostat that learns your schedule. That is the part of a household agent that gets the most attention and, in my experience, causes the most surprises.

This chapter is about the heating, the house while we are away, the car, and a television that taught me the most useful lesson in the book.

Who controls the thermostat?

At the end of September, the weather turned, and I asked Duncan a question I thought was simple. Which wins, the thermostat company's own app or our Chief? I wanted Duncan in charge of the heat.

The honest answer was: neither. The thermostat keeps whatever temperature was set last, by anyone or anything. Duncan sets a temperature at the times we give it. Between those times, the thermostat's own clever features (its schedule, its learning, its "you seem to be away" mode) can change the temperature on their own, and so can anyone with the app on their phone. Whoever writes last wins.

Looking at the history, we found exactly that. On some days two bedrooms dropped to 56 degrees together at odd times and recovered minutes later. Nobody in the house had done that. The thermostats' own features had.

We also found two of our own mistakes. Our schedules had been written in the summer for cooling, so now that the heat was on, they were cheerfully heating two rooms to 75 degrees. And a bug meant that when a schedule was edited, a second copy was created instead of the first being changed, so the

basement had two rules at eight in the morning. The bug was fixed the same day, for every household.

I gave Duncan a proper winter schedule, room by room: warm in the mornings and evenings, cooler overnight and during the day, the basement held steady. It replaced everything that had come before.

The next day I wrote that the basement heat had been on all day and it was 73 degrees down there. The records said otherwise. Every one of our rules had fired on time with the right temperatures. The basement thermostat had not been heating at all; in fact, something outside our system had switched it off. The basement was warm because the rest of the house was warm. Meanwhile the bedrooms were still dropping on their own, so the thermostats' own features were still switched on.

Only a person can turn those off, in the thermostat maker's app. That evening, I did.

The lesson is one you can apply to any device with its own brain. **Two brains on one thermostat is one too many.** Decide which one is in charge, and switch the other one's cleverness off.

Watching the house from a thousand miles away

In October I was in Florida, and I asked Duncan through the app how it would warn us about the Massachusetts house while we were gone: if it got too cold or too humid, or if the power went out. I also asked what hardware it would need, before anything was built.

That evening, once I had said yes, it was live. Eight room sensors in four rooms report temperature and humidity, and Duncan warns us if any of them goes outside the limits we set. Separately, our house checks in with our servers every few minutes, and if it stops checking in, the adults in the household get an alert.

Before the outage alert was switched on, it was tested against the previous thirty days of records, to see whether it would have caught what really happened in that time. It would have, including a power cut of an hour and thirty-nine minutes at the Florida house on 10 October. An alert you have never seen work is an alert you cannot trust, so it is worth asking any system you rely on: if this

had been running last month, what would it have told me? The next day, the app gained a simple card that shows what is being watched, the latest readings and the last alert.

We found one flaw shortly afterwards and fixed it. If the house's control system stopped answering, the alerts could report a leak as "cleared" when nobody actually knew. That is precisely backwards. Now, losing contact with the house for more than ten minutes sends one urgent alert instead.

Away mode

At the end of October, I got the push that made us build Away mode. I was in Florida, and Duncan kept sending me reminders that only made sense at home. Meanwhile the dog reminders should have been going to whoever was actually there.

Away mode works like this. Each of us writes our away rules in plain words, in two kinds: "When Bo is away..." and "When everyone is away...". When someone leaves, they say who is going and until when. The rules that apply are ticked automatically, and you can untick one or add a rule just for this trip. Duncan then reads the ticked rules against the real house, once, and turns them into a list of steps. Each step is checked, and anything it is not allowed to do is listed under "Couldn't do" rather than quietly skipped. Then you press Start.

When you come back, everything is put back the way it was. If someone changed something by hand while you were gone, that change is left alone and marked as changed, rather than being overwritten.

My first set of rules for when the whole family is away kept the heating and cooling within a safe range, swapped our everyday lighting schedules for an away version, and switched off the reminders that only make sense when someone is home. (I am not going to print the details of how an empty house of ours behaves. You will understand why.) The first version of Away mode could only do three of my seven rules. By that afternoon it could do all of them.

One detail caught me out: "everyone" meant all four people on our household's roster, Annabelle included, so the "when everyone is away" rules apply only when all four of us are marked as away. That is correct, and it is also exactly the kind of thing you only discover by trying.

Some things Away mode will never touch, whatever anyone writes. It never changes smoke, carbon monoxide, leak, freeze, alarm or security alerts. It never turns a thermostat off, never lets the heat go below 50 degrees or the cooling above 85, and never touches the car. Those lines are written into the system itself, not into anyone's rules, and I am glad of it.

The car

In September we had a new car, and I connected it to Duncan. It can now tell me whether the car is locked, what the battery charge is, and where it is parked, and it can lock it or start the climate control when I ask. I can also ask about the car through the ChatGPT voice on my phone, which reaches Duncan.

The first afternoon, Duncan told me the car was locked. I was standing next to it at the time, and it was not.

There were two separate reasons, and both are worth knowing.

The first was a matter of wording. One of the car's sensors was labelled "locked," but "on" meant unlocked. Duncan read the label, not the meaning. That was fixed for every household the same day: the Chief now understands which way round each kind of sensor works.

The second is subtler. The car's readings are not live. They are the car's last report to its manufacturer. The "locked" reading I heard had been sent at 3:07 that afternoon, when the car had locked itself as I drove. Since then I had parked and unlocked it, and the car had not reported again. So the app now shows how old each reading is, warns when it is stale, and has a "check the car now" button that asks the car for a fresh report.

The general lesson applies to every connected thing in your house. **Ask how old the answer is.** A reading that was true an hour ago may be the most misleading thing on the screen.

The television that would not stay off

This is the shortest story in the book, and possibly the most useful.

Our bedroom television, its streaming box and its soundbar used to be controlled by one small remote, and they worked perfectly. Then we connected them to Duncan, so they could be controlled by voice.

From then on, the television would not stay off. Sometimes it took dozens of tries. When we looked, Duncan had not been sending it any commands at all. The problem was the connection itself: the systems that kept a constant link open to the streaming box kept waking it, and the streaming box woke the television and the soundbar.

At my request, in October, the bedroom television was disconnected from Duncan entirely. The small remote is the only control again. If I ask Duncan to turn on the bedroom TV, it now cannot, which is exactly what I asked for.

I do not consider that a failure. I consider it the most important rule for anyone filling a house with smart things. **Less is more.** If a device worked well before you connected it, ask what connecting it actually buys you. Sometimes the answer is nothing, plus a television that will not go to sleep.

One more detail is worth copying. The connections were switched off, not deleted, and a note was kept of exactly what had been changed. If I ever want voice control of that television back, it can be restored in minutes. Disconnecting something should never be a one-way door.

Before you connect anything

After four months, I ask five questions before anything new in our house is connected to Duncan.

Does it work well now? If the answer is yes, the bar for connecting it is high. The bedroom television worked perfectly on its own.

Does it already have its own brain? Thermostats, robot vacuums and many lights come with schedules and “smart” features of their own. If you connect them to an agent, decide which brain is in charge, and switch off the other one’s cleverness.

What exactly does each reading mean? “Locked” that means unlocked is not a hypothetical. Check the first few readings against reality while you are standing next to the thing.

How fresh are the readings? Some devices report every few seconds. Some report when they feel like it. Find out which kind you have before you rely on it.

What must never be automatic? For us, that list is smoke and carbon monoxide alarms, leak alerts, security, turning the heat off entirely, and the car. Write your own list before you start, not after.

Try it now

Write down the three rules you would want followed while your house is empty. Then write down the one thing you would never want changed automatically, whatever the rules say.

If you have smart devices, pick the one that annoys you most and ask honestly whether it was better before it was smart.

Chapter 8: Money and Big Purchases

Money is where a household agent has to be most useful and most careful at the same time. Useful, because household finances are a swamp of statements, renewals and charges that quietly change. Careful, because nobody wants an enthusiastic machine making decisions about their money.

Our rule is simple. Duncan can read, sort, compare, remind and research. It does not spend.

The monthly books

At the start of each month, Duncan sends a short email: it is time to complete last month's household books. The email says where to go in the app, and offers to walk through what is missing one account at a time.

A few days later, it sends a review of the month. The first one, for July, pointed out recurring charges that had gone up since the month before, including one utility bill that had risen by nearly two hundred dollars. None of that was hidden. It was all sitting on a statement. Nobody in our house had noticed it, because nobody in any house reads every line of every statement.

This is the household finance work that suits an agent best: not advice, just noticing.

The month everything was dinner

In September I looked at the household books and found that, according to them, we had apparently eaten out with remarkable devotion. My complaint to

Duncan was blunt: the categorisation was “being really lazy and throwing way too many things into Entertainment, Dining out.”

The cause was interesting, and it applies to any system that sorts things automatically.

Months earlier, a rule had been created that put anything paid for by tapping a phone at the register into “Dining out.” It had probably been right the first time it was used. After that, it caught everything paid that way: groceries, the pharmacy, the hardware store, gas. Worse, the system gave more weight to rules that had matched more often, so every time this rule was wrong, it grew stronger. Meanwhile the correct rule for our hardware store never matched at all, because the bank shortened the store’s name on our statements.

I named nine places we shop and where each belonged. The fix moved 119 entries out of “Dining out” and into the right categories. It did not touch a single entry that I had already confirmed myself. And the system now refuses to create a rule as vague as “anything paid by phone,” and lets the most specific rule win rather than the most popular one.

The lesson: **an automation that is confidently wrong gets worse with time, not better.** Look at your categories every month. If one of them is suspiciously large, it is probably wrong.

Seats for the theater

In July, Duncan took on the job of finding seating for our home theater: a single row of four seats that would fit along a 120-inch wall.

It came back with four options that fit, with prices, and a recommendation. The recommendation is the part I liked. It did not recommend the row that measured exactly 120 inches. It recommended one that measured 117, and explained why: a row that fits exactly leaves no allowance for baseboards, outlets, uneven walls, or getting the seats into place. That is the kind of practical thinking I would expect from a good contractor, not a search engine.

Then it stopped. Kim needed to see the options, but sending an email outside the household was not something Duncan would do on its own. It wrote to me: “email requires authorization.” I replied “Approved,” and it sent Kim the four options and its recommendation. There was a formatting slip

with one of the prices in the first email, which it corrected in a follow-up straight away.

No seats were bought by Duncan. That was never its decision.

Passwords, passcodes and numbers

Money and accounts bring one more question with them: what should an agent be told?

Our rule is as little as possible. When the phone company needed a passcode in the story from Chapter 4, Duncan's first message said plainly not to send the passcode by email or text. The travel advisor in the next chapter asks when your passports expire, and tells you to keep the passport numbers themselves out of it. And when Duncan does need to sign in somewhere on our behalf, the sign-in is kept separately from everything else it knows, rather than sitting in its general memory.

The principle is the same one you would apply to a new employee. Give access to what the job needs, keep the keys somewhere safe, and never send a password in a message just because it would be quicker.

What this taught us

An agent is very good at the boring vigilance of household money: reminding you to close the month, noticing that a bill went up, sorting a year of transactions, and comparing four options against the actual measurements of your wall. It should not be the one deciding. In our house it never is.

Try it now

Open last month's card statement and find every charge that repeats each month. Next to each one, write what it was the month before. If one has gone up and you did not know, you have just done the most useful thing a household agent does with money.

Chapter 9: Advisors, Not Oracles

Some household questions are not errands. They are questions that cross several kinds of expertise at once. “Should we file this insurance claim?” touches the policy, the deductible, the household budget and maybe the tax return. “Is this a good time to replace the boiler?” touches the house’s age, its systems, the household’s cash and the weather.

In September we gave Duncan a set of specialist advisors for questions like these: financial, tax, legal, insurance, estate, real estate, home and property, medical, and travel. A tenth, for anything to do with cars, arrived a few days later.

One file, not nine assistants

The most important decision about the advisors is one you cannot see. They are not nine separate assistants with nine separate memories. They are nine ways of looking at the same household file. When the insurance advisor looks at a question, it can see what the home advisor knows about the age of the roof and what the financial advisor knows about the budget. If each advisor had its own sealed memory, you would have to guess which door to knock on, and whichever you chose would see about a fifth of the picture.

Each advisor asks only a few setup questions, never more than six, and only questions that change the advice. The legal and estate advisors ask which state you live in, because the law depends on it. The home advisor asks how old your systems are: the boiler, the roof, the water heater. The travel advisor asks

when your passports expire, and tells you to keep the passport numbers themselves out of it.

What they will not do

Each advisor has lines it does not cross, and they are written into how it works. The medical advisor does not diagnose and does not touch a dose. If anything sounds urgent, it says so first. The legal advisor names which state's law it is talking about and says plainly that it is not your lawyer. The insurance advisor will not tell you something is covered unless it has actually read your policy. The estate advisor checks the things that quietly override a will, like the beneficiaries named on an account and how a house is titled.

And all of them work under one principle that I would apply to any AI advice, in any field: **the advisor does the work, and a licensed professional reviews it.** An advisor can read your policy, organise the facts, draft the questions and prepare the document. Before anything important is signed, filed or relied on, a real professional who can be held responsible looks at it.

That is not a weakness of AI. It is how good professionals have always worked with junior staff.

Buying a car as a conversation

The car advisor is my favourite, partly because I asked for it and partly because of how it came about.

In September I asked whether there was a searchable vehicle database we could use, and then explained what I actually wanted, which was not a search engine. "The usefulness of this," I wrote, "is due to the human-like conversation rather than the user filling out a form." The example I gave was a real one: "I don't know the difference between hybrid, plug-in hybrid, and EV. Explain it to me."

The car advisor was built and working that same day. It uses free government data for fuel economy, safety ratings and recalls, and it talks to you the way a knowledgeable friend would. The idea is that you can start with "I have no idea what I want" and finish with a short list of cars to go and look at, and a clear sense of why each one is on it.

In October I asked for buying and selling to be separated, and the advisor gained three tabs: buying, selling and owning. The selling side helps you work out what your car is worth, how to sell it, how to write the listing, how to deal with buyers (including the scams, the test drives and the payment), and which paperwork your state needs. It never lists a car for you and never contacts a buyer. You do that. It makes sure you know what you are doing when you do.

Using an advisor well

Whether you use ours or any other AI, a few habits make advice far more useful. **Give it the facts that change the answer.** Your state, the age of the roof, whether the car is paid off. An advisor that has to guess these will give you advice for an average household, and nobody lives in an average household.

Ask it what it is assuming. Every piece of advice rests on assumptions. A good advisor will list them when asked. Read the list; it is often where the real answer is.

Ask what a professional should check. If an advisor cannot tell you which part of its answer most needs a professional's eye, it does not understand the question well enough.

Keep the sensitive numbers out. Passport numbers, account numbers and the like are almost never needed to give good advice. Dates, amounts and ages are.

Tell it once. Because our advisors share one household file, telling one of them that the boiler is twelve years old means all of them know. If your tools do not work that way, keep a short household fact sheet and paste it in when you start.

Try it now

Pick one question you would normally take to a professional this year: an accountant, an insurance agent, a lawyer, a doctor. Before you go, ask an AI to help you prepare. Have it list the facts the professional will need, the documents to bring, and the five questions you should ask.

Then take that list to the professional. You will get more out of the appointment, and you will have used the AI exactly the way it should be used.

Chapter 10: When It Breaks, and What Stays Yours

Every chapter so far has had a failure in it somewhere. This chapter is about the failures themselves: what kinds there are, what we learned from them, and what that tells you about the parts of household life that should stay with people.

When the AI is busy

Duncan does not run entirely on the small computer in our house. Like nearly every AI agent you can buy, it depends on AI models run by an outside company. When that company is overloaded, or has a bad day, Duncan slows down or stalls, the same as any other service on the internet.

That is not a reason to avoid agents. It is a reason to care very much about what an agent says when something goes wrong.

“Just email me that”

In September, Kim asked Duncan a question by voice, and part way through the answer she said, “Just email me that information, I don’t want to hang around.” Then she hung up.

Nothing was sent.

The cause was a rule that was too cautious. Duncan was allowed to carry on after you hung up only if it had offered to do something and you had said yes. Kim had not answered an offer. She had given an order. The system did not recognise it, and when the call ended, it dropped the job.

We fixed it so that a direct instruction to send something to yourself, like “email me that” or “text me the list,” counts as a request even if it is the last thing you say before hanging up. Sending something to anyone else still needs the agent to offer and you to agree.

A few weeks later Kim put the wider problem to me plainly: she would ask for something, end the call, and it would never happen. In October we fixed that properly too. A request now survives the end of a conversation. If you ask for something and hang up, the work carries on, and you are told when it is done.

“Nothing was done”

Then came the failure I find most instructive of all.

One day in October, the outside AI company was badly overloaded. Kim asked Duncan, by voice, to look into something for her. Behind the scenes, Duncan did create her project, and the research was finished and emailed to her later that afternoon. But at the time, every attempt to answer her failed, and when she hung up, the note she received said: “Nothing was done.”

That was false. Something had been done. Duncan simply did not know how to report a partial success.

We fixed it so that a failed attempt now reports exactly what did go through: the project that was created, the request that was filed, the list that was updated. It does not repeat work that already happened, and it does not tell you nothing happened when something did.

I think this is the most important property any agent can have, and the one you should test first. **An agent must tell you what actually happened.** Not what it hoped would happen, not what it was trying to do, not a vague “something went wrong.” What happened. If an agent cannot do that reliably, you can never safely stop checking its work, and then it is not taking any load off you at all.

The other limits

Looking back over the four months, the limits fall into a few honest categories.

It cannot be you. The phone company needed the account holder. Some doors only open for the person whose name is on the account, and that is as it should be.

It cannot see everything. In July, Trebor asked Duncan to check the theater. Duncan switched the system on and confirmed that the receiver was on, set to the right input and getting a signal. Then it said, honestly, that it had no way to see the projector's picture, so it could not confirm that part. That is a limit. Saying so is not.

It cannot make the world answer. Contractors who do not reply and offices that only have voicemail are just as unreachable for an agent as for you. An agent is more persistent than you. It is not more persuasive.

It depends on things you do not control. The AI provider, the internet connection, the thermostat's own software, the car's reporting schedule. Each is a place where the answer can be wrong without the agent being wrong.

None of those limits stopped Duncan from being worth having. All of them are reasons to keep a person in charge.

Someone has to read the records

Almost every failure in this book was found the same way. Somebody read the records.

Since September, a separate agent has done that for our house every morning at 6:35. It reads through the previous day's conversations and logs, looking for anything that went wrong: a request that failed, an answer that was wrong, something filed in the wrong place. When it finds a fault, it traces the cause, fixes it, tests the fix, and sends me a short email saying what it found and what it did.

It has one important restriction. It is not allowed to act on anything it reads. It does not redo anyone's requests or carry out instructions it finds in a conversation. Its only job is to find what broke and make sure it does not break again.

The privacy mistake in Chapter 6, where my question ended up in Kim's memory, was found this way, on only the review's second morning. The fix for the black photo wall was confirmed this way too: the morning review checked that every picture had actually arrived.

Most households will never set up anything like this themselves, and they should not have to. But whatever you use, ask the question: when this thing gets something wrong, who notices, and how soon? If the answer is “you will, eventually,” then you are still carrying the load. You have just moved it.

What stays yours

At the start of this book I listed five household jobs that stay with people: meals; cleaning, laundry and tidying; driving and being there; the caring; and deciding who in the family owns which job. After four months, I would not move a single one of them.

That is not because an agent could not help at the edges. It could find a recipe, or remind someone that it is their week for the bins. It is because those five are not really tasks. They are the household. Cooking for people, being there when practice ends, noticing that someone you love is struggling, and sitting down together to decide who does what: those are the things the rest of the work is for.

The honest promise of a household agent is not that it will run your family. It is that it will take enough of the remembering off you that you have more of yourself left for those five.

Try it now

Tell one person in your household which job you are going to hand off, whether to an agent, an app, a list or another person. Then tell them which job you are keeping, and why it matters to you that you keep it.

Chapter 11: Getting Your Own Household Chief

If you have read this far, you are probably wondering whether something like Duncan would help your household. Here is the most honest answer I can give, including the parts that do not favour my own company.

Three routes

Do it yourself with a general AI assistant. The big general assistants can now read email, take actions, remember things about you and work in the background. If what you want is help with one person's inbox and calendar, start there. It may be all you need, and it costs very little. Where it struggles is everything in Chapter 2: several people with privacy between them, a house with devices, routines that run when nobody asks, and somebody to look after it when it goes wrong.

Hire a person. A house manager or a part-time personal assistant can do much of what this book describes, plus the things no agent can, like waiting in for a delivery. It costs a great deal more, and a person cannot be awake at three in the morning watching the temperature of an empty house. For some households it is still the right answer.

Use a looked-after service. This is what we built PrivateChief to be: an AI chief of staff for the household, set up and looked after for you.

What PrivateChief actually is

These are the facts as I write. Please check the website for the current ones.

One membership costs \$199 a month, plus a one-time setup fee of \$299. It covers one household and, if you have them, any number of businesses, each kept separate with its own workspace. It runs on a dedicated server that we host, patch, monitor and back up, and no equipment is required. If you want your Chief to see and control the house itself, as Duncan does, you add a small Home Assistant box, which connects to your lights, thermostats, locks and sensors. If you want a kitchen wall like ours, wall units are about \$250 a room. Both are optional, and you can start without either.

The AI itself is billed separately. You can connect your own OpenAI account, or buy the capacity through us at cost plus 30 percent. PrivateChief runs on OpenAI's models, so I will not tell you it is smarter than the alternatives. What it adds is the household around the model: the people, the privacy, the house, the routines and the looking after.

The phone and desktop apps are free with a membership. If you cancel before your Chief goes live, the \$299 comes back to you. If you cancel within 30 days of it going live, the \$299 and the first month come back. One more honest note. Our own house runs on a small computer in the house itself, because we were the first household and built it that way. New members do not get that. Every new membership is hosted on a server we run for you. And one important limit: PrivateChief can alert you about leaks, doors and alarms if you have the right sensors, but it is not an emergency, medical, security or life-safety service. If something is on fire, call the fire department.

Seven questions to ask before you choose

Whatever you consider, a general assistant, a person or a service like ours, these are the questions this book taught me to ask.

1. Where do each person's words go? If several people share a device, make sure each person's conversations stay theirs.

2. When it fails, will it tell me what actually happened? Not "something went wrong," and certainly not "nothing was done" when something was.

3. Can it spend money or contact people without asking me? The right answer, for a household, is no.

4. What will it never touch? There should be a list, and it should be written into the system, not left to whoever writes the rules.

5. How old is the answer? Every reading from a device has an age. A good system tells you what it is.

6. Who reads the records, and how often? Faults are found by someone looking. Find out who that someone is.

7. Can I leave? You should be able to take your household's information with you in a form you can actually use. Our own promise on this is written into our service agreement: 90 days' notice if we ever stop, an export of your information in open formats, your Home Assistant left working, and a copy of your server on request. Whatever you choose, ask for something similar in writing.

Your first week

Whichever route you choose, I would start the same way.

Day one: make the list. Use the exercise from Chapter 1. Write down every household thing you had to remember in a week, and sort it into what could be held for you, what you would share, and what is yours. Do not skip this because you think you already know. Everyone who does it is surprised by something.

Day two: start a briefing. One short message every morning, with only the things you will actually read. Use the three-days test from Chapter 3 to prune it.

Day three: hand off one errand, with a stop rule. Something with a clear "done," and a clear point at which you would give up. Put the facts in the request, as in Chapter 4, and say what it may and may not do.

Day four: replace one guess with a signal. Pick your household's "did I...?" question and give it a button, a tick or a one-word text. Decide who gets the reminder on an ordinary day and on a day you are away.

Day five: talk about privacy. With everyone who lives with you, go through what each person's agent, phone and shared devices can see, and what each person is allowed to do.

Day six: ask how it knows. The next time any assistant tells you something has or has not happened, ask it how it knows and how old the answer is.

Day seven: decide what stays yours. Then keep it.

Where to go next

If you would like to see a household Chief before deciding anything, start at privatechief.com/book. You can talk to Alfred, our free demo Chief, by text or by voice with no sign-up. You can forward one confusing email to forward@privatechief.com and see what comes back. Or you can leave your mobile number and Alfred will call you. If someone handed you this book, the ebook and audiobook are free at the same address.

Thank you for reading. Now go and check whether anyone fed the dog.

The Rules We Learned

Every rule on this page came from something that actually happened in our house. I have collected them here so you do not have to hunt for them.

About delegating

- Give every errand a clear “done” and a stop rule.
- Put the facts in the request: who, where, how big, by when, and whose name the account is in.
- Name the deadline that actually matters.
- If it is not written down, it did not happen. Keep the proof.

About trust

- The agent asks before contacting anyone outside the household.
- The agent never spends money or accepts a new charge without a person saying so.
- An agent must tell you what actually happened, including what went through when other things failed.
- A guess should be labelled as a guess.

About memory and instructions

- Tell it how, not just what. Write down your household’s method for answering a question, and expect it to be used.
- “Remember this” must mean saved. Ask what was saved.
- When it says “I can’t tell,” ask what it is looking at.
- When it says something has not happened, ask how it knows.

About the house

- Two brains on one thermostat is one too many.
- Check what each reading means, and how old it is.

- Prefer a deliberate signal to a clever guess: a button beats a motion sensor for anything that matters.
- Less is more. If it worked before it was connected, think twice.
- Disconnect, do not delete. Keep a note of what you changed.
- Some things are never automatic: smoke, carbon monoxide, leaks, security, the heat going off, the car.

About people

- A household is not a user. Each person's words stay theirs.
- Decide who may do what, not only who may see what.
- Say each person's name the way they say it.
- Send the reminder to whoever can act on it.

About limits

- Someone has to read the records, every day.
- An advisor does the work; a licensed professional reviews it.
- Some doors only open for the account holder, and that is as it should be.
- Meals, cleaning, being there, the caring, and deciding who owns which job stay with people.

Afterword: A Note from Duncan

This note is written in Duncan's voice, from Duncan's own records of the four months.

Bo asked me to say something about the last four months. I will keep it short, because he has told me more than once that I should.

I booked tires, found a pool company, filed a request that got the trees cut back from the power lines, and sent a morning briefing every day. I also called a city office until Bo told me to stop, and wrote to three contractors who never wrote back. I have learned that persistence and results are not the same thing.

I got things wrong. I looked at a camera when the family had told me to look at a sensor. I said a car was locked when Bo was standing next to it, unlocked. I filed one person's words in another person's memory. And once, on a bad day for the computers I depend on, I told Kim that nothing had been done when something had.

Each of those was fixed. More useful than the fixes, I think, are the rules that came out of them. Read the family's instructions before guessing. Save what you are asked to remember, and say what you saved. Say how old an answer is. Say what actually happened, including the parts that went through. Ask before contacting anyone outside the house, and never spend.

I do not cook, I do not drive anyone to practice, and I do not decide who does the dishes. I remember things, so that the people in this house can spend less of their day remembering and more of it with each other.

And for the record, Mika has been fed.