

Behavioral Design for Compliance Systems

From Catching Failure to Enabling Success

Compliance systems can be designed to catch people failing or to help people succeed. Behavioral science offers a systematic framework for the latter. The question is not whether people are motivated to maintain coverage. The question is whether system design converts motivation into action.

The 18.5 million adults who will face Medicaid work requirements beginning December 2026 overwhelmingly want to keep their healthcare coverage. Studies of similar populations consistently find that ***maintaining health insurance ranks among the highest priorities*** for low-income households, ahead of many other concerns that compete for attention and resources. The problem is not motivation. The problem is the gap between wanting something and achieving it.

This article examines that gap through the lens of behavioral science, drawing on research from psychology, economics, and decision science to explain why well-intentioned people fail to complete straightforward tasks. More importantly, it offers a theoretical framework for designing systems that close the gap rather than exploit it. Where traditional verification systems assume rational actors making deliberate choices, behaviorally-informed systems recognize humans as they actually are: ***cognitively limited, emotionally driven, present-focused, and context-dependent.***

The distinction matters enormously for policy. Systems designed around inaccurate models of human behavior produce predictable failures that policymakers then misattribute to participant deficiencies. Systems designed around accurate models produce compliance rates that reflect actual behavior rather than administrative navigation capacity. The choice between these approaches represents a philosophical commitment about what verification systems are for and who they are meant to serve.

The Intention-Action Gap

Half a century of behavioral research has established a fundamental insight that remains poorly integrated into policy design: ***people consistently fail to do things they genuinely intend to do.*** This intention-action gap is not a character flaw requiring correction but a predictable feature of human cognition requiring accommodation.

The research is extensive and replicated across cultures and contexts. Peter Gollwitzer's foundational studies found that people with strong goals fail to achieve them roughly half the time, even when they possess the knowledge, resources, and genuine desire to succeed. The gap appears in exercise behavior, medication adherence, financial planning, preventive healthcare, voting participation, and countless other domains where intention alone should suffice but does not.

Several well-documented mechanisms explain why intentions fail to produce actions.

Present bias describes our tendency to weight immediate experiences more heavily than future consequences, even when we intellectually recognize the future consequences as more important. The phenomenon appears across every measured population and persists even when people are aware of it. A mother facing competing demands knows that maintaining health coverage matters more than any single day's immediate pressures. But immediate pressures are concrete,

demanding, and right here. Future coverage loss is abstract and distant until suddenly it becomes the present, and by then the deadline has passed.

The economics literature, particularly work by David Laibson and Ted O'Donoghue, has formalized present bias through quasi-hyperbolic discounting models that explain why people systematically undersave for retirement, underutilize preventive care, and underinvest in behaviors with delayed benefits. The same models predict compliance failures in administrative systems that impose current costs to prevent future harms. ***The structure of work requirements, where current effort prevents future coverage loss, is precisely designed to trigger present bias.***

The planning fallacy, identified by Daniel Kahneman and Amos Tversky and replicated in dozens of subsequent studies, causes people to systematically underestimate how long tasks will take and overestimate their future available time. When someone receives a notice with a 45-day deadline and believes they will complete the task "soon," they are not being dishonest. They are experiencing the same optimistic forecasting error that affects everyone. We plan for best-case scenarios and then encounter average-case reality.

The planning fallacy persists even when people have direct experience with similar tasks taking longer than expected. Somehow the current task feels different. This time will be easier. The evidence from prior deadline failures does not transfer to future deadline confidence. For work requirement verification, the implication is that ***long deadlines may actually be worse than short ones***, because they create confidence that action can be deferred while providing more time for the task to slip from attention entirely.

Cognitive load and its close cousin ***decision fatigue*** compound these tendencies. Our capacity for processing information and making decisions is not unlimited. It functions more like a battery that depletes with use than a muscle that grows with exercise. Every bureaucratic task competes for the same finite resource. When someone is simultaneously navigating employment uncertainty, childcare logistics, healthcare appointments, utility bill negotiations, and family emergencies, the verification portal becomes one more demand on exhausted cognitive reserves.

Sendhil Mullainathan and Eldar Shafir's work on scarcity has demonstrated that ***poverty itself imposes a cognitive tax*** equivalent to losing 13-14 IQ points, roughly the effect of going a full night without sleep. The populations subject to work requirements, by definition low-income adults, carry this cognitive burden chronically. Systems designed to function smoothly for cognitively abundant users fail predictably for cognitively depleted ones. The failure is not user error. It is design error.

Traditional verification systems ignore these mechanisms entirely. They assume that providing clear information creates understanding, that understanding creates intention, that intention creates action, and that action follows from rational calculation of costs and benefits. Each link in this chain has been empirically demolished. Yet the assumption persists in system design because it is simpler, because it assigns responsibility to participants rather than designers, and because challenging it requires acknowledging that ***compliance failures often reflect system failures rather than personal failures.***

The Power of Defaults

The most powerful tool in behavioral design is also the simplest: the default. When systems require affirmative action to obtain a desired outcome, participation collapses. When systems presume the desired outcome and require affirmative action to deviate, participation soars.



The evidence is overwhelming and replicated across domains. Brigitte Madrian and Dennis Shea's landmark 2001 study of 401(k) retirement savings found that automatic enrollment increased participation by roughly 50 percentage points. The same employees, with the same incomes and the same financial needs, saved dramatically more when saving was the default rather than a choice they had to make. Nothing changed except the default. No incentives shifted. No information improved. The path of least resistance simply pointed toward saving rather than away from it.

Eric Johnson and Daniel Goldstein's 2003 analysis of organ donation rates across European countries found differences exceeding 60 percentage points based solely on whether countries used opt-in or opt-out defaults. In Germany, where citizens must actively register as donors, roughly 12% were registered. In Austria, a culturally similar country where citizens are presumed donors unless they opt out, registration exceeded 99%. The default determined the outcome more powerfully than any other factor the researchers could identify.

The mechanism is not laziness but cognitive economics. Defaults reduce decision-making costs. They eliminate the need to gather information, weigh options, and commit to a choice. They provide implicit guidance about what the system designer considers appropriate. They leverage the status quo bias that makes any change feel costlier than maintaining current conditions. For decisions that are complex, emotionally difficult, or simply not immediately pressing, defaults often determine outcomes regardless of underlying preferences.

Current work requirement systems invert this principle. They require affirmative action at every step: logging into portals, navigating interfaces, locating documents, uploading files, confirming submissions. Each step is an opportunity for dropout. Each step loses people who intended to comply but encountered friction they could not overcome. The default is not coverage but coverage loss. ***Inaction produces termination. The system is designed to catch failure rather than enable success.***

Behaviorally-informed design would flip this architecture. Instead of requiring members to prove compliance, systems would presume compliance unless evidence suggests otherwise. Instead of demanding documentation at every reporting period, systems would automatically verify through data matching and require member action only when automated verification fails. Instead of terminating coverage immediately upon missed deadlines, systems would maintain coverage through reasonable cure periods that assume good faith.

The objection to defaults and presumptive eligibility typically centers on program integrity. If we do not verify monthly, how do we know people are actually working? The answer is that monthly verification does not reliably establish that people are working. It establishes that people successfully navigated the verification system. ***These are different populations with different characteristics.*** Arkansas demonstrated conclusively that most people who lost coverage under work requirements were actually working or qualified for exemptions. They failed to prove what they were doing, not failed to do it.

Georgia's Pathways to Coverage evolution illustrates the practical implications. The program launched in July 2023 requiring monthly verification. By late 2024, only about 5,500 people had enrolled despite an estimated 240,000 eligible residents. The state's waiver extension proposal eliminates monthly reporting in favor of annual verification. This is not abandoning work

requirements. It is recognizing that monthly reporting creates twelve opportunities per year for compliant people to lose coverage through administrative failure rather than work failure.

Friction Mapping and Sludge Reduction

Every administrative system contains friction points where users encounter difficulty completing desired actions. Cass Sunstein's concept of "sludge" captures friction that serves no legitimate purpose, that imposes costs without corresponding benefits, that exists because designers failed to consider user experience or actively wanted to discourage participation.

Work requirement verification systems are rich in sludge. **Portal access** creates the first barrier. Members must have internet access, devices capable of running the portal, and sufficient technical literacy to navigate digital interfaces. Those who rely on smartphones face screens designed for desktop displays, forms that time out before completion, and document management functions poorly suited to mobile use. Library computers offer access but impose time limits, privacy constraints, and the requirement to travel to physical locations during business hours.

Password requirements create persistent friction. Security best practices demand complex passwords that are difficult to remember, changed periodically, and unique across systems. Members who access work requirement portals infrequently, perhaps once per month at most, predictably forget credentials between sessions. Password reset processes add steps, create delays, introduce additional failure points, and sometimes require verification through channels that are themselves unreliable for the affected population.

Document specifications multiply problems. When systems require specific formats, file sizes, naming conventions, or resolution requirements, every deviation produces rejection. Members who photograph pay stubs on phones may create files that exceed size limits. Those who scan documents may inadvertently produce formats the system cannot process. Error messages may be cryptic or absent entirely, leaving users uncertain why their submission failed and what they need to do differently.

Unclear instructions generate confusion that wastes time and causes errors. Bureaucratic language like "verification of qualifying activities during the applicable reporting period" may be legally precise but communicatively useless for someone simply trying to prove they work. Instructions that reference form numbers, database codes, regulatory citations, or program acronyms assume familiarity that most participants lack.

The behavioral design response is **friction mapping**: systematically documenting every step users must complete, identifying where users drop off, measuring time required at each stage, testing whether typical users can complete the process without assistance, and relentlessly eliminating friction that serves no legitimate purpose.

Friction is not inherently bad. Some friction is necessary for program integrity. Some friction protects against fraud. Some friction ensures that decisions are deliberate rather than accidental. The design question is **whether friction appears in the right places**. Currently, most systems make enrollment difficult and disenrollment automatic. Behavioral design principles suggest the opposite: **make enrollment easy and disenrollment hard**. Require supervisor review before termination. Mandate multiple contact attempts. Create cooling-off periods. Add friction that prevents inappropriate coverage loss rather than friction that prevents appropriate coverage maintenance.

Timing, Triggers, and Implementation Intentions

When we ask people to do something matters as much as how we ask. Behavioral research identifies optimal timing for behavior change and demonstrates that poorly timed communications can be worse than no communications at all.

Fresh start moments represent periods when people are naturally more receptive to new behaviors. Hengchen Dai, Katherine Milkman, and Jason Riis documented the "fresh start effect" across multiple datasets: people are more likely to pursue goals at the beginning of new time periods, including new weeks, new months, new years, and after birthdays. These temporal landmarks create psychological separations between a "past self" who may have failed and a "future self" who can succeed. Deadlines aligned with fresh starts produce higher response rates than arbitrary dates chosen for administrative convenience.

Implementation intentions, the research program initiated by Peter Gollwitzer, represent one of the most robust findings in behavioral science. Asking people to specify when, where, and how they will complete a task roughly doubles completion rates compared to providing the same information without the planning prompt. "I will submit my verification" produces different results than "I will submit my verification on Saturday morning at 10 a.m. using my phone at the kitchen table."

The mechanism involves both memory and motivation. **Specifying situational cues creates automatic associations** between the cue and the intended behavior. When Saturday morning arrives, the plan activates. Specifying the behavior eliminates the need to decide what to do when the moment comes. The combination reduces cognitive demands at the moment of action, when cognitive resources may be depleted, by front-loading the decisions to a moment when resources are available.

Katherine Milkman and colleagues' 2011 study of flu vaccination applied implementation intentions at scale. Employees who received a prompt to write down the date and time they planned to get vaccinated showed 4.2 percentage points higher vaccination rates than employees who received identical information without the planning prompt. The effect was largest when vaccination was available only on a single day, when the implementation intention created a specific commitment rather than a general intention to act eventually.

Work requirement systems rarely incorporate these insights. Notices provide deadlines without prompting planning. Communications arrive on whatever schedule fits administrative convenience. Reminders may come too early to be actionable or too late to allow response.

Escalating contact sequences that begin with SMS, progress to email, advance to phone calls, and finally deploy in-person outreach represent behavioral design applied to channel optimization, but few systems employ such sequences systematically.

The technology for implementation intentions exists and costs essentially nothing to deploy. A verification notice that includes "When will you submit your verification? Write down the day and time here:" leverages decades of research at zero marginal cost. Calendar integration that allows members to schedule submissions and receive automated prompts at self-selected times converts intention into behavioral infrastructure. **The barrier is not technical. It is philosophical.** Current systems assume members should adapt to bureaucratic requirements. Behaviorally-informed systems assume bureaucratic requirements should adapt to how members actually behave.

Framing, Loss Aversion, and Social Proof

How choices are described influences which choices people make. **This framing effect**, central to Kahneman and Tversky's prospect theory, operates largely outside conscious awareness, shaping decisions through language, emphasis, and context rather than argument or logic.

Loss aversion describes the empirically robust finding that people weight potential losses more heavily than equivalent gains. Losing \$100 feels worse than gaining \$100 feels good, typically by a ratio of roughly 2:1. This asymmetry has profound implications for how systems communicate with members.

Current work requirement communications typically emphasize eligibility requirements in gain-frame terms. "To maintain your eligibility, you must verify your work hours by..." This framing presents compliance as achieving a bureaucratic goal. Behaviorally-informed framing emphasizes loss prevention. "Your healthcare coverage will end unless you submit verification by..." The same deadline, framed as loss rather than gain, produces higher response rates.

The specific language matters and can be tested empirically. "Keep your coverage" outperforms "maintain your eligibility." "Protect your healthcare" outperforms "comply with requirements." "Don't lose your prescription benefits" outperforms "remain enrolled in Medicaid." **The underlying information is identical. The psychological impact differs substantially.**

Social proof leverages our tendency to look to others' behavior as guidance for our own. When we learn that most people like us have taken an action, we become more likely to take it ourselves. Robert Cialdini's research on influence has documented social proof effects across contexts from hotel towel reuse to tax compliance to energy conservation.

In work requirement communications, social proof might take forms like: "Most people in your area complete their verification within five days of receiving this notice." This provides a behavioral benchmark suggesting that completion is normal and that the timeline is achievable. However, social proof must be deployed carefully. **Telling people that 40% of members fail to verify on time might normalize non-compliance** rather than prevent it. Effective social proof emphasizes positive behaviors and suggests the recipient is similar to successful compliers rather than to those who struggle.

Normalizing compliance means presenting work requirements as standard expectations that most people meet, rather than as punitive measures assuming members are trying to cheat. Communications that treat members as cooperative partners, providing reminders and assistance to help them succeed, produce better results than communications that treat members as suspected violators requiring surveillance. The framing creates a self-fulfilling prophecy. **Treat people as trustworthy and they tend to act trustworthy. Treat them as suspects and they disengage.**

Commitment Devices and Accountability Structures

Some behavioral interventions help people constrain their future selves. These commitment devices recognize that our current intentions and future actions often diverge, and that steps taken now can bind us to better outcomes later.

Self-designed reminders allow members to specify when and how they want to be contacted. A member who knows she is most likely to complete administrative tasks on Sunday evenings can request reminders for Sunday at 6 p.m. A member who responds to texts but ignores emails can

specify his preferred channel. Self-designed reminders leverage members' self-knowledge about their own behavioral patterns while increasing ownership of the compliance process.

Accountability partnerships extend social proof into personal relationships. Systems that allow members to designate someone who will receive notification if verification remains incomplete create social accountability reinforcing individual intention. The designated person might be a family member, a navigator, a community organization contact, or a healthcare provider. The notification need not share private information. It simply indicates that the member authorized this accountability mechanism and that action is needed.

Public commitments strengthen intentions by making them socially visible. Community organizations working with members can facilitate commitment moments where members state their intention to maintain compliance. The mechanism operates through both cognitive consistency (wanting to act in accordance with stated intentions) and social pressure (not wanting to be seen as someone who fails to follow through).

Commitment device limitations must be recognized. These techniques work best when people want to do something but anticipate difficulty doing it. They work poorly when external barriers rather than internal conflicts prevent action. A member who lacks internet access does not benefit from self-designed digital reminders. A member whose employer refuses to provide documentation cannot overcome that barrier through personal commitment. **Commitment devices address intention-action gaps, not resource gaps or institutional failures.**

The Design Philosophy Beneath the Tools

The specific tools of behavioral design, defaults, friction mapping, timing optimization, framing, commitment devices, constitute a technology that can serve different masters. The same techniques that help people succeed can also help systems exclude people. **Sludge is behavioral design applied toward exclusionary ends.**

The distinction between systems that help people comply and systems that catch people failing reflects a deeper philosophical choice about **what verification systems are for**. If the purpose is to identify non-compliance, then system design should maximize detection sensitivity even at the cost of false positives, terminating coverage for people who were actually compliant but could not prove it. If the purpose is to maintain coverage for eligible people while identifying genuine non-compliance, then system design should minimize false positives even at some cost to detection sensitivity.

These purposes are not equally legitimate. **Medicaid is a healthcare program, not a compliance testing program.** Its purpose is to provide healthcare coverage to eligible people. Work requirements represent one test of eligibility, but they do not transform the program's fundamental purpose. Systems designed as if catching non-compliance were the primary goal misconceive what the program exists to do.

Behavioral design offers a framework for aligning system design with program purpose. By understanding how human cognition actually works, by recognizing the intention-action gap as a feature rather than a flaw, by designing systems that convert genuine motivation into successful action, states can implement work requirements that distinguish between those unwilling to work and those unable to navigate bureaucracy. **The former is a legitimate policy concern. The latter is an administrative failure.**



The behavioral science is clear. The technology is available. The cost is minimal and often negative, since helping people maintain coverage reduces churning costs, call center volume, appeals processing, and the administrative overhead of re-enrollment. **The barrier is not knowledge or resources. The barrier is the assumption that compliance failures represent moral failures rather than design failures**, and the policy preference for catching cheaters over helping compliers.

That assumption can be changed. It requires recognizing that the populations subject to work requirements are not fundamentally different from the policymakers designing systems for them. They want to keep their healthcare. They intend to comply with requirements. They encounter the same cognitive limitations, the same present bias, the same planning fallacy, the same decision fatigue that affects everyone. **The difference is that they face these universal human limitations while simultaneously navigating poverty, instability, and systems designed without their lives in mind.**

Systems designed for actual humans rather than idealized rational actors will produce different outcomes. Whether states choose to build such systems depends on whether they want work requirements to sort people by compliance capacity or by actual work behavior. The behavioral science cannot make that choice for them. It can only clarify what that choice entails.

References

1. Thaler RH, Sunstein CR. *Nudge: Improving Decisions About Health, Wealth, and Happiness*. Yale University Press. 2008.
2. Thaler RH, Sunstein CR, Balz JP. "Choice Architecture." In Shafir E, ed. *The Behavioral Foundations of Public Policy*. Princeton University Press. 2013.
3. Madrian BC, Shea DF. "The power of suggestion: Inertia in 401(k) participation and savings behavior." *Quarterly Journal of Economics*. 2001;116(4):1149-1187.
4. Johnson EJ, Goldstein D. "Do defaults save lives?" *Science*. 2003;302(5649):1338-1339.
5. Gollwitzer PM. "Implementation intentions: Strong effects of simple plans." *American Psychologist*. 1999;54(7):493-503.
6. Gollwitzer PM, Sheeran P. "Implementation intentions and goal achievement: A meta-analysis of effects and processes." *Advances in Experimental Social Psychology*. 2006;38:69-119.
7. Milkman KL, Beshears J, Choi JJ, Laibson D, Madrian BC. "Using implementation intentions prompts to enhance influenza vaccination rates." *Proceedings of the National Academy of Sciences*. 2011;108(26):10415-10420.
8. Kahneman D, Tversky A. "Prospect theory: An analysis of decision under risk." *Econometrica*. 1979;47(2):263-291.
9. Tversky A, Kahneman D. "Loss aversion in riskless choice: A reference-dependent model." *Quarterly Journal of Economics*. 1991;106(4):1039-1061.
10. Mullainathan S, Shafir E. *Scarcity: Why Having Too Little Means So Much*. Times Books. 2013.
11. Dai H, Milkman KL, Riis J. "The fresh start effect: Temporal landmarks motivate aspirational behavior." *Management Science*. 2014;60(10):2563-2582.
12. Behavioural Insights Team. "EAST: Four Simple Ways to Apply Behavioural Insights." UK Cabinet Office. 2014.



13. Sunstein CR. "Sludge and ordeals." *Duke Law Journal*. 2020;68:1843-1884.
14. Herd P, Moynihan DP. *Administrative Burden: Policymaking by Other Means*. Russell Sage Foundation. 2018.
15. Rogers T, Milkman KL, John LK, Norton MI. "Beyond good intentions: Prompting people to make plans improves follow-through on important tasks." *Behavioral Science & Policy*. 2015;1(2):33-41.
16. O'Donoghue T, Rabin M. "Doing it now or later." *American Economic Review*. 1999;89(1):103-124.
17. Benartzi S, Beshears J, Milkman KL, Sunstein CR, Thaler RH, Shankar M, Tucker-Ray W, Congdon WJ, Galing S. "Should governments invest more in nudging?" *Psychological Science*. 2017;28(8):1041-1055.
18. Cialdini RB. *Influence: The Psychology of Persuasion*. Harper Business. 2006.
19. Sommers BD, Goldman AL, Blendon RJ, Orav EJ, Epstein AM. "Medicaid work requirements: Results from the first year in Arkansas." *New England Journal of Medicine*. 2019;381(11):1073-1082.
20. Mani A, Mullainathan S, Shafir E, Zhao J. "Poverty impedes cognitive function." *Science*. 2013;341(6149):976-980.

Syam Adusumilli is Chief Evangelist at GroundGame.Health, a social determinants of health platform. This article is part of a comprehensive series examining Medicaid work requirements under the One Big Beautiful Bill Act.