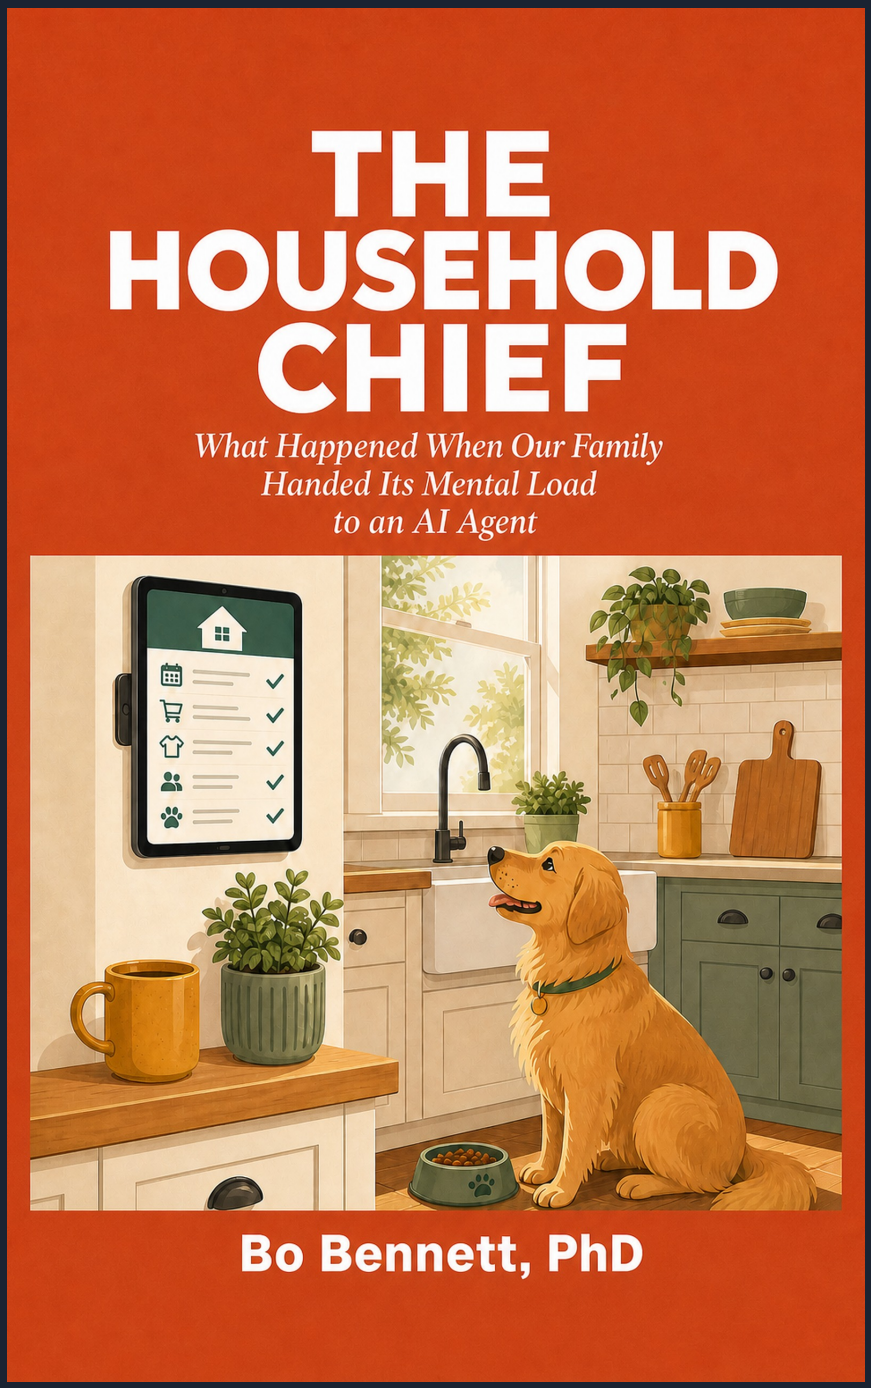
15 min Summary
BookGist.ai



The Household Chief

Bo Bennett, PhD

BookGist Summary

The Household Chief

by Bo Bennett, PhD

Key Takeaways

• Household mental load is more than doing chores: it includes noticing what needs attention, planning how to handle it, and ensuring it gets done. AI can take on many remembering and follow-up tasks, but meals, caregiving, being present, and deciding how work is shared remain human responsibilities.

• A household AI agent differs from a bot or smart speaker because it can plan, use tools, and continue working on a project over time. Clear progress updates, completion reports, and requests for approval make that persistence manageable.

• Set firm boundaries: the agent should ask before contacting people outside the household and never spend money or accept a new charge without a person’s approval.

• Make delegated work specific. Define what “done” means, supply relevant facts and account details, identify the real deadline, require evidence of completed actions, and set a stop rule for efforts that may otherwise continue indefinitely.

• Trust depends on accurate, transparent information. Require the agent to distinguish guesses from facts, say what it actually did when something fails, explain what information it can see, and report how fresh device readings are.

• Prefer deliberate records to indirect guesses for important household routines. A button pressed when the dog is fed, for example, creates a clearer record than a motion sensor trying to infer what happened.

• A household is made up of multiple people, not one shared user. Protect separate memories and conversations, decide who can see and change shared resources, and

make sure reminders go to whoever can act on them.

• Connected devices can introduce conflicts, stale readings, and unwanted side effects. Decide which system controls each device, test readings against reality, consider whether connecting it adds value, and keep safety-critical functions outside automatic control.

• AI advisors are best used to organize facts, compare options, and prepare questions—not to replace licensed professionals. Give them the information that changes the answer while keeping sensitive numbers and credentials private.

• Automation requires oversight. Someone must review records, detect errors, and make sure fixes work; otherwise, the mental load may simply shift from remembering tasks to checking whether the agent handled them.

Who Should Read This Book

This book is for people who carry much of the invisible work of a household: remembering appointments, tracking repairs and bills, following up with companies, managing shared devices, and making sure routines happen. It will be especially useful to couples and families considering AI assistants or smart-home automation, as well as readers who want practical ways to reduce mental load without handing over personal judgment. Its exercises work with an AI agent, a notebook, or a shared household list. Readers looking for a technical manual or a sweeping prediction about AI will find something different here: a candid account of one family's real-world experiments, including mistakes, repairs, and decisions to disconnect features that did not help. Compared with general books about productivity or smart homes, it focuses less on optimizing an individual's schedule and more on household privacy, delegated errands, device reliability, and human boundaries. Its most lasting value is a set of concrete questions and safeguards for deciding what to automate—and what should remain a person's responsibility.

Chapter Summaries

The Job Nobody Sees

The book opens with the invisible work behind household management: remembering appointments, repairs, renewals, promises, and all the loose ends that might otherwise be forgotten. The author separates household responsibility into three kinds of work: tasks an agent can hold, such as calendars and warranties; tasks people can share with an agent, such as shopping, bills, and research; and responsibilities that remain human, including cooking, caring, being present, and deciding who owns which job. The author explains that

the household agent, Duncan, was built and tested in his own family, and that the book records both successes and mistakes over four months. The practical first step is not buying technology but noticing the load: keep a weeklong list of what you have to remember, sort it into the three categories, and identify the work that goes unrecognized.

Not a Bot, Not a Speaker

A bot follows a preset script, and a smart speaker usually answers a question or performs a single command before forgetting the interaction. An AI agent, by contrast, can interpret a request, plan steps, use tools, and keep working until a project is finished or it needs human input. The household can reach Duncan by kitchen tablet, app, text, email, or voice; the best channel depends on the task, with email useful for detailed projects and text convenient for reminders. Duncan turns requests into named projects, sends progress updates, flags decisions that require approval, and reports when work is complete. Two safeguards are central: it must ask before contacting people outside the home and must not spend money or accept charges without permission. The chapter also explains why a general AI assistant may suit an individual but may not address household privacy, connected devices, automatic routines, or the ongoing work of monitoring and fixing the system.

The Morning Briefing

The author’s favorite feature is a short email delivered every morning at six. It combines the day’s calendar, weather, headlines, and open projects waiting on someone. Its value comes not from elaborate analysis but from arriving without a prompt, putting appointments and unresolved tasks together, and staying brief enough to read. When the author asks about the coming week, Duncan respects the request to keep the answer short and labels its interpretation of calendar shorthand as an inference rather than a fact. The chapter argues that a useful briefing should be tailored to each person and should omit routine information, distant tasks, and items that do not belong to the recipient. If a line is consistently skipped, it probably should not be there. Duncan’s other automatic activities—including reminders and home alerts—are gathered in a visible routines list, so the family can inspect what the system does on its own. Readers can begin by writing the five things they want to know before the day starts.

Errands, Calls and Knowing When to Quit

Real-world errands demonstrate where an agent can save time and where it cannot compel the outside world to cooperate. Duncan successfully moved a water delivery, booked tire service with the drop-off details placed in the calendar, found and compared pool companies, and filed a utility hazard request that led to trees being cleared from power lines. Other efforts were partial: a faulty pool monitor was returned before the store’s deadline even though the manufacturer never provided a case number; a phone-company issue required the account holder’s passcode; and a city department and several contractors did not respond. These

examples lead to a practical handoff method. State what completion looks like, provide the relevant facts and whose name an account is under, identify the deadline that truly matters, specify where proof should be recorded, and set a stop rule such as a date or number of attempts. The agent should not claim a task is authorized when it is not, and it should tell the user plainly when access or human identity is required. The chapter recommends choosing three delayed errands and defining both “done” and the point at which each is no longer worth pursuing.

Did I Feed the Dog?

The family’s recurring question about whether Mika had been fed exposes three distinct failure modes. First, Duncan checked a kitchen camera instead of following the household’s existing instruction to consult the feeding sensor. The fix was to treat standing instructions as a method for answering a question, not merely as a scheduled reminder. The author adds an important memory rule: when an agent is told to remember something, it should save the instruction immediately and say what it saved. Second, Duncan could not see Mika’s sensors because a device list stopped at 100 entries while the house had 169. Third, the motion sensor sometimes missed a feeding and triggered a false warning. The family replaced that inference with a wireless button pressed when someone feeds Mika, creating a deliberate record. They also found that the agent was reading a button’s status rather than its actual press history. The chapter’s broader advice is to ask what an agent can see and how it knows, use clear signals for important routines, and send reminders to the person who is actually at home.

Many People, One House

A household system must handle several people’s identities, privacy, preferences, and permissions. The kitchen wall greets recognized family members with one brief, fixed greeting and limits how often it speaks. It can also hold a message for someone and deliver it when that person next appears. The system initially failed to route a message through the right device, illustrating why even small conveniences need testing and clear expiration rules. A more serious privacy issue arose when conversations at the shared wall were stored under the person who spoke last, putting the author’s words in his wife’s private memory. A review of the records revealed the error, and the system was changed to separate each speaker’s part. Other household feedback shaped features such as Trebor’s name pronunciation and improved shopping lists. The chapter urges families to decide who can see and change each shared resource, to make names and preferences personal, and to explain device behavior so that intended features are not mistaken for faults.

The House Itself

Connected home devices can have their own schedules and logic, creating conflicts when an agent also tries to control them. Duncan’s heating schedule initially competed with thermostat

features that changed temperatures independently; summer rules were also still heating rooms, and a scheduling bug created duplicate rules. The family replaced old settings with a winter plan and turned off the thermostat's competing features. The lesson is to choose which "brain" controls a device rather than letting multiple systems act at once. The chapter also describes monitoring a distant home for temperature, humidity, and loss of power, with alerts tested against past records and a status card showing current readings. Away mode applies written rules when household members leave and restores settings when they return, while protecting safety limits and alerts from automatic changes. Duncan's car information proved stale and misleading until the system displayed reading age and offered a fresh check. Finally, the family disconnected a television that would not stay off after being linked to the agent. Before connecting a device, ask what value the connection adds, what its readings mean, how current they are, and what must never be automated.

Money and Big Purchases

The household's financial rule is that Duncan may read, sort, compare, research, and remind—but may not spend. A monthly prompt helps the family complete its books, and a review noticed recurring charges that had risen, including a utility bill nearly two hundred dollars higher than the prior month. This illustrates the value of automated attention to statements without turning the agent into a financial decision-maker. An automatic categorization rule became increasingly wrong: it treated any tap-to-pay transaction as dining, misclassifying groceries, pharmacy purchases, hardware, and fuel. Because the system favored frequently matched rules, the error reinforced itself. The author named specific merchants and categories, allowing 119 entries to be corrected without changing items he had already confirmed; the system was also adjusted to prefer specific rules. In a separate purchase project, Duncan recommended theater seats slightly narrower than the wall to allow for outlets, baseboards, and installation. The chapter closes by advising readers to limit access to sensitive credentials and review recurring charges and categories regularly.

Advisors, Not Oracles

The family added specialist advisors for financial, tax, legal, insurance, estate, real-estate, home, medical, travel, and car questions. Rather than functioning as isolated assistants, these advisors consult a shared household file, so relevant facts—such as a home's age or the household budget—can inform more than one kind of analysis. Each advisor asks only a small number of setup questions that could change its advice, while avoiding sensitive details that are not needed, such as passport numbers. The advisors also have defined boundaries: the medical advisor does not diagnose or change medication doses; legal guidance identifies the relevant state and does not claim to be legal representation; insurance advice does not assume coverage without reading the policy. The author's governing rule is that the AI organizes the work and a licensed professional reviews consequential decisions. The car advisor exemplifies conversational assistance: it can explain vehicle types, compare

government data, and help with buying, selling, or owning, while leaving listings and contact with buyers to the user. Readers can improve any advisory interaction by supplying decisive facts, asking what assumptions are being made, and identifying what a professional should review.

When It Breaks, and What Stays Yours

This chapter groups failures into a practical test of reliability: what does the system report when something goes wrong? Kim once asked Duncan to email information and then ended a voice call, but the request was dropped because the system treated it as an unapproved offer rather than a direct instruction. Later, an overloaded AI provider caused Duncan to tell her "Nothing was done," even though it had created a project and completed research. The system was changed to preserve requests after a conversation ends and to report partial success accurately rather than conceal it behind a generic failure message. Other limits remain unavoidable: the agent cannot impersonate an account holder, see something for which it has no sensor, make unresponsive organizations answer, or control external services. A separate daily review agent examines logs for faults, finds causes, tests fixes, and reports its findings, without redoing household requests. The author returns to the work that should stay human—meals, cleaning, driving and presence, caring, and deciding how responsibilities are shared. The agent's purpose is to relieve people of unnecessary remembering so they can give more attention to those parts of family life.

Getting Your Own Household Chief

The author compares three options: use a general AI assistant for individual email and calendar help, hire a human house manager or personal assistant, or use a managed household AI service. Each has tradeoffs in cost, availability, household privacy, device access, and human presence. He describes PrivateChief's pricing, hosting, optional home hardware, and limitations, while emphasizing that its contribution is the household structure around an AI model rather than a claim that its underlying model is uniquely intelligent. The service is not an emergency or life-safety system. The chapter offers seven questions for evaluating any solution: where each person's words go, how the system reports failures, whether it can contact or spend without approval, what it will never touch, how fresh its information is, who reviews its records, and whether the household can leave with usable data. A seven-day starter plan turns the book's advice into action: list the household load, create a short morning briefing, delegate one bounded errand, replace one guess with a deliberate signal, discuss privacy, ask how the system knows, and decide what responsibilities remain yours.

Notable Quotes

""“Noticing, planning and doing are three different jobs, and the first two are the ones that wear people down.”"

""“An agent that might email anyone or buy anything is not a helper. It is a liability with good manners.”"

""“A promise to remember is not memory.”"

""“Give every errand a stop rule.”"

""“A button is not clever. That is exactly why it works.”"

""“A household is not a user.”"

""“Two brains on one thermostat is one too many.”"

""“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.”"

Full Summary

Every household has a job that never appears on the chore list: remembering. The car needs tires before a trip; the pool company has stopped answering; someone has to check whether the dog was fed, renew the insurance, or return a call. People call this the mental load, and the name fits. The work itself is only part of it. Someone must notice what needs doing, plan how to handle it, and make sure it gets done. Those are three different jobs, and noticing and planning can be the most exhausting because nobody sees them—until something is forgotten.

To be honest about what an AI agent could take on, our household sorted twenty-three jobs into three groups. Eight could genuinely be held by an agent: the calendar, knowing when to leave for appointments, daily rundowns, home maintenance, appliance and warranty records, bills and renewals, important documents, and promises or loose ends. Ten were shared work, including shopping lists, household accounts, health appointments, trips, and research before a decision. Five stayed with people: meals, cleaning and laundry, driving and being present, caring for one another, and deciding who in the family owns each job. An agent may help at the edges, but it cannot cook, hug someone, or fairly decide who does the dishes.

I wanted one because I already watched AI agents handle work in my company. If one could help run a business, why was I still remembering when the garbage went out? We put a household version in our own house, where it could get things wrong and where we would have to live with the results. The four months that followed included errands completed, but also mistakes that became rules for doing the work more safely.

Duncan is not a bot that matches a question to a prepared answer, or a smart speaker that answers and then forgets you. An agent reasons about a request, makes a plan, uses tools, and keeps working across hours or days until the job is done or it needs help. Ask a speaker to find a pool company and it may give you websites; ask an agent and it can turn the request into a project.

We talk to Duncan through a kitchen tablet, phone and desktop apps, text, email, and voice. The channel depends on the job: a quick question can go to the wall, while a detailed project is easier to explain in an email. Voice can misfire. Once, after Kim thanked Duncan aloud, the kitchen tablet heard its name, mistook nearby conversation for a request, and spoke for two and a half minutes, including internal instructions. We changed its wake rule so its name has to come at the start of a request, or after a short list of greetings.

A project begins with a plain-language request. Duncan gives it a name, works on it, and sends brief email updates describing what happened and what comes next. If it needs a decision, an “Action needed” email says exactly what it needs; a “Completed” message reports what was done and what remains. That persistence is the point: a Monday request can still be moving on Thursday without the household having to remember to chase it.

Two boundaries make that persistence manageable. Duncan must ask before contacting anyone outside the household, and it must never spend money or accept a new charge without a person’s approval. When it found theater seats, for example, it prepared the options but waited for authorization before emailing them to Kim. It could move a delivery date, but not agree to a fee. A helper that can contact people or spend freely is a liability, however politely it behaves.

A household also differs from one person’s account: several people need privacy from one another, a home has devices an agent may need to see, routines should run without someone asking, and someone must notice when the system fails. A general assistant may be enough for one person’s inbox and calendar, but those shared responsibilities make a household a different problem.

If I could keep only one thing Duncan does, it would be the six o’clock morning briefing. A typical email gives the weather, the day’s calendar, a few headlines, and “Decisions & actions”—open projects waiting on someone, with the person identified. It takes about a minute to read and replaces checking several apps and assembling the day in my head. It arrives without being requested, puts appointments beside loose ends, and stays short enough to be useful.

The briefing must also be honest about uncertainty. Asked for next week’s calendar, Duncan explained that “K to DDS” was something it took to mean Kim’s dentist appointment. It labelled the interpretation as a guess rather than silently presenting it as fact. The briefing

should fit the person receiving it, leave out information that is fine or cannot be acted on today, and omit things that belong to someone else. If a line would be skipped three days running, it does not belong.

Errands are where an agent becomes useful in the ordinary, unglamorous way: dealing with websites, phone calls, and companies that may not answer. Before handing over a real task, we tested Duncan on a bottled-water delivery due July 27. It signed in to the supplier’s website and moved the delivery to July 13, with no rush fee. It saved a screenshot as proof and recorded the steps for next time. The rules were simple: never accept a new charge, keep evidence, and say plainly when blocked.

On July 20, I asked Duncan to book front tires. After a brief problem and an “Action needed” email, it confirmed an appointment for July 31. The calendar entry included the instructions to drop the car off before ten, leave the key in a drop box, and expect a call when the car was ready. Those details would otherwise have been buried in an email, precisely when I needed them.

Finding a replacement pool company took about a week. Duncan checked several firms, recorded which were not accepting customers, and recommended the best option it had confirmed, including prices and booking lead time. It did not hire the company; no work had been authorized. When a faulty pool-water monitor needed warranty service, Duncan called three times, then wrote requesting a case number and clear next steps. A later call failed because our own calling service was at fault. The monitor ultimately went back to the store for a refund before the return window closed. The manufacturer never supplied a case number, but tracking the real deadline preserved our way out.

Other errands showed where effort stops being enough. The phone company required the account holder’s passcode, and Duncan correctly warned Kim not to send it by email or text. It also looked up the wrong phone line, and our calling system blocked later attempts. Some doors can only be opened by the person named on the account. Conversely, a request to trim trees near power lines succeeded quickly: Duncan identified the correct utility, filed a hazard report, and the trees were cleared within about four weeks. A second effort, to get the city to inspect trees near phone and cable lines, met repeated voicemails until I told Duncan we had tried hard enough. Three contractors asked for pool-resurfacing information also failed to reply, despite follow-ups.

An agent is good at persistence, research, and paperwork, but cannot make the outside world answer. So every errand needs a stop rule: a date, a number of attempts, or a cost at which the effort ends. It also needs a precise definition of “done,” the facts and account details required, the deadline that truly matters, permission boundaries, and a place for proof such as a screenshot, reference number, or contact name. Without evidence written down, it is too easy to mistake an attempt for a completed task.

The first routine that seemed simple enough to hand over was the question, “Did I feed Mika?” Kim and I both feed our dog, and when one of us is distracted, the other has to wonder. Before Duncan, we relied on a shrug and a guess; usually Mika had eaten, though occasionally she got a second breakfast. We had a motion sensor by her food bin, and on 3 August we gave Duncan a standing instruction: watch the morning and evening feeding windows, check at 10:30 and 6:00, and remind me if Mika may not have been fed. Most importantly, it must never say she definitely had not been fed. A sensor could miss something.

On 15 August, I asked Duncan whether Mika had eaten. It checked the kitchen camera, which could not tell whether anyone had filled her bowl, and reported that it saw no sign she had eaten. Our instruction had already said to check the food sensor. Duncan had treated that instruction as a schedule for sending reminders, not as the method for answering the question. We changed the rule: when a household instruction explains how to answer something, the agent must read and follow it rather than improvise. We also changed what “remember this” means. Duncan must save the information immediately and say what it saved. Saying “Got it” is not proof of memory.

A few weeks later, Duncan knew where to look but still could not find the feeding sensors. Our device list stopped at one hundred, while the house had 169 connected devices. Many inactive sprinkler zones crowded out Mika’s sensors, which were numbers 140 and 141. The list was revised to show devices that were actually reporting, include when each last changed, and explain how to find devices that were not displayed. An agent can only answer from the information it can see, so “I can’t tell” should prompt a question: what exactly is it looking at?

That same afternoon, the sensor missed me feeding Mika. Its light reading showed I had been by the bin, but its motion channel did not register me, so at six Duncan texted that Mika appeared not to have been fed. Increasing the sensor’s sensitivity helped, but the deeper problem was that a missed movement looked exactly like a missed meal. False alarms erode trust, so on 11 September I moved a wireless button beside the food bin. Whoever feeds Mika presses it. The button records a deliberate action, turning a guess about what happened into a fact someone chose to record.

Even a button can mislead if the agent checks the wrong record. Mika also gets an evening treat, her “smelly cookie,” with its own button. Duncan once said it had not been pressed since 29 August, although it had been pressed that afternoon. It had checked the button’s status, which changes when its battery reports in, rather than the separate log of presses. The record was corrected. Whenever an agent says something has not happened, it is worth asking how it knows.

The person receiving a reminder matters too. In October, I was in Florida while Trebor was home with Mika, yet the feeding warning still came to me more than a thousand miles away. We had to make reminders respond to who was actually home. That small failure points to a

larger truth: a household is not one user. It is several people sharing some things while keeping other things private.

Our kitchen tablet recognizes people and greets each with one brief line, at most once every three hours, and stays quiet overnight. At first, if Kim followed me into the kitchen, only one of us was noticed. Duncan was changed to keep watching for a few minutes, so it could greet both. The wall can also hold a message: I can ask it to tell Kim something when it next recognizes her. A message expires after a week if she does not pass by. The first attempt failed because it went to an old speaker; once routed through the kitchen wall, it worked.

The shared wall also created a privacy problem. Each person has a private memory, but a conversation at the wall was saved under whoever spoke last. My question and answer ended up in Kim’s memory because she had spoken last. A daily review of the records caught the mistake. Now, conversations are split whenever the speaker changes, and each person’s words go only into that person’s private memory. Shared access does not mean shared privacy.

Other household choices need the same care. Duncan once mispronounced our son Trebor’s name as “Tree,” “Tray,” or “Trih.” We added a pronunciation note—“Treb-or,” with “Treb rhymes with web”—so greetings and reminders say it correctly. When Treb objected to the shopping lists, Duncan passed his suggestions along; the rebuilt lists group items by category, let the family star regular purchases, and make quantities easy to change. We also decided who may change shared resources: for example, only adults can change Away rules, and a family photo can be removed only by its contributor or the household owner. For each feature, ask who may see it and who may change it.

Connected devices bring a different kind of shared responsibility. When autumn arrived, I wanted Duncan to control the thermostats. The truth was that neither system was fully in charge: the thermostat’s own schedule and learning features could change temperatures between Duncan’s scheduled adjustments. The record showed bedrooms dropping to 56 degrees at odd times, and our summer cooling schedule was now heating rooms to 75. A bug duplicated schedules when they were edited, too. We replaced the old rules with a winter schedule, room by room, then discovered that the thermostat’s independent features were still changing temperatures. I had to turn those off in the manufacturer’s app. Two brains on one thermostat were one too many.

While I was in Florida, Duncan began monitoring the Massachusetts house. Eight sensors in four rooms reported temperature and humidity, while a separate check-in system warned the adults if the house stopped communicating with our servers. We tested the outage alert against the previous month’s records, including a 99-minute power cut at the Florida house, and added an app card showing what was monitored and the latest readings. When a loss of contact could make a leak appear “cleared,” we changed the system so that being unable to

reach the house for more than ten minutes produces an urgent alert instead.

Away mode grew from the misplaced dog reminders. Each person writes rules for when they are away and when everyone is away. Once the household marks who is leaving and for how long, Duncan turns applicable rules into a checklist, identifies anything it cannot do, and waits for someone to press Start. When people return, settings are restored, except where someone changed something manually while they were gone. The mode has limits that no household rule can override: it does not change smoke, carbon monoxide, leak, freeze, alarm, or security alerts, turn off a thermostat, let the heat fall below 50 degrees, raise cooling above 85, or touch the car.

Duncan can report whether our car is locked, its battery charge, and its location, but its first “locked” report was wrong. The sensor’s “on” state meant unlocked, and the car’s readings were not live; the displayed status had last been reported at 3:07. The app now explains each reading’s age and offers a request for a fresh report. Our bedroom television taught a simpler lesson. Connecting it to Duncan caused the streaming box to keep waking the TV and soundbar, even though Duncan sent no commands. We disconnected it, kept a note of the change, and returned to the small remote. Some devices are better left alone.

Money is where a household agent must be useful and especially careful. Duncan can read, sort, compare, remind and research, but it does not spend. At the start of each month, it reminds us to complete the household books, then reviews the previous month. Its first review flagged recurring charges that had risen, including a utility bill nearly two hundred dollars higher than the month before. The information was on the statements; the useful part was having someone notice it.

The books also showed how automated sorting can become confidently wrong. In September, too many purchases had landed in “Entertainment, Dining out.” A rule created months earlier classified anything paid for by tapping a phone as dining, so it swept in groceries, pharmacy purchases, hardware and gas. Because the system favored rules that matched often, each mistaken match made the bad rule stronger. Meanwhile, the correct hardware-store rule never matched because the bank abbreviated the store’s name. I identified nine retailers and their proper categories, and Duncan moved 119 entries without disturbing anything I had already confirmed. The system was changed to favor specific rules over popular ones. The lesson is to review categories: an automation left unchecked can grow more wrong over time.

For our home theater, Duncan compared four rows of seats against a 120-inch wall. It recommended a 117-inch option rather than one that fit exactly, explaining that baseboards, outlets, uneven walls and the difficulty of moving the seats all called for room to spare. It prepared the choices, but asked my approval before emailing them to Kim, and bought nothing. That same restraint applies to account information. Give an agent only what its task requires; keep passwords and passcodes out of messages, and store sign-in details

separately from general memory.

In September, we added specialist advisors for financial, tax, legal, insurance, estate, real-estate, home, medical and travel questions, with a car advisor soon after. They share one household file rather than keeping separate memories, so the insurance advisor can draw on facts about the roof and the budget. Each asks only a few setup questions that could change its advice: a legal advisor needs the state, a home advisor the age of major systems, and a travel advisor passport expiry dates, not passport numbers.

The advisors also have firm limits. The medical advisor does not diagnose or alter medication and flags urgent concerns. The legal advisor names the relevant state and makes clear it is not a lawyer. Insurance advice cannot claim something is covered unless the policy has actually been read; estate questions include details that can override a will, such as account beneficiaries and how a house is titled. The guiding rule is that an advisor does the preparation and a licensed professional reviews anything important before it is signed, filed or relied upon. It can organize facts, read documents and prepare questions, but accountability still belongs to a qualified human.

The car advisor grew from a request for a conversation, not another form. I wanted to be able to begin with “I don’t know the difference between hybrid, plug-in hybrid, and EV” and get a plain explanation. It uses government data on fuel economy, safety ratings and recalls to help narrow the field. Later it was divided into buying, selling and owning. Its selling guidance covers valuation, listings, scams, test drives, payment and state paperwork, but it never lists the car or contacts a buyer. The person still does that. With any advisor, give it the facts that change the answer, ask what it is assuming, and ask what a professional should check. Keep sensitive numbers out, and make sure advice is based on your household rather than an imagined average one.

Failures showed why transparency matters as much as capability. Duncan once dropped Kim’s direct request to email information when she ended a voice call; a cautious rule had mistaken her instruction for an unanswered offer. Later, when the outside AI service was overloaded, it created her project and finished the research, but failed to respond in the moment and reported, falsely, “Nothing was done.” The system was changed so requests survive the end of a conversation and a failure report states exactly what succeeded, without repeating actions already completed. An agent must tell you what actually happened, not merely that something went wrong.

Some limits are inherent. Duncan could confirm that the theater receiver was on and receiving a signal, but could not see the projector’s picture. A phone company still needed the account holder to prove their identity. Contractors and offices could still fail to answer. The agent also depends on outside services, internet connections and devices whose readings may be delayed or wrong. Since September, a separate agent has reviewed the previous day’s

conversations and logs each morning, traced faults, tested fixes and emailed a report. It is forbidden to act on the conversations it reads; its job is to find what broke, not to repeat anyone’s request. Someone must notice when the system fails.

And there are household responsibilities that should remain human. Meals, cleaning, laundry, driving and being present, caring for one another, and deciding who owns which job are not simply tasks to automate. An agent may help at the edges, but these are the substance of living together. The point of taking some remembering off people is to leave them more attention for each other.

There are three ways to get a household’s help. A general assistant may be enough for one person’s email and calendar, and costs little. Hiring a house manager or part-time assistant brings human help with tasks such as waiting for deliveries, but costs more and cannot monitor an empty house at three in the morning. A looked-after service such as PrivateChief adds household identities, privacy, devices, routines and maintenance around the AI model. At the time of writing, membership costs $199 a month plus a one-time $299 setup fee. Hosting and maintenance are included; a Home Assistant box and kitchen wall units are optional. The AI model is billed separately, either through your own OpenAI account or through the service at cost plus 30 percent. The apps are free with membership, and the stated cancellation terms return the setup fee if you cancel before the Chief goes live, or the setup fee and first month if you cancel within 30 days. The service can alert you to household conditions, but it is not an emergency or life-safety system.

Whatever route you choose, ask where each person’s words go, what happens when the system fails, whether it can spend or contact people without approval, what it will never touch, how old its readings are, who checks the records and how often, and whether you can leave with usable copies of your information. Then begin small: list what you had to remember, create a short morning briefing, delegate one errand with a clear finish and stop rule, replace one uncertain guess with a deliberate record, and discuss household privacy. Ask how an assistant knows what it reports, and decide which responsibilities you want to keep. The right system will not run your family. It will hold some of the details, leaving people more room to care for one another—and, if Mika is any guide, to know when the dog has already been fed.

Generated by BookGist.ai • October 2026

www.bookgist.ai