

BRAKED

by EPISTEMIST LABS

A behavioral microtraining instrument
for urge and impulse control



PRODUCT WHITEPAPER
October 2026

BRAKED

Delay as a Trainable Capacity

A design and evidence account of a cognitive-behavioral impulse training instrument for self-directed and practitioner-supported use alike, built for complete user privacy and autonomy



by **EPISTEMIST LABS**

Onur Yavuz, Psychologist (MSc University of Zurich, Switzerland)

Product Whitepaper · October 2026 · V. 1.3

Table of Contents

1. A Note from the Author	1
2. What “Evidence-Based” Currently Means in This Market	5
3. What BRAKED Is	8
4. The Training Model: Delay as a Trainable Capacity	14
5. From Construct to Feature	21
6. Seven Domains, Seven Trained Skills	26
7. What the Coaching Material Actually Looks Like	32
8. Architecture as Position	36
9. Using BRAKED Alongside Professional Work	38
10. Safety and Scope	41
11. What I Have Chosen Not to Claim	43
12. Open Questions and What Validation Would Require	46
13. Closing Note	48
References	49
Imprint and Contact	56

1. A Note from the Author

Every approach to behavioral change I have worked with, taught, or been taught eventually asks the same thing of the person doing the change. Whatever the school, whatever the presenting problem, whether or not a professional is in the room: notice the behavior while it happens and write it down, then look at the record with honesty a few weeks later.

This is the least glamorous part of the work and the part that most reliably decides whether it succeeds. The evidence is not marginal. Pooling 138 studies and close to 20,000 participants, Harkin et al. (2016) found that prompting people to monitor their position toward a goal moved them closer to reaching it, with a larger gain when the monitoring produced a physical record rather than remaining only in memory. In a meta-regression covering 122 intervention evaluations and a combined sample of 44,747, Michie et al. (2009) examined which intervention technique accounted for the differences between studies. Self-monitoring accounted for more of that variation than any other technique examined, and interventions that combined self-monitoring with at least one additional self-regulatory technique performed better than the other interventions, at 0.42 versus 0.26.

In private counseling and in organizational work alike, this is where I kept watching things quietly fall apart.

It was rarely unwillingness to change. The people I worked with wanted it. What defeated them was the gap between the moment of impulse and the tools available to record it. Pen and paper, a notes app, a timer in one place and a journal in another. By the time someone is in the moment that actually matters, they are not reaching for any of it. Even where the urge management behavior itself is exercised, the entry gets made later, from memory, or not at all. The memory version, made presumably out of shame, is the worst of all: a record now exists, and it is wrong.

I long suspected this was widespread. It turns out to be measurable, and the measurement deserves its own spotlight.

In a study by Stone et al. (2002), researchers gave patients paper diaries to track pain. The binders looked ordinary, but each had a photosensor hidden in the spine that recorded whenever the cover was opened. Patients reported having followed the schedule 90% of the time. The sensor data told a different story: only 11% of entries fell within the target time window. Entries appeared in the record for at least one day on which the diary was actually never opened, a pattern found in three quarters of the patients.

Likely nobody was lying. Reconstruction can feel like remembering right up until someone checks. But a reconstructed record causes more harm than a missing one. It creates a fictitious history, and decisions are then made from it.

The obvious remedy is to move the capture onto the device already in the person's hand. The same study included a second arm using a palmtop computer, which timestamped each entry as it was made. In that arm, the reported and actual figures converged at around 94%. This solves the reconstruction problem but introduces another one, and it is one I could not in good conscience ask a client to accept.

Huckvale et al. (2019) assessed 36 top-ranked apps for depression and smoking cessation, finding that 33 transmitted user data to third parties, including 29 that sent data to Google or Facebook. Only 12 of the apps accurately disclosed this practice in their privacy policies. Their recommendation to clinicians was direct: privacy policies do not reliably indicate what an app actually does, so the safer assumption is that data leaves the device. In 2023, the Federal Trade Commission brought an enforcement action to an online counseling provider that resulted in \$7.8 million in a settlement and barred it from further disclosure of health data for advertising after charges that the provider had passed users' email addresses, IP addresses, and health questionnaire answers to advertising platforms while assuring users on the same pages that their answers were private (Federal Trade Commission, 2023).

For a psychologist, this converts a question of convenience into a question of data sovereignty. I did not want to recommend a tool whose usefulness depended on a client's willingness to be surveilled.

There is a third problem, and it is the one I find most avoidable. Many tracking apps motivate usage through streaks, badges and nudges, which work in a narrow sense. But a broken streak might imply something about "personal performance" rather than about the action, and that distinction carries weight. Guilt, technically speaking, means feeling bad about a specific behavior; it tends to be adaptive and points toward repair. Shame means feeling bad about the self (Tangney et al., 2005; Tangney et al., 2011). Across three separate samples, Dearing et al. (2005) examined shame and guilt in relation to substance use and found that proneness to shame was generally associated with more *substance-use problems*, while proneness to guilt was associated with fewer problems or with none. In a meta-analysis of 18 samples, shame was associated with *substance-use problems* at $r = 0.16$, an effect size comparable to the established link between depression and substance use (Luoma et al., 2019).

The same review found no association between shame and the amount of substance use itself across 14 samples ($r = 0.00$), so the link is specific to substance-related problems rather than to use as such.


A prospective study of adults newly sober from alcohol found that nonverbal displays of shame, but not shame reported about oneself, predicted a return to drinking, the severity of that return, and declining health over the following months (Randles & Tracy, 2013). Shame also tends to prompt withdrawal and concealment, where guilt tends to prompt repair (Tangney et al., 2005; Tangney et al., 2011). The studies link shame to worse *substance-use outcomes*; applying that link to the behaviors BRAKED addresses is my own inference in the absence of other evidence.

A design that turns a missed day into a verdict on the person's performance or discipline is building on the less useful of the two, and it recruits the same mechanism that risks a single lapse into abandonment of the whole effort.

The other answer on offer is to remove the choice altogether: block the app, lock the screen, make the behavior impossible. I understand the appeal and I would not claim these never work. My objection is that they train the wrong thing. A person who was prevented from acting has not practiced anything. There is no performance to draw confidence from afterwards, and direct successful performance builds a sense of one's own capability more powerfully than any other source (Bandura, 1977).

What I wanted was narrow and specific: something that kept the homework while making it survivable and genuinely useful. An instrument available in the actual moment when an urge arises, rather than reconstructed hours later. Something that stored what it recorded on the person's own device and nowhere else. Something that treated a lapse as information rather than a verdict. And something that left the decision, every time, with the person making it.

BRAKED is that instrument, and it stays narrow by design: one mechanism, practiced the way a single movement is trained, rather than a general answer to how someone feels. It also stays narrow within its category. Instead of the *everything app* the market tends to produce, with mindfulness content, affirmations, breathing exercises, and habit trackers bundled into one interface that solves no single problem particularly well, it targets a fixed set of behaviors, each with its own graded rungs, and asks for almost nothing to get started.



The first screen of the onboarding process sets the instrument scope for the user

The training instrument is free and stays free: delay timers, training statistics, session history and calendar view, printable private reports, unlimited profile resets and training data purges, and unlimited export and import of your own data. No account, no login. A subscription does not unlock the core mechanism. It buys professional capacity: weekly coaching material written by a psychologist at EPISTEMIST for the behaviors someone is actively training and, within the annual plan, a personally written “Deep Dive” psychological consultation of up to 1,300 words, included once per subscription year, on a main topic of their choosing. Annual subscribers also receive a price advantage on priority-tier individual live-online consultation sessions for the active plan duration. Both the Deep Dive and priority-tier live sessions can be booked under a pseudonym. The split is deliberate. A mental health professional colleague should be able to recommend this to a client without recommending a purchase.

Two things happen within that one mechanism. The first is training: holding a delay, rung by rung. The second is building a factual, judgment-free baseline of your own practice in whichever domain you are training, so that a session next month can be read against your own starting point rather than anyone else’s benchmark. In theory, such data can be assembled manually, with enough discipline and enough paper. Most people do not, though, which is exactly why BRAKED exists. What the instrument provides to turn the theoretical into reality is friction removal through convenience: it keeps the record automatically at multiple levels of detail, entirely in private, and returns it in a form a person can actually read.

I wrote this for three readers. For fellow practitioners, psychologists, psychotherapists, psychiatrists and other physicians, as well as pedagogical counselors, who may consider recommending BRAKED and are entitled to see the reasoning first. For colleagues in research, who will want to know which claims rest on evidence, which rest on clinical convention, and which are decisions I made and am prepared to defend as decisions. And for the end user of the app, who deserves the same account rather than a simplified one.

What I offer is a qualified judgment formed over more than a decade of professional practice as a psychologist, and grounded in the research set out below: a graded delay, practiced consistently, is a sound way to train this capacity; even more successfully when combined with domain-specific behavioral strategies involving episodic future thinking (Atance & O’Neill, 2001), connectedness to the future self (Ersner-Hersfield et al., 2009), and construal level manipulation (Trope & Liberman, 2003). I built BRAKED on that judgment, put my name on it as the accountable party, and expect it to deliver positive change for most people who use it as intended. What follows sets out what BRAKED does, where the evidence was found, where it was not, and how I have used professional judgment in the absence of specific peer-reviewed research and said so.

Onur Yavuz, Psychologist (MSc University of Zurich, Switzerland)
Founder at EPISTEMIST
Developer of the BRAKED App | EPISTEMIST LABS

2. What “Evidence-Based” Currently Means in This Market

Search any app store for a mental health tool and the promises arrive within the first screen of results: “scientific”, “CBT”, “mindfulness”, “depression test”, “developed by experts”; then there are those that package CBT and (AI) psychoanalysis (for dream interpretation) in one. Each term is intended as a quality and authority signal and most people treat it as one.

The question is what that signal is currently worth.

Marshall et al. (2020) searched the major app stores for products presented as treatments for anxiety or depression and identified 293 candidates. Of those, 162, a little over half, stated in their store descriptions that they were based on a recognized therapeutic framework. The authors then searched for published evidence that these apps actually helped anyone and found such evidence for only 10 of them. Thus, 6.2% of the apps making the claim could point to supporting evidence, representing around 3% of the category as a whole.

The same group's companion paper, drawing on the same set of apps, adds a detail that matters: of the ten apps with supporting research, only three were backed by research conducted independently of the people who developed or profited from them (Marshall et al., 2019). Larsen et al. (2019) searched app stores for 1,435 candidate apps and coded the 73 top-ranked consumer titles for scientific claims, finding that scientific terminology was widely used in store marketing without the supporting basis that would justify those claims.

A more recent study using a different sample reached a related conclusion. Rickard et al. (2022) applied the American Psychiatric Association's own app-evaluation model to the 100 most visible depression, anxiety, and mood apps across both major stores. They found that store visibility was associated with popularity and commercial signals rather than the model's criteria for safety or evidence, and only one of the sampled apps met all five of the model's dimensions.

The content itself becomes thinner when examined against established treatment components. Wasil et al. (2019) assessed 27 widely downloaded depression and anxiety apps against a list of components found in protocols known to work. The app with the most components contained 12. Some contained none. The midpoint of the distribution was three components.

Wasil et al. (2020) repeated the analysis, this time weighting each app according to the number of people who opened it. Two products dominated the sample, accounting between them for 90% of monthly active users. Once weighted by actual use, the pattern reverses. Mindfulness material appeared in 37% of the apps but reached 96% of users. Cognitive restructuring appeared in 22% of the apps but reached 2% of users. Exposure therapy, the primary active ingredient in treating anxiety, was present in only 7% of the reviewed apps and reached a mere 0.0004% of total users.

Hence, the techniques with the strongest support in the treatment literature are, in practice, the ones almost nobody is receiving. The market has largely failed to select for these techniques, while store ranking algorithms favor downloads and reviews rather than clinical content.

None of this means the underlying intervention class is inert. A meta-analysis of 176 randomized controlled trials found that standalone apps produced small but significant effects on depression ($g = 0.28$) and generalized anxiety ($g = 0.26$) against control conditions (Linardon, Torous, et al., 2024). The market failure documented above is therefore not that apps cannot work, but that the apps shown to work and the apps people download are frequently not the same apps. A 2026 review found that, among 81 apps with trial evidence, only 42, just over half, were publicly available at all, with no relationship between an app's demonstrated quality and whether the public could reach it (Bell et al., 2026).

I assume that people who install any app that requires regular check-ins without incentive or telemetry-induced nudges will stop using it within a few weeks, a trend consistent with empirical observations in the literature (e.g., Baumel et al., 2019). BRAKED is designed with features in response to this reality. Specifically, Section 8 details the architectural choices made to remove friction and support retention, though the utility of any behavioral tool ultimately remains conditional on active use.

This requirement is universal across all change modalities. Neither a digital instrument nor traditional consultation can operate as a passive treatment: without intrinsic motivation and deliberate practice, sitting in a mental health practice for an hour a week yields no meaningful outcome either. The tool (digital or human) can structure the work, but it cannot perform it.

Taken together, the findings in this section describe a market in which the claim of being evidence-based carries close to no information. The claim is cheap to make, expensive to verify, and nobody is checking. A reader encountering it on a store listing has no way to distinguish an app built on a treatment protocol from one built on a phrase.

A newer segment of this market falls outside the older studies because it did not exist when they were conducted. Hua et al. (2025) reviewed 160 studies of AI mental health chatbots published between 2020 and 2024. They classified the systems by architecture, rule-based, machine-learning-based, and LLM-based, and recommended a three-tier framework for evaluating them: bench testing, pilot feasibility testing, and clinical efficacy testing. Their conclusion is direct: large language models “remain largely untested in high-stakes mental health contexts,” and the field needs standardized evaluation and benchmarking aligned with medical-AI certification before deployment can be considered safe, transparent and ethical (Hua et al., 2025).

A companion meta-analysis addressed the possibility of harm and found that apps rarely made users worse off on average, with a pooled deterioration rate of 6.7% that was not significantly different from control conditions. The problem is that adverse events are seldom assessed: only 55 of the 171 identified clinical trials reported them at all (Linardon, Fuller-Tyszkiewicz, et al., 2024).

BRAKED’s answer to the gap highlighted particularly by Hua et al. (2025) is primarily architectural. It contains no generative model, no chatbot, and no conversational surface. Every word a user reads was written in advance by a named, credentialed psychologist (namely the author of this whitepaper and the developer of the instrument, in personal union) and fixed at publication; nothing is automatically generated in response to what the user types anywhere in the app.

There is another aspect of LLM-based or other chatbot use in the mental health context that goes beyond product certifications in terms of safety, transparency, and ethics. It is the aspect of general suitability of the tool for the designated job, regardless of how safety is defined and implemented. An LLM is (however advanced) an assisting tool that can accumulate, organize, and present information in varying degrees of accuracy, and can give different answers to the same question from one attempt to the next, following its underlying statistical next-token prediction methods and any guardrails built in. It cannot, and is not made to, understand human behavior, the individual and social backgrounds that shape and drive it, individual emotional nuances and biographic gradients of a personal history.

As such, the risks of entrusting mental health behavior and outcomes to an advanced auto-complete assistant, according to my professional judgment, outweigh the likely benefits. These risks likely become even more accentuated in our era of increasing loneliness in modern urban societies (Ghebreyesus et al., 2025; Office of the Surgeon General, 2023). Accordingly, BRAKED consciously does not employ any LLM-chatbots.

That is the situation this document is written into, with the hope of setting a positive example in professional rigor, research-based utility, transparency, quality and privacy in this space. What follows gives a source for each claim drawn from research, marks claims that rest on clinical convention as such and labels the decisions I made without evidence as decisions of practicability and utility.

3. What BRAKED Is

The interaction the whole application is built around takes about 15 seconds.

An urge arrives. The user opens BRAKED, taps the template matching the behavior, chooses a duration from a list, and starts the timer. Then they put the phone down and do something else until it finishes. When it does, they answer three short questions about how it went and, if they want, type a note. That is the loop. Everything else in the app exists to support it or to show the user what it has accumulated into.

The brevity is deliberate. Section 1 described the central practical failure of behavioral homework, which is that the recording happens hours after the moment or does not happen at all. An instrument that takes 15 seconds and requires no login has a chance of being reached for at 11 at night with a phone already in hand. One that requires locating a notebook does not.

The training profile

BRAKED ships with seven behavior templates: Doomscrolling and Social Media, Eating, Sweet Treats, Shopping, Trading, Gaming, and Anger Reaction. A user may also build a custom template by naming a personal impulse and setting their own duration. Any running timer can be cancelled before it completes; cancelled sessions are recorded as such in the personal statistics. The real elapsed hold time before cancellation is recorded alongside it, not merely the fact of cancellation.

A profile can have up to three templates active at any one time. The active templates determine which behaviors receive the weekly coaching material, so this distinction applies only to subscribers; the free, core instrument, has no coaching to filter and is unaffected by the templates selected as active. Any of the seven templates, as well as a custom one, can be used to start a timer whether or not it is active, and personal statistics include every session run, regardless of whether its template is active. Section 5 sets out the reasoning behind the cap.

The ladders

Each template offers a fixed set of durations. The lowest rung is five minutes in every predefined template. The ceilings differ by template, and the intervals between rungs are not uniform, because they are calibrated to the behavior rather than generated by a formula. Rung selection is not formal progression-based: Any timer rung can be chosen at any time as the user sees fit.

The ladders (continued)

Table 1:
Behavior templates and their timer rungs

Template	Available durations	Ceiling
Doomscrolling and Social Media	5m, 10m, 15m, 30m, 1h, 2h, 4h, 6h, 12h, 1d, 2d	2 days
Eating	5m, 10m, 15m, 30m, 1h, 2h	2 hours
Sweet Treats	5m, 10m, 15m, 30m, 1h, 2h, 4h, 6h, 1d, 3d, 7d	7 days
Shopping	5m, 10m, 15m, 30m, 1h, 2h, 4h, 6h, 1d, 3d, 7d	7 days
Trading	5m, 15m, 30m, 1h, 4h, 1d, 3d, 7d	7 days
Gaming	5m, 10m, 15m, 30m, 1h, 2h, 4h, 6h, 12h, 1d	1 day
Anger Reaction	5m, 10m, 15m, 30m, 1h, 2h, 4h, 6h, 1d	1 day
Custom	Free entry in minutes, hours, days or weeks	None

Several of these choices may look irregular at first, but each follows the logic of the behavior it is designed for. Trading follows the most widespread chart timeframes a retail daytrader already reads throughout the day, which is why there is no two-hour rung as a design choice. Sweet Treats and Shopping move from six hours to one day because these are momentary behaviors rather than continuous ones, unlike doomscrolling, gaming, or trading, which can continue for hours at a stretch. In practice, the hours between the six-hour and one-day rungs will be likely spent asleep, and a delay started in the afternoon should not end with a notification at three o'clock in the morning. Eating stops at two hours for reasons that are clinical rather than technical. Section 6 takes each of these choices in turn.

The post-session review

When a timer completes, the screen reads *Well done. You successfully delayed your impulse*, and then asks three questions:

- **Difficulty**, on a 10-point scale
- **Satisfaction**, on a 10-point scale
- **Did the impulse go away?**, answered yes or no

A free-text note is optional. The 10-point format belongs to the post-session review alone. Where the coaching material described in Section 7 asks a user to rate an urge, it uses a 0 to 10 scale, which is the convention that reads most naturally inside written guidance. The two ranges look similar, but they remain separate instruments: one is a structured, recorded item shown after every completed delay and feeding the app's own statistics; the other is a freeform prompt inside written text that BRAKED never records. Neither feeds into the other.

The order matters and is discussed in Section 4. The user is told they succeeded before they are asked whether the impulse went away, because completing the delay is the thing being trained and the fate of the impulse is a separate question.

What the app shows the user

The Home tab presents, in order: total sessions completed and session completion rate (whether the delay was held to the end or cancelled early, a separate measure from whether the impulse itself resolved); a recap of the last session giving date, time, template, duration chosen, difficulty, satisfaction and whether the impulse resolved.

A calendar in which any day can be tapped to reveal that day's statistics as described above, the breakdown by behavior, the individual session lengths and the session count; the current week's weekday distribution and time-of-day breakdown; and finally the same two distributions again, cumulative across the whole history.

A separate stats subsection gives average difficulty and average satisfaction, both on the 10-point scale and both template and timer agnostic, alongside a cumulative count of resolved sessions against total sessions; each of the three is shown for the current week and again for the lifetime total. Beneath that, Avg. Evaluations per Behavior (Lifetime) breaks the same three figures out by template, limited to templates with at least one completed session.



These are not decorative counters. A weekday distribution showing almost everything falling on a Friday, or a time-of-day breakdown showing half of all sessions occurring at night, gives the user meaningful information that memory could not reliably provide. Both the current week and lifetime totals are shown because a pattern only becomes meaningful when compared with a baseline.

This is where the Section 1 finding that recorded tracking outperforms remembered tracking becomes a design feature rather than simply a citation (Harkin et al., 2016). The app is not interpreting anything. It shows the user their own record in a form that makes the shape of it visible. An Export PDF (All Statistics) button, present on the Home tab, in the inbox's session history and overall stats views, and in the data section of Settings, compiles all of this into a single report: every timer-related figure the app currently holds, broken out by month, the current week and the lifetime total shown live in the app, drawn from whatever training history a purge has not removed, since a fuller history produces a fuller report.

The report closes with a Session History table listing every recorded session with no limit on count, sorted newest to oldest, with an on-page note explaining that order so that anyone printing the report can knowingly trim the older pages if they choose to. The report excludes inbox management state, meaning which coaching items were received, kept, or have expired, along with user settings and device information; those are handled separately by Export/Import User System Data, the function built for moving a full profile between devices, which remains local with no server or cloud component. The PDF is built for a human reader instead: the user, a family member or friend the user designates as their *attachment figure*, or a mental health professional. Who sees it, and whether it is shared at all, is the user's decision; it stays on the device until they choose otherwise.

A separate Raw Data Export (CSV) button, in the same data section of Settings, exports the complete underlying session records and onboarding history as a plain CSV file. This sits apart from both the PDF report and Export/Import User System Data, and is built for a different reader: someone who wants their training data in a spreadsheet and prefers to work it themselves rather than read the version BRAKED assembles.

A third, outbound-only button appears twice in the app, directly on the Home tab and in Settings/Data: a button linking to a dedicated feedback form hosted externally at the EPISTEMIST website. Submission requires no login, no IP-logging and no contact information, and is entirely voluntary; nothing about using BRAKED depends on it. The form itself retains nothing; entries are transmitted directly to EPISTEMIST's SSL-secured mail server rather than stored on the website. A code from a large, randomly assigned pool can be copied into the form from inside the app; it functions as a checksum-validated proof that a submission originated from within BRAKED, not as an identifier of who sent it. A missing or invalid code does not block submission. The entry is still received, only flagged as validation-missing for separate handling.



A selection of “Home Tab” elements

Privacy architecture

There is no account and no login. Nothing from the app is transmitted anywhere. All data is written to the device and stays there, and the user can export it, import it, or purge it entirely at any time. The weekly coaching content is bundled directly into the application, with new material authored by a psychologist at EPISTEMIST shipped regularly via standard store updates.

Profile resets and data purges are separate actions. Resetting the profile changes which behaviors are active and leaves the history intact. Purging erases the record, and it happens only when the user deliberately chooses it. Applications in this category commonly collapse the two, so that changing what you are working on costs you the record of what you have already done. Keeping them apart means a user can restructure their training without paying for it in lost history.

Both actions are unlimited. Both are free.

In the app, the profile reset is called Reset Onboarding. It only provides an opportunity to redesign the training focus and the currently perceived baseline intensity of impulses if so desired. It does not affect a subscription's entitlement, the training history, or coaching material already delivered, archived, or favorited in the inbox. For a subscriber, its main use is redirecting the coming week's coaching material toward a different set of behaviors. The Onboarding history can be seen on the respective submenu item of the inbox tab and is exported into the PDF report section alongside other user statistics.

What costs money

The delay timers for all seven templates and the custom template, the statistics, session history, onboarding history, and calendar, unlimited profile resets, and unlimited export, import, report generation and training data purge are always free and remain so.

A subscription adds weekly research-based behavioral coaching material, written by a psychologist at EPISTEMIST and delivered to an in-app inbox, keyed to whichever behaviors are currently in the user's profile. Section 7 gives four complete examples, written to the same specification, so that the material can be assessed rather than taken on description. The annual plan offers one month off and adds one personally authored "Deep Dive" psychological consultation per subscription year, up to 1,300 words on a topic of the user's choosing and available under a pseudonym, together with an ongoing discount on separately bookable priority-tier 1:1 live consultation sessions through EPISTEMIST for the entire duration of the annual plan.



Coaching material delivered to the inbox expires 28 days after delivery and is removed automatically if not archived or favorited, regardless of whether it was read. The 28-day countdown and the option to keep the item are shown under its title for as long as it remains in the inbox.

Manually archived or favorited coaching articles persist forever and are included in Export/Import User System Data to be seamlessly available on a new device.

A market datum worth setting beside this. A 2025 review of AI-branded mental health apps, a child-focused sample, but the pricing pattern is not specific to that age group, found that only seven of the 27 apps evaluated were free; the remaining 20 required a subscription or one-time payment, averaging \$20.16 a month (Yang et al., 2025).

In sum, the core training instrument is entirely free to use and requires no account creation. Subscribers pay for the authored weekly coaching content only. The annual subscribers get a distinct structural advantage besides the weekly coaching content for a full year with one month off: This plan includes an “Annual Bonus” on top; a personalized Deep Dive consultation as described above, whose standalone catalog value at EPISTEMIST independently exceeds the cost of the entire BRAKED annual subscription.

The division is between the mechanism and the authored material. Nothing that makes the app work as a training instrument is behind the subscription, and nothing about it expires.



Accordion menu elements and a selection of unread weekly coaching content on the “Inbox Tab”

4. The Training Model: Delay as a Trainable Capacity

A delay timer is a simple object. The reasoning for why it should work is less simple, and it needs setting out in full, because several plausible accounts of why it works are wrong, and at least one of them, telling someone to stop wanting something, is wrong in a way that does active damage rather than simply falling short. The next few pages work through each account in turn, starting with that one.

Two words do different work in what follows, and it makes sense to fix them here. My working definition of an urge is the wanting itself, a desire at whatever intensity, to do the thing. An impulse, on the other hand, is what an urge becomes once it turns into pressure to act on it immediately, before the moment passes. Most urges may never make that turn. This is why the post-session question asks whether the impulse went away rather than the urge, and why the desire literature below is read as a description of the raw material delay training works on, not yet the pressure to act that a delay interrupts.

The distinction has a home in the dual-process literature. Hofmann et al. (2009) describe self-control as a contest between an impulsive system, which reacts to immediate cues, and a reflective system, which represents longer-term goals; Strack and Deutsch (2004) set out the underlying model. An impulse, in that framework, is what wins when the impulsive system outpaces the reflective one before it can weigh in.

What the model has to explain

Hofmann, Baumeister, et al. (2012) gave 205 adults beepers for a week and collected 7,827 reports of wanting something. Three findings from that dataset shape everything below.

Desires are frequent and mostly ordinary. Most instances of wanting are not a self-control problem and do not need to be treated as one. Where a desire conflicts with a person's other goals, they generally do try to resist it. And resistance succeeds unevenly, which is where the useful distinction lies, because the gap between trying and succeeding is the only place an intervention has anything to work with. The companion analysis of the same data is instructive on its own (Hofmann, Vohs, et al., 2012). Desires for tobacco and alcohol had the lowest average strength of any category measured, despite being the substances everyone assumes produce the most overpowering wanting.

Intensity and addictiveness appear to be separate variables. Meanwhile, media use had the highest rate of self-control failure of any domain in the study. The desire involved was not the strongest measured, yet people were most likely to fail to resist it.

That finding is the reason Doomscrolling and Social Media is the template most users are likely to start with.

The wrong answer that sounds right

The intuitive response to an unwanted urge is to try to stop wanting. This is the worst available instruction and the evidence against it is old and solid.

Wegner (1994) described what happens when someone deliberately suppresses a thought. Two processes engage: an effortful one searching for distraction, and an automatic one monitoring for the very thought being suppressed. The monitor is cheap to run and the operator is expensive, so under any cognitive load the monitor wins and the suppressed content becomes more accessible rather than less. Wang et al. (2020) pooled 31 studies using the original paradigm and found the rebound effect robust regardless of load.

Applied to urges, this rules out a whole family of designs. Telling a person not to want something is a suppression instruction and it will produce rebound. Telling them to want it and wait is not a suppression instruction at all, because nothing about the desire is being contested. The intervention sits between the urge and the action rather than between the person and their own wanting.

BRAKED therefore never asks a user to stop wanting anything, never frames the urge as an adversary, and does not block or restrict access to the behavior. The coaching material is written under the same constraint.

What the urge itself does, and what nobody has measured

Urges are self-limiting. In the absence of reinforcement they rise and fall rather than escalating indefinitely, which is the expected behavior of a conditioned appetitive response and the reason waiting works at all.

What has not been established is how long that takes. Any figures in circulation, that cravings peak at 15 to 20 minutes, or pass in 3, can be clinical rules of thumb. I am not aware of a study that has tracked individual urge trajectories at a resolution fine enough to fix a duration; the sampling intervals in this literature are spaced in hours, and an instrument that samples every few hours cannot describe the shape of something that lasts minutes.

This matters beyond the academic. Telling someone their urge will be gone in 3 minutes, and then watching it run past 15, is actively counterproductive. When the promised number fails, the person has been handed evidence that their craving is abnormal or that they are uniquely bad at this. The defensible statement is that it will pass, that the timing varies, and that the skill is getting through it without feeding it. That statement survives a 20-minute urge. A stopwatch promise does not.

So the wave explains why a delay works. It does not explain why any particular duration was chosen. That is a separate question with a separate answer.

The ladder

BRAKED's durations are a graded load and do not describe urge durations in any way. Five minutes is just the lightest weight on the rack, rather than a claim about how long the particular urge lasts.

Three lines of evidence converge on graded practice as the right structure.

The first is shaping. Complex behavior is established by reinforcing progressively closer approximations to the target, which is foundational operant methodology (Skinner, 1951, 1953) and standard practice in behavioral activation, where a distant goal is broken into a staircase of steps small enough to be completed.

The second is graded exposure, though the analogy needs a qualifier. Graded exposure hierarchies, a structure with a long history in behavioral practice dating to Wolpe (1958), share the surface form of a delay ladder (ascending steps, one attempted once the last is manageable) but the mechanism they target is different: extinguishing a threat appraisal, not building tolerance for wanting. Current exposure practice has in fact moved away from strict graduation for exactly this reason: Abramowitz et al. (2019) argue that requiring anxiety to subside on an easier step before a harder one reinforces the belief that lower-intensity distress is safer, and favor confronting items in whatever order is most productive. BRAKED's ladder is an inference drawn from the older, structural resemblance to graded exposure and from shaping (Skinner) and mastery experience (Bandura), rather than a direct application of current exposure methodology.

The third and most important is Bandura (1977) on where a sense of one's own capability actually comes from. Of the four sources Bandura identified, direct successful performance is the most potent, ahead of watching someone else succeed, being persuaded by others, and reading one's own physiological and emotional state; what builds the sense of capability is doing the thing and finding that one did it. Bandura also specified the conditions under which a success builds anything: it has to be non-trivial, it has to be attributable to the person's own effort, and the difficulty has to increase over time. A success that cost nothing teaches nothing.

A recent review surveyed the animal-learning and human literatures for interventions shown to build self-control directly, effort exposure, reward discrimination, reward bundling, interval reinforcement schedules, impulse-control training, and mindfulness training among them.

Most of these increased self-control on measures taken after training; mindfulness training was the exception (Smith et al., 2019). The finding most relevant here is the general one: self-control, across most of the paradigms tested, behaves like something that can be built through structured practice. Most of this evidence is domain-specific and comes from non-human subjects, so it is read here as convergent support for the premise only.

A more targeted line of work trains delay tolerance directly through repeated experience with delayed reinforcement. In a study of 36 rats, researchers randomly assigned animals to either delay-exposure or immediacy-exposure training before assessing impulsive choice. The rats then chose between one immediately available pellet and three delayed pellets across two assessments, with delays increasing from 0 to 32 seconds within the first assessment and from 8 to 4 seconds between sessions in the second. Delay-exposure training reduced impulsive choice in female rats during the initial five sessions, although the effect did not persist beyond that period. This finding extends a pattern previously observed seven times in male rats while also showing that the effect may differ by sex and may not be durable in female rats (Peck et al., 2022).

BRAKED's ladder applies those principles by construction. Five minutes is not trivial for someone who has never deliberately delayed an urge in their life. The completion is unambiguously the user's own, because the app does nothing except hold time and refuse to make the decision for them. And the ladder ascends according to the user's choice.

Self-efficacy is not merely correlated with behavior change. Sheeran et al. (2016) found that experimentally raising self-efficacy produced change in behavior and not only in intention. It functions as a lever rather than as a symptom.

The working term I use for this internally is a **mental plank**: an isometric hold, kept for as long as it can currently be kept, then kept longer next time. Nobody expects to hold a plank indefinitely on the first attempt, and nobody concludes from a one-minute plank that they are constitutionally weak.

The specific rung values are a design judgment. I know of no study establishing 5, 10, 15 and 30 as optimal, and I make no claim that they are. The principle of graded progression is well supported. The particular staircase is mine.

The plan that has to exist before the urge does


The second mechanism is separable from the first.

Gollwitzer (1999) drew a line between naming a desired end state and specifying in advance the situation that will trigger a particular response: if Y happens, I will do X. Pooling 94 independent tests of the second kind, Gollwitzer and Sheeran (2006) put the effect on reaching the goal at $d = .65$.

A BRAKED template can serve as the then-part of such a plan. A user who decides at a calm moment, during onboarding for example, that the next shopping urge will be met with a Shopping timer has formed one: if I feel the pull to buy, then I start a delay. At the moment of the urge, the app offers the response already chosen, so no deliberation is needed. The meta-analytic effect belongs to the planning. The app supports carrying the plan out and has not been tested as a planning tool itself.

The three-template cap, explained in Section 5, governs coaching bandwidth only. Any template can serve as a planned response, whether or not it is one of the three currently active.

There is a further reason to build on a decision made in advance rather than on resistance in the moment. People who score well on trait self-control report struggling less, not more, and their advantage over everyone else runs through the routines they have established rather than through the force they apply when tempted (Galla & Duckworth, 2015). Experience sampling points the same way. Milyavskaya and Inzlicht (2017) followed students through a semester and found that how much progress they made on their goals tracked how often temptation arose in the first place, while the effort they put into resisting it was unrelated to whether they got there. This is a caution for any tool built on in-the-moment resistance, BRAKED's timer included. It is the main reason the coaching material gives as much attention to arranging the environment, so that temptation arises less often, as to the moment of the urge itself. Whether repeated delays also reduce how often temptation arises is untested and belongs with the open questions in Section 12.



A daily summary of impulse delay statistics across multiple domains, selectable from the Home Tab calendar

What counts as success

This is the point on which BRAKED differs most sharply from the category, and it rests on a single study that most of the category reports incorrectly.

Bowen and Marlatt (2009) randomized 123 undergraduate smokers in a cue exposure paradigm, giving one group brief urge surfing instructions and the other nothing. The two groups reported wanting a cigarette to much the same degree. Over the week that followed, the group that had been taught the technique smoked significantly less. What the authors take from this is that the instruction left the wanting where it was and changed what people did while it was happening.

The null result is the finding. I observe particularly in commercial treatment marketing, a tendency to convert it into a large craving reduction with an invented percentage attached, which reverses its meaning.

BRAKED is built on the version that was actually published. The proposition is that what happens between the urge and the action will change, rather than that urges will weaken. This is why the completion screen tells the user they succeeded before it asks whether the impulse went away, and why answering *no* to that question does not make it a failed session. The delay was held. That was the exercise. When a timer ends, the decision returns to the user. Acting on the urge after a completed delay is a legitimate outcome. The training lies in the wait itself and, over time, in occasions of acting that become fewer and further apart.

The objection a psychologist will raise

Any document that invokes delayed gratification has to address the marshmallow test; leaving it out would be an obvious omission.

Watts et al. (2018) conducted a conceptual replication using a larger and more diverse sample. The bivariate association between preschool waiting and adolescent achievement was roughly half the size reported in the original study, and controlling for family background, early cognitive ability, and home environment reduced it by about two thirds. Most of the predictive power came from the first 20 seconds of waiting: what separated children was whether they could wait at all, not how long the longest waiters lasted. A later follow-up found that the test did not reliably predict adult functioning (Sperber et al., 2024). The debate is live rather than settled, and critics have argued that some of the covariates plausibly capture constituent parts of delay ability rather than confounds, which would attenuate the association mechanically (Doebel et al., 2020). None of this undermines what BRAKED claims, because BRAKED does not make the claim under attack.

The finding that has weakened is that childhood delay capacity is a stable trait predicting distant life outcomes. BRAKED asserts nothing about traits and nothing about life outcomes. It asserts that a specific situational skill can be practiced by an adult who wants to practice it, which is a different construct resting on the evidence set out above.

Watts et al. (2018) make a point in their own discussion that supports the distinction. An intervention that changes delay ability without changing broader cognitive and behavioral capacities should be expected to have small effects on distal outcomes. That is an argument against treating delay as a predictive trait, yet it does not imply anything against delay as a trainable capacity.

5. From Construct to Feature

Section 3 described what BRAKED does. Section 4 set out why delaying an urge should work at all. This section connects the two, feature by feature, and states how strong each connection is.

The strength matters as much as the connection. Not every decision in a product like this rests on the same category of evidence, and presenting them as though they did is the exact failure Section 2 documented. Some features implement findings with meta-analytic support. Some implement conventions that are standard in clinical practice but have never been measured directly. Some are decisions derived from practical behavioral reasoning, grounded in pragmatic application mechanics rather than empirical trials.

I use four labels throughout the rest of this document.

T1, established. Meta-analytic support or multiple independent trials, replicated, with effect sizes reported. May be stated as a finding.

T2, convergent. Consistent results across several studies, but without meta-analysis, or with meta-analysis showing heterogeneity or non-significant follow-up. Stated as evidence suggesting, not as evidence showing. Where a single principal source is given for a T2 row, it anchors a wider convergent literature rather than standing as the only study behind the claim.

T3, clinical heuristic. Widely used in cognitive-behavioral practice and theoretically grounded, but not directly measured. May be described as something clinicians do. Never as something studies show.

T4, design judgment. A decision made on my own practitioner judgment and operational constraints, where the literature does not specify the parameter in question. Stated explicitly as a decision.

The derivation

Table 2:
BRAKED Features Mapped to Their Psychological Constructs and Evidence Tiers

Feature	Construct it operationalizes	Tier	Principal source	What may be claimed
The delay timer	Delay of enactment; practice in tolerating a wait	T3	Peck et al. (2022); Bowen & Marlatt (2009)	A structured practice in holding a delay. The timer on its own has not been tested; Bowen & Marlatt tested the same wait paired with urge-surfing instruction.
Selecting a template	Implementation intention: a cue class bound in advance to a response	T1 construct, T4 as implemented	Gollwitzer & Sheeran (2006)	Meta-analytic effect size of the underlying construct ($d = 0.65$), serving as theoretical rationale rather than a trial of this specific interface.
The ladder	Graded practice: shaping, exposure hierarchy, mastery experience	T2 principle, T4 values	Wolpe (1958); Bandura (1977); Sheeran et al. (2016)	Empirical support for progressive exposure principles, rather than optimization of these specific step intervals.
Difficulty rating	Perceived effort; the progression signal for the hierarchy	T3	Analogue of distress ratings (SUDS) in graded exposure (Abramowitz et al., 2019)	Utility as an intra-individual progression signal, strictly not to be interpreted as a validated clinical scale.
Satisfaction rating	Mastery experience	T2	Bandura (1977)	Theoretical alignment with mastery as a driver of self-efficacy, serving as a subjective reflection prompt rather than a psychometric measurement.
“Did the impulse go away?”	Impulse outcome, recorded and never scored	T2	Bowen & Marlatt (2009)	Descriptive outcome logging. Bowen & Marlatt found no group difference in urges; this item records the closest in-app counterpart, the impulse.

The derivation (continued)

Table 2 (continued)

Feature	Construct it operationalizes	Tier	Principal source	What may be claimed
Session count	Frequency of deliberate practice	T3, by extension	Apolzan et al. (2017)	A process measure of practice frequency. Apolzan links lower consumption frequency to lower craving; the link to this count is inference (Section 12).
Calendar, weekday and time-of-day distributions	Self-monitoring; recorded rather than remembered information	T1	Harkin et al. (2016); Michie et al. (2009); Stone et al. (2002)	Support for self-monitoring when recorded or visible (Harkin, Michie), and for the real-time-vs-reconstructed gap (Stone).
Maximum of three templates	Weekly coaching bandwidth	T4	None	A deliberate constraint on how many behaviors receive weekly coaching content, not on how many can be trained; the training instrument itself carries no such limit.
Weekly content cadence	Repetition requirement of mastery practice	T4	None directly	A deliberate structural choice supporting the repetition required for mastery practice.
Seven separate templates	Cue reactivity and extinction specificity	T2	Carter & Tiffany (1999); Price et al. (2010)	Rationale for domain-specific context matching, cited as inference rather than direct application: both sources concern laboratory extinction only.
On-device storage, no account	Not a psychological construct	T4	Not applicable	An explicit architectural choice prioritizing privacy, with trade-offs detailed in Section 8.

Four of these need more than a table row.

Frequency rather than magnitude

Most self-control applications measure the size of what was avoided. Minutes not spent, money not spent, calories not consumed. BRAKED counts completed delays instead, and the reason comes from a study that had nothing to do with apps.

Apolzan et al. (2017) followed 367 adults with overweight or obesity through a two-year weight loss trial, measuring food cravings with the Food Craving Inventory at baseline, 6 months and 24 months, alongside actual food intake. They asked two questions. Did changes in *how often* people ate a craved food correspond to changes in how much they craved it? And did changes in *how much* they ate correspond to changes in craving?

The first was associated at 6 months and again at 24. The second was not associated at all. Reducing the amount consumed did nothing to cravings; reducing the frequency of consumption did.

This supports a conditioning account, in which each pairing of cue and consumption maintains the craving, so fewer pairings mean less craving over time. What matters is how often the pairing happens, far more than how large each occasion is.

Two qualifications travel with the finding. The effect size was small, and the authors report it as such. The data are correlational, and the authors say so explicitly. Hence, this is a T2 finding.

Apolzan et al. (2017) studied food consumption frequency rather than digital delay timers. Extending their finding to BRAKED's completed-delay count is an inference on my part, consistent with the conditioning model the authors themselves invoke.

What it justifies is the emphasis on frequency. A completed delay pushes an occasion of acting later or removes it, so a user who practices regularly will usually be acting on the urge less often. The session count records that practice; whether it tracks reduced craving is untested (Section 12).

Three templates and no more

The cap on active coaching is a T4 decision. I am not aware of a study establishing three necessarily as an optimum and I do not imply one.

The cap applies to which behaviors receive the weekly coaching material, not to which templates a user can train. All seven templates and the custom template accept timers at any time, independent from whether they are part of the active three, and personal statistics reflect every session run regardless. What the cap limits is coaching bandwidth.

A person can realistically take on new material for perhaps three behavioral strategies in a week; material for seven arrives faster than anyone could act on it, and material nobody acts on is noise rather than coaching.

Why the coaching material arrives weekly

Also T4, and the reasoning is short.

Mastery requires repetition rather than exposure. A technique read once and never attempted has produced no performance, and Section 4 set out why performance is the thing that matters. A week gives a reader room to meet the situation the material describes several times, try the technique, fail at it once, and try again before the next piece arrives.

A daily cadence would deliver more material and less practice. It would also compete for attention with the notifications the app exists to help people resist, which would be an odd thing for it to do.

The post-session review is not a psychometric instrument

The three questions after each completed delay map onto recognizable constructs. Difficulty functions as the progression signal that governs when to attempt a longer rung, in the same way distress ratings govern movement through an exposure hierarchy. Satisfaction is a mastery rating, of the kind used in behavioral activation. The impulse question records the in-app counterpart of the outcome that returned null in Bowen and Marlatt (2009), who measured urges; that is why it is recorded and never scored.

None of these are validated scales; they are bespoke items on a bespoke 10-point format, written for this application. The record they produce is personally meaningful, and deliberately so: a user reads their own trend against their own accumulated baseline, which is a measurement in the within-person sense and in that sense only. Nothing derived from it should be treated as an objective measurement.

I state this here rather than in the limitations section, as it belongs next to the description of what the items are for.

6. Seven Domains, Seven Trained Skills

The seven templates are not one timer with seven labels. Each has its own ladder, its own ceiling, and in some cases a design decision that only makes sense for that behavior. This section takes them in turn.

Four of them also carry constraints that follow from the behavior itself. Two, Eating and Sweet Treats, limit what the coaching material may say. Three, Eating, Trading and Anger Reaction, limit what the template is meant to address. Each constraint appears with its template here, and Section 10 collects all of them in one place.

Each predefined template exists because it clears a specific bar: a ladder ships only once the coaching material behind it is written and the evidence supporting that material has been assembled. Tobacco, alcohol and other legal but harmful substances are not templates for this reason, not because they sit outside the model. Building either properly, the literature review and the coaching library, is its own undertaking, and it is on the list rather than done.

Doomscrolling and Social Media

Ladder: 5m to 2d. **Construct:** cue-driven checking maintained by intermittent reward.

This is the template most users likely select first, and the evidence supports that choice better than intuition would. In the experience sampling study described in Section 4, media use produced the highest rate of self-control failure of any domain measured, and it did so without producing the strongest desire (Hofmann, Vohs, et al., 2012).

The mechanism is old and well characterized. A reward arriving after an unpredictable number of responses holds behavior more strongly, and outlasts the withdrawal of reward for far longer, than one arriving every time. Skinner (1953) made the point with the gambling house, which keeps a patron handing money over by returning some of it on exactly that schedule; the laboratory basis for the claim is Ferster and Skinner (1957). A feed that is sometimes interesting and sometimes not runs on the same schedule, and the checking behavior it produces is the expected result rather than a personal failing. The coaching material says this explicitly, because a user who understands they are responding normally to a designed system is in a better position than one who concludes they are weak.

What is not claimed. I am aware of no empirically supported withdrawal timeline. Circulating accounts of what happens across a social media break, in which cravings subside at one week and disappear at six, are not ones I have been able to trace to the peer-reviewed literature, and thus BRAKED does not publish a timeline either. The honest position is that the timing varies between people and is not currently established.

Eating

Ladder: 5m to 2h. **Construct:** urges to eat that arise outside of hunger; often affect-driven.

The template addresses the pull to eat when the body is not asking for food, whether to regulate a feeling or in response to a cue. Hunger itself is a separate problem, and the template does not address it.

This distinction has direct support in the clinical literature: a narrative review of 115 articles on food craving in medical weight management concludes that craving is a hedonic phenomenon distinct from homeostatic hunger, and that managing hunger alone is insufficient without also addressing craving (Kakoschke et al., 2024).

The ceiling is the design decision that needs explaining, and it is a subtraction. The architecture is capable of longer durations and every other template uses them. Eating stops at two hours because past a certain point a delay before eating stops being a training exercise and becomes restriction, and an application that structures restriction for a user who may already be inclined toward it has done harm regardless of intent.

Two hours is not a measured threshold but my professional practitioner's judgment. BRAKED is a domain-focused behavioral microtraining application, distinct from intermittent fasting and from calorie counting; it helps the user delay the urge and then reflect on the experience, building personal mechanisms toward the urge management they are after. What is not defensible is offering 12 hours or longer here because the software could.

Binding constraint on content. No coaching material for this template contains calorie targets, portion prescriptions, meal timing rules, weight-loss framing, or any numeric dietary guidance whatsoever. This is not a matter of tone. It is a rule, and it is enforced at the content preparation stage. Scope constraint. Eating difficulties severe enough to require specialized treatment fall outside what a delay timer can address. This template addresses everyday affect-driven eating, which is a narrower problem.

Sweet Treats

Ladder: 5m to 7d. **Construct:** hedonic craving for a discretionary food.

A reader who has just read the Eating section will notice that this template allows a seven-day delay for something also eaten, and the difference is deliberate.

The Eating template works close to food a person physiologically needs, since a long delay before eating in general can end up postponing a meal. Sweet Treats concerns a specific discretionary food nobody needs. Nothing compels a delay before a piece of chocolate to end, and no physiological signal is being overridden by extending it. The two templates are calibrated to which construct is in play rather than to a category called food.

The long ceiling also implements the finding described in Section 5 more directly than any other template. Apolzan et al. (2017) found that reducing the *frequency* of consuming a craved food was associated with reduced craving for it while reducing the *amount* was not. The mechanism is fewer occasions, each one intact. This is delay followed by permission rather than delay followed by denial, and the coaching material for this template is written to that framing without exception. A seven-day wait ending in the thing you waited for is the intended shape. A seven-day wait framed as abstinence would be a different intervention with a different risk profile.

The same prohibition on numeric dietary guidance applies here.

Shopping

Ladder: 5m to 7d. **Construct:** delay discounting; arousal at the point of purchase.

Impulsive buying is the clearest everyday expression of steep delay discounting (also called temporal or time discounting in parts of the literature; Scholten et al., 2019), in which an immediate reward is valued far above a larger one further away. The intervention is straightforward in principle: put time between the wanting and the transaction, and let the arousal that inflated the immediate option subside.

One finding shapes the coaching material for this template specifically. Episodic future thinking (Atance & O'Neill, 2001), in which a person vividly imagines a concrete personal future scene before an intertemporal choice, reduces discounting in laboratory choice tasks. In a meta-analysis of 47 studies, the pooled effect was $g = 0.52$, and positively valenced imagery produced the largest effect, at $g = 0.64$, while neutral or negatively valenced imagery combined produced none, at $g = -0.03$ (Ye et al., 2022). Two limits apply: the studies used hypothetical choices with delays of months or years rather than hours or days,

and the valence finding is not fully settled, since at least one large replication found negative imagery effective too (Bulley et al., 2019).

A systematic review of behavioral manipulations of discounting covers much of the same underlying literature (Scholten et al., 2019). The coaching material for this template asks the user to picture something specific and positive they are working toward. I leave out the common advice to picture the regret or the credit card statement: that kind of consequence imagery is a different construct from what these studies tested, and I prefer not to build coaching on fear regardless.

Trading

Ladder: 5m to 7d, following standard chart timeframes. **Construct:** loss chasing and impulsive execution.

The ladder here departs from every other template and does so on purpose. The rungs follow the timeframes a daytrader already reads all day: 5, 15 and 30 minutes, then 1 hour, 4 hours, 1 day, 3 days and 1 week. There is no two-hour rung because the two-hour chart is a far less standard timeframe than the ones around it. A duration expressed in a vocabulary the user already thinks in requires no translation at the moment it is needed.

The nearest evidence base is gambling disorder, and it deserves to be reported completely rather than selectively. An umbrella review synthesizing moderate- to high-quality meta-analyses (31 studies, $N = 4,023$) found that cognitive-behavioral treatment reduced gambling disorder severity at $g = -0.91$, frequency at $g = -0.52$, and intensity at $g = -0.32$ against minimal or no treatment at post-treatment (Pfund, Ginley, et al., 2023). Those are substantial effects.

They are also, on the same authors' primary meta-analysis of 29 studies and 3,991 participants, accompanied by three things that must travel with them. Nothing survived to follow-up at a level that reached significance. Publication bias showed up in the analysis and heterogeneity between studies was high. And the authors themselves read the post-treatment figure as flattering, and doubt that the techniques can be counted on to work for everyone who presents for treatment (Pfund, Forman, et al., 2023).

A whitepaper quoting only the first paragraph would be doing what Section 2 described.

Scope constraint. Trading at an intensity that meets criteria for gambling disorder requires clinical care, and a delay timer cannot provide that. This template addresses impulsive execution, which is a narrower problem.

Gaming

Ladder: 5m to 1d. **Construct:** tolerance for interruption.

Most tools in this space treat gaming as a quantity problem and respond with time limits. The evidence, however, points somewhere else.

Interrupting a game elicits markedly higher anger and anxiety than interrupting a control task, and the magnitude of that anger during interruption predicts individual differences in problematic gaming, mediated by urges to return to the game (Bao et al., 2025). What distinguishes problematic from unproblematic gaming appears to be less about hours accumulated and more about what happens when the session is interrupted and whether obligations are displaced.

A separate line of evidence points the same way. Across two independent samples totaling more than 500 participants, hours spent playing showed no meaningful relationship to procrastination or delay discounting on its own; what predicted difficulty was the reason someone gave for playing, escaping or de-stressing carried more risk than playing for entertainment or connection (Nordby et al., 2019).

That points the training at interruption tolerance rather than at total time, which is what a delay timer trains and what a time limit does not. A hard limit removes the interruption from the user's control and therefore removes the thing being practiced, which is the same objection Section 1 raised against blockers generally.

Anger Reaction

Ladder: 5m to 1d. **Construct:** arousal reduction before re-engagement.

This template draws on one of the largest evidence bases in the document.

Kjærvik and Bushman (2024) registered their predictions in advance and then pooled 154 studies, 184 independent tests and 10,189 participants, sorting every activity by whether it raised or lowered physiological arousal. Bringing arousal down reduced anger and aggression at $g = -0.63$, in a confidence interval running from -0.82 to -0.43 . The result held when the authors tested it against publication bias and against outliers, and it did not vary meaningfully by gender, age, race, culture, criminal-offender status, intellectual disability, or whether the study ran in a laboratory or in the field. One qualification applies directly to this template: in the studies where participants had just been provoked, the effect was smaller and its confidence interval reached zero, although the authors read the result as holding under provocation. Taking a timeout appears on their list of activities that work.

Arousal-increasing activities showed little to no benefit and in some cases made things worse. This includes jogging, cycling, hitting a bag, and venting.

The coaching material for this template therefore tells users what not to do during the delay as specifically as it tells them what to do, which is counterintuitive advice and correct.

Recovery from a meaningful provocation is slower than people expect and varies considerably between individuals, so a bounded pause before re-engaging is reasonable. The specific durations are design judgments. The one-day ceiling is deliberate: a pause in an interpersonal conflict is supposed to end, and a multi-day delay is avoidance rather than de-escalation.

Scope constraint. This template addresses everyday conflict rather than intimate partner violence, which requires specialized professional support and falls outside what a delay timer can address.

7. What the Coaching Material Actually Looks Like

A subscription adds one thing: in-house written material, authored by a psychologist at EPISTEMIST, updated and expanded regularly with ongoing literature review, delivered weekly to an inbox inside the application, keyed to whichever behaviors are currently in the user's onboarding profile. Someone training Doomscrolling and Anger Reaction receives items for those two and nothing else.

Each template holds an annual set of 52 items, each with domain-specific behavior regulation strategies, which is where the weekly cadence comes from. Section 5 gives the reasoning for that interval rather than a faster one. New sets are added year by year while the existing sets are regularly updated, so a subscriber in a second year draws on a larger library rather than repeating the first.

The material is written to a fixed internal shape, based on behavioral change mechanics supported by evidence. An item opens with a short framing of what is happening in the body, the attention or the environment. It then instructs a specific delay on a rung that exists in that template's ladder. It names something to do during the delay. It closes by asking the user to rate the urge again and choose a next step, which may be acting, making a request, or starting a longer timer. Each item is written in accessible language for a general audience, 100–170 words, aiming for minimal cognitive friction and quick practicability. It frames the phenomenon through one concrete everyday example, chosen so the underlying pattern generalizes easily to other situations.

That last element is what connects the material to the instrument. The coaching does not sit alongside the timers as general advice. It sends the reader into a timed session and meets them on the other side of it.

The four items below were written for this document rather than taken from the subscription library, which stays with subscribers. They are written to the same specification, by the same hand, and at the same length, and they are representative of what arrives in the inbox.

Three properties of the writing are visible in them and easy to mistake for stylistic preference. The material never asks the reader to stop wanting anything. It never promises that an urge will pass within a stated interval. It never treats a lapse as evidence about the person. Section 10 sets out the full set of constraints and where each comes from; what matters here is that they are constraints on the writing itself rather than a disclaimer attached to it.

The material is written in English and German, complete across all seven templates. The app detects the device's locale on launch and whenever it returns to the foreground, and displays in German automatically when that locale is German; English is the default fallback for any other device language, and a manual language switch is available in Settings. The Annual Bonus, a one-time, free, personalized, "Deep Dive" consultation, is available externally via EPISTEMIST, currently in English, German and Turkish.

Doomscrolling and Social Media, 10 minutes

The Empty Checks Are the Point

Think about the last ten times you opened the feed. Most of them gave you nothing much. That is not a sign the habit is weakening. Rewards that arrive unpredictably hold attention far better than rewards that arrive every time, which is why a slot machine holds people and a vending machine does not. The empty checks are doing the work. Rate the pull from 0 to 10, then start a 10 minute delay timer in BRAKED and leave the phone where it is. During the timer, do one thing with a guaranteed payoff: refill your water, step outside, or finish one small task you have been circling all day. When the timer ends, rate the pull from 0 to 10 and decide whether you actually want to open it.

The point of that item is the reframe, and it is doing more work than the instruction. Checking behavior maintained by an unpredictable reward schedule is the expected result of the schedule rather than a defect in the person (Skinner, 1953). A user who understands that is in a better position than one who concludes they lack discipline, and Section 1 gives the reasoning for keeping shame out of the design.

Shopping, 1 day

Picture What You Are Saving For

The standard advice is to picture the regret: the statement arriving, the thing sitting unused in a cupboard. Research on money decisions points somewhere more useful: imagining something specific and good that you are working toward makes the later reward weigh more. Rate the urge from 0 to 10, then start a 1 day delay timer in BRAKED and leave the item in the cart. Sometime during the day, take one minute and picture what this money is for, in real detail: where you are, who is there, what it looks like when it happens. Keep it concrete and keep it yours. When the timer ends, rate the urge from 0 to 10 and decide with that picture still in mind.

The substantive content of that item is the positive picture. Episodic future thinking (Atance & O'Neill, 2001) reduced delay discounting most when the imagined scene was positive, at $g = 0.64$ (Ye et al., 2022). Leaving out the regret picture is my design choice, for the reasons given in Section 6.

Sweet Treats, 3 days

Three Days, Then the Thing You Want

Eating less of something at each sitting seems to do little to change how much you want it. Eating it less often is the part that seems to matter. So this is a frequency exercise, and it ends with the thing you were waiting for. Rate the urge from 0 to 10, then start a 3 day delay timer in BRAKED and decide now what you will have when it finishes. Write it down. During the three days, eat the way you normally eat; nothing else changes. When the urge arrives, picture the person who gets to have it: you, three days from now, sitting down without hurry and actually tasting it. When the timer ends, rate the urge again from 0 to 10, then have what you chose, or start a longer timer if you would rather keep waiting. Sit down for it, take your time, and pay attention to it. The waiting is the training, and the treat is how it finishes.

The last three sentences are the template's whole design position. Reducing the frequency of consuming a craved food was associated with reduced craving for it, while reducing the amount was not (Apolzan et al., 2017). A delay that ends in denial would be a different intervention with a different risk profile, and Section 6 sets out why this template is written to permission instead. The middle instruction draws on a second, independent mechanism. People who perceive their future self as similar to, and connected with, their present self make more choices that benefit that future self, an effect demonstrated by pairing self-continuity ratings with saving behavior (Ersner-Hersfield et al., 2009). Framing the wait as something done for a future version of oneself, rather than to a present one, recruits that connection directly, and it fits this template specifically because the future self is not being asked to sacrifice here. It is the one who gets the reward.

Anger Reaction, 15 minutes

What Not to Do With the Energy

Anger fills you with energy and the energy wants out. The common advice is to burn it off with a run, a punching bag, or a long vent to a friend. The general evidence points the other way. Start a 15 minute delay timer in BRAKED. During the timer, do the slow things. Sit down, let your out-breath run longer than your in-breath, loosen your jaw and your hands, and keep your eyes on one fixed point. Do not replay the argument or draft the message, and do not tell the story to anyone yet. If your mind pulls toward the details anyway, zoom out instead of in: ask why this actually matters rather than what exactly was said, and picture how someone outside the situation, someone who knows neither of you, would size it up in one sentence. When the timer ends, rate your anger again from 0 to 10 and choose one next step: a single clear request, a single question, or a longer timer.

The instruction to avoid the run and the punching bag is the substantive content there, and it is the opposite of what most anger advice says. Arousal-decreasing activities reduced anger at $g = -0.63$ across 154 studies, while arousal-increasing activities showed little benefit and in some cases made things worse (Kjærvi & Bushman, 2024). The zoom-out instruction draws on a second, independent mechanism. Trope and Liberman (2003) describe how construing a situation at a more abstract, psychologically distant level, asking why instead of what exactly, produces a qualitatively different representation than construing it at a concrete, close-up level. Kross et al. (2005) tested this distinction directly on anger-eliciting memories and found that a distanced, why-focused perspective produced significantly less emotional reactivity than an immersed, what-focused one. No item cites its source, because a reader in the middle of an argument is not going to read a citation. Both findings are nonetheless the reason each instruction is what it is. The coaching content for this domain excludes an initial 0-10 emotional self-rating invitation.

What the material does not do

It does not automatically respond, adapt, or personalize. An item is written for anyone training that behavior, and the application holds no information about the individual reader with which to tailor it. Hence the standard weekly article structure, which packages brief psychoeducation inside an exemplary everyday situation for concrete framing, complemented by a generalizable behavioral strategy for mastery. The privacy-first approach is the visible cost of the architecture described in Section 8. A product receiving usage data could select the next item on the basis of what the user has actually been doing. BRAKED selects it from the sets active for the templates in the profile, and that is the whole of the mechanism.

8. Architecture as Position

BRAKED has no account system, no login, and no server that receives anything, and Section 3 describes what that means for the user in practice. Section 3 noted that no login lowers the barrier to reaching for the app in the moment. The same absence is also a retention choice: a required account is a plausible deterrent for a privacy-aware user asked to log a record of their own impulsive behavior, and an occasional re-authentication prompt is a plausible dropout point in its own right. This is a straightforward privacy position and Section 1 gave the reasons for taking it. What is less often said by anyone taking this position is what it costs, so I will say it here.

What I do not know

I know how many people have downloaded BRAKED, because Apple tells me. I do not know whether any of them has ever completed a delay.

That is not an exaggeration for effect. No telemetry means no completion rates, no drop-off curves, and no ability to detect that a feature is confusing because nobody is using it. Every design judgment described in Section 5 will remain, at its core, a judgment: I cannot see the actual distribution of durations chosen in practice, and I cannot find out whether the difficulty rating predicts who advances, since that would require linking within-person behavior across time; something no anonymous, one-shot instrument can provide without compromising the privacy guarantee itself. That is the knowledge cost of wanting to provide uncompromising privacy.

A voluntary, anonymous feedback form partially narrows some of these gaps, with important caveats. Respondents are directly asked whether the three-template cap feels restrictive, about right, or generous; whether the weekly coaching cadence feels too slow, appropriate, or too fast; and whether they perceive themselves choosing longer or shorter delays than when they started (as a self-report proxy for whether users climb the ladder toward longer delays or stay on the lowest rung). Which templates tend to be selected together is also directly observable from the form's multi-select response pattern. None of this is telemetry: it is self-report, filtered through whoever chooses to volunteer it, subject to recall bias, and drawn from a self-selected sample of engaged users rather than the full user base; including, definitionally, no one who disengaged before ever seeing the form. It narrows some of the design judgments described above from pure inference to informed inference. However, it does not close the underlying gap or open the foreclosed research path.

Section 12 sets out what a validation study would have to look like instead, and the answer is that it would require building something separate, with consent, outside the product.

The trade-off is absolute

I cannot optimize features based on aggregated real usage data, nor run automated within-subject trials across the user base. That operational trade-off is deliberate: I chose absolute privacy guarantees over product analytics.

Why the trade is defensible

Section 2 described the market this trade sits in. The mental health category is not full of applications carefully accumulating evidence and publishing it, and the ones making evidence claims mostly cannot support them. Being unable to produce data is a different failure from being able to produce it and choosing not to, but the practical position of a reader trying to evaluate either is not very different.

The same position, applied to money

The training instrument stays free permanently and requires no account, indefinitely and with no trial period attached.

The only thing being sold is authored labor and personally delivered expertise, never access to the mechanism or the user's data.

I chose this model out of structural alignment, not commercial convention. An application collecting no data cannot monetize data, removing the most common business model in this category and leaving a clean exchange: selling the individual professional capacity of a psychologist rather than user activity. This produces a result that matters more to me than maximizing revenue: a colleague in mental health care can point a client toward BRAKED without pointing them toward a purchase, because the part that does the training costs nothing. Section 9 takes up what that makes possible.

9. Using BRAKED Alongside Professional Work

This section is addressed to colleagues who might consider mentioning BRAKED to a client and want to know what they would be mentioning.

The record belongs to the client, and only to the client

Most tracking applications that position themselves for clinical use do so by offering the clinician a dashboard. The client's data flows to a server and the practitioner is given a view of it.

BRAKED has no dashboard, because it has no server. If a client wants you to see their record, they show you their phone, or they generate the Export PDF described in Section 3 and share it with you. You are never a party to the transfer unless the client makes you one, on the occasion they choose, with the portion they choose.

The practical consequence is easy to state plainly. Recommending BRAKED creates no data-processing relationship, no controller or processor question, no agreement to put in place, and nothing that can be breached on your side because nothing arrives on your side. A client who stops seeing you keeps their entire history and you retain none of it. For practitioners working under GDPR who have thought carefully about what recommending a tracking application actually commits them to, this is the point that matters most.

Recommending it also commits the client to nothing: the training instrument is free and permanent, with no trial period that expires into a paywall.

What is in the record

If a client brings the Export PDF described in Section 3 to a session, what you are looking at is: total and cancelled sessions and the completion rate (each cancellation showing the real elapsed hold time before it, not just the fact of cancelling); perceived average difficulty, average satisfaction and the resolved-session count, each for the current week, the month, and for the lifetime; the same three figures broken out by behavior, limited to behaviors actually trained; onboarding history; a 10-point difficulty and satisfaction rating and the resolution outcome for each individual session; a line chart of perceived difficulty and satisfaction over the last ten sessions for any trained behavior; any free-text notes; and distributions of sessions by weekday and by time of day.

The distributions are usually the practically interesting part. A client who could not previously say when their practice happens can now point at a chart showing that half of their sessions start at night, or that Friday accounts for most of a month. The record holds the delays a client chose to run; urges acted on without opening the app do not appear in it, which is worth keeping in mind when reading the pattern. That is the kind of observation that is difficult to recover from memory and easy to work with once it exists, which is the finding from Harkin et al. (2016) discussed in Section 1 arriving in a usable form.

The notes field is unstructured and unprompted. The app does not ask leading questions and does nothing with what is written there beyond storing it for the user's own referencing and reflection at will.

What the app does not do

It does not diagnose, screen, score, or flag. It does not classify a pattern as concerning. It does not interpret the notes, generate summaries, or produce anything resembling a formulation. It does not tell the user what their data means about them. If you are weighing whether the app might say something to a client that you would then have to undo, it will not.

The one thing it does assert is the framing of a completed session, and that is the item most useful to raise with a client directly.

What to tell a client to expect

An honestly completed delay in which the impulse did not go away is a successful session.

This follows from the evidence in Section 4, but a client is unlikely to arrive at it on their own. The app records the impulse question after congratulating the user precisely because the completion is the exercise, but someone who answers *no* three sessions running may still conclude they are doing it wrong. If you mention the app, mention this.

Two smaller points. The lowest rung is five minutes and it makes sense to start there regardless of how confident a client feels, since the ladder depends on early completions being completions. And if the client subscribes, coaching material only reaches the three templates active in the profile, which in practice means an early conversation about which few behaviors matter enough to receive it; the training instrument itself has no such limit, and a client can run timers against any of the seven templates, or a custom one, from the start.

Boundaries

BRAKED is not clinical treatment and does not become treatment when a clinician recommends it. It is a self-monitoring and delay-practice instrument, of the kind described in Section 1 as the homework component, to be used as a self-training tool or in conjunction with professional mental health support.

Four of the templates carry constraints that follow from the behavior rather than from general policy; Section 10 states them in one place.

Finally, a limitation that matters more in your hands than in a user's. The 10-point difficulty and satisfaction items are not validated psychometric scales. They are useful as a conversation prompt and as a within-person trend a client can see for themselves. They should not be scored, or treated as objective measurement.

10. Safety and Scope

BRAKED is a digital self-monitoring and delay-practice instrument. Every operational claim in this document is bounded strictly by that functional scope.

The author holds a Master of Science in Psychology. Psychologist (MSc) is a legally protected academic and professional title across Switzerland, Germany, and most EU member states, denoting a recognized European professional qualification.

Crisis

An acute crisis is not a timer problem. Anyone experiencing ideation of harm to self or others, severe withdrawal, psychotic symptoms, or a medical or psychiatric emergency should contact a qualified healthcare provider or emergency services.

Age

BRAKED carries an 18+ rating on the App Store as of October 2026. At initial submission, Apple's age rating questionnaire automatically calculated a 4+ rating. The current 18+ rating is my own choice, made because behavioral self-training content of this kind, and the priority-tier individual psychological consultations via EPISTEMIST, are better suited to users old enough to apply them of their own accord.

Template-specific constraints

Four templates carry restrictions that follow from the behavior itself rather than from a general policy. They arise in Section 6 and are stated here so they sit in one place.

Eating. The ladder is capped at two hours. Beyond that point a delay before eating stops being practice and becomes restriction. No coaching material for this template, or for **Sweet Treats**, contains calorie targets, portion prescriptions, meal timing rules, weight-loss framing, or numeric dietary guidance of any kind. Eating difficulties severe enough to require specialized treatment fall outside what a delay timer can address.

Trading. Trading at an intensity that meets criteria for gambling disorder requires clinical care. The template addresses impulsive execution, which is a narrower problem than disordered gambling and should not be confused with it.

Anger Reaction. This template addresses everyday conflict rather than intimate partner violence, which requires specialized professional support. The one-day ceiling exists because a pause in a conflict is meant to end.

Constraints on all material

Four rules hold across every template, and each follows from evidence set out earlier rather than from a wish to be careful.

No material asks a user to stop wanting something, because suppression instructions produce rebound.

No material promises that a given urge will weaken or pass within a stated interval, or that the delay itself removes wanting. As stated in the opening chapters, the supported claim concerns what happens between the urge and the action.

No material uses aversive future imagery. The current literature focuses more frequently on the benefits of positive valence, and frightening someone is in my view a less effective way to fortify personal strengths.

No material frames a lapse as a verdict on the person. Shame risks concealment rather than correction, and an application that converts a missed day into a character judgment is working against itself.

11. What I Have Chosen Not to Claim

The preceding sections make a number of claims. This one states what they do not amount to.

The distance between construct and product

Everything in Sections 4 through 6 is evidence that the constructs BRAKED operationalizes have support in the literature. BRAKED itself is neither built to, nor aiming to, generate and deliver research evidence of its own.

These are different claims and the distance between them is the most important thing in this document. Implementation intentions have a meta-analytic effect of $d = .65$ across 94 tests. Arousal-decreasing activities reduce anger at $g = -0.63$ across 154 studies. Neither finding was, however, produced by anyone using this application. The app follows the path opened by peer-reviewed research, but it has no means of testing and verifying its own effectiveness, not least because of its uncompromising privacy promise.

Thus, BRAKED has no efficacy data of its own, and Section 8 explains why it cannot generate any as currently built.

The numbers are judgments

The specific ladder values are design decisions. The rungs at 5, 10, 15 and 30 minutes are defensible starting points on a graded hierarchy, chosen pragmatically to maximize practicability and utility. The same applies to the ceilings, the three-template coaching maximum, and the weekly content interval. Section 5 labels each of these T4 for exactly this reason.

Nothing in the app should be read as implying that a duration corresponds to a measured property of an urge. Section 4 sets out why I do not believe such a measurement is currently available.

What the literature itself does not settle

Three gaps sit in the evidence rather than in the product, and they constrain what could be claimed even by an application with perfect data.

Behavioral targets are under-studied relative to substances. Most of the research on craving dynamics and delay concerns smoking, alcohol and opioids. Shopping, gaming and social media rest on smaller samples and fewer controlled trials, and inferences from one to the other are inferences.

Durability is poorly characterized. Most intervention studies report at post-treatment or at three to six months. The clearest illustration is in Section 6: cognitive-behavioral techniques produce substantial effects on gambling outcomes at post-treatment and no significant effect at follow-up (Pfund, Forman, et al., 2023). Whether anything trained is retained is, across most of this literature, not established.

The functional scope of cue specificity. Section 5 references cue reactivity and domain specificity to justify providing seven distinct templates rather than a single generic timer. The retrieval-extinction literature often cited alongside this rationale has shown inconsistent replication, and BRAKED does not implement a reconsolidation-based intervention. The templates function as a practical organizing framework matching context to behavior, not as a claims-based implementation of extinction protocols.

Concepts deliberately not used

Willpower does not appear anywhere in BRAKED as a depletable resource, and no material refers to reserves, running out, or having spent one's self-control for the day. A large-scale, preregistered multi-lab replication found no reliable evidence for the ego depletion effect (Hagger et al., 2016), though the finding remains contested rather than settled; building a training model on the resource account either way seemed unwise given the uncertainty.

What the model uses in its place is the process account, on which a period of self-control shifts what a person attends to and what they feel like doing next, rather than draining a supply of anything (Inzlicht & Schmeichel, 2012; Milyavskaya & Inzlicht, 2017). That account carries the same practical consequence, which is that the environment a person arranges in advance matters more than the effort available to them in the moment, and it carries it without a metaphor the evidence no longer supports.

SMART goals, specific, measurable, achievable, relevant, time-bound, are not used either, and the omission is deliberate. The framework is built for projects with an external target and a deadline, and it works well there. An urge is neither. Imposing a measurable target on an appetitive or emotional impulse turns a lapse into a missed metric, which is precisely the pass-or-fail framing Section 1 already rejects for the reasons Dearing et al. (2005), Luoma et al. (2019), and Randles and Tracy (2013) give. BRAKED replaces the external target with the person's own starting point: a record of where they began, read against where they are now, with no goalpost borrowed from a framework built for something else.

The retention premise applies here

Section 2 established that unguided digital tools very likely face rapid user drop-off in the absence of continuous artificial incentives or telemetry-driven nudges. BRAKED is architected directly in response to this reality. Several design choices mitigate these common friction points: eliminating logins removes the highest-friction onboarding barrier; strictly local storage removes the disclosure risk any transfer to a third-party server would create; and a five-minute entry rung requires minimal cognitive effort on the first attempt. These structural choices lower friction as far as software can, though (as noted earlier) no instrument can substitute for intrinsic user practice and deliberate engagement.

12. Open Questions and What Validation Would Require

A study of BRAKED would have to be constructed separately, for the reasons Section 8 gives. This section sets out what that would involve and which questions it would settle.

What a study would have to look like

The architecture makes the design problem specific rather than impossible. A validation study would require a separate build carrying instrumentation, used only by participants who had consented to it, with the public application unchanged. Data would belong to the study rather than to the product, and nothing would flow back into the version anyone else is using.

The outcome measure is the part that needs care. The obvious choice, whether urges diminish, is a poor primary outcome. In the short term, the closest controlled evidence found no difference in urges alongside a significant difference in behavior (Bowen & Marlatt, 2009). Over months, craving may fall if acting on it becomes less frequent (Apolzan et al., 2017), so urge intensity belongs among the secondary measures, assessed at follow-up.

A more appropriate primary outcome is the frequency of the target behavior itself, measured before and after a defined training period. Secondary measures would sensibly include a general self-control instrument such as the Brief Self-Control Scale (Tangney et al., 2004) and a domain-specific measure matched to the template being trained. The completion count would then function as a process variable rather than as an outcome, which is what it is.

A pre-post design with a waitlist control group represents a sound methodological baseline, providing an initial empirical benchmark where, to my knowledge, none currently exists.

What is open to ask

Six questions follow directly from decisions this document labels as judgments. Each is answerable, and answering any of them would replace inference or design judgment with direct evidence.

Does step size matter? Whether the spacing between 5, 10, 15 and 30 minutes is well judged is testable, and so is whether users who skip rungs complete less often than users who climb them.

Does the difficulty rating predict advancement? It is treated throughout as the natural progression signal. Whether it actually forecasts who successfully attempts a longer duration is an empirical question and currently an assumption.

Does completion frequency predict anything durable? Section 5 made an inference from a finding about how often people eat craved foods to a metric counting how often people hold a delay. It is a reasonable inference and it has never been tested in this form. This is the study I would most like to see.

Is three the right number of templates to coach at once? Two might be better. Four might be worse or vice versa. The reasoning in Section 5 supports a small number for coaching bandwidth and does not identify which small number.

Does delay-then-permit outperform abstinence framing? The Sweet Treats template is built on the first, following the frequency finding. A direct comparison against the second would test the design decision rather than the underlying claim.

For whom does any of this work? Trait impulsivity, attentional difficulties, age and cultural context are all plausible moderators, and I am not aware of work examining any of them in this context.

An invitation

If you work in this area and any of the above is interesting, the instrumented build described here is something I would construct for a properly designed study. Every source behind the claims made here is listed in the reference list, with a DOI wherever one exists.

Correspondence on any of this reaches me at support@epistemist.de.

13. Closing Note

A five-minute delay is an intentionally low-friction commitment. Someone stands in a kitchen, or an aisle, or at the end of a bad conversation, taps a button, puts the phone down and waits. Nothing happens. Then they decide again, perhaps as they would have before the delay, hopefully differently and more constructively than they would have without it.

This focus on simplicity is deliberate and central to the architecture. Section 1 described what defeats behavioral homework, which is not necessarily motivation but practical and temporal reach: the distance between the moment that matters and the tools available in it. An instrument asking for 15 seconds and no account has a chance of being present in that moment. One asking for more does not.

What I have tried to do is give every reason it is built the way it is, why I believe it can work, name the source behind each one, and say plainly which decisions rest on my own academically qualified judgment. This is a professional judgment rather than a hope, built on the research evidence gathered here and on more than a decade of watching this kind of work succeed and fail in practice, and I stand behind it.

That is what I would want to see from anyone asking me to recommend their application to a client or to a friend. It seemed reasonable to produce it before asking.

Onur Yavuz, Psychologist (MSc University of Zurich, Switzerland)

Founder at EPISTEMIST

Developer of the BRAKED App | EPISTEMIST LABS

References

- Abramowitz, J. S., Deacon, B. J., & Whiteside, S. P. H. (2019). *Exposure therapy for anxiety: Principles and practice* (2nd ed.). Guilford Press.
- Apolzan, J. W., Myers, C. A., Champagne, C. M., Beyl, R. A., Raynor, H. A., Anton, S. A., Williamson, D. A., Sacks, F. M., Bray, G. A., & Martin, C. K. (2017). Frequency of consuming foods predicts changes in cravings for those foods during weight loss: The POUNDS Lost Study. *Obesity*, 25(8), 1343–1348. <https://doi.org/10.1002/oby.21895>
- Atance, C. M., & O'Neill, D. K. (2001). Episodic future thinking. *Trends in Cognitive Sciences*, 5(12), 533–539. [https://doi.org/10.1016/S1364-6613\(00\)01804-0](https://doi.org/10.1016/S1364-6613(00)01804-0)
- Bandura, A. (1977). Self-efficacy: Toward a unifying theory of behavioral change. *Psychological Review*, 84(2), 191–215. <https://doi.org/10.1037/0033-295X.84.2.191>
- Bao, Y., Zhou, H., Geng, F., & Hu, Y. (2025). The relation between game disorder and interruption during game is mediated by game craving. *Frontiers in Psychology*, 16, 1579016. <https://doi.org/10.3389/fpsyg.2025.1579016>
- Baumel, A., Muench, F., Edan, S., & Kane, J. M. (2019). Objective user engagement with mental health apps: Systematic search and panel-based usage analysis. *Journal of Medical Internet Research*, 21(9), e14567. <https://doi.org/10.2196/14567>
- Bell, I. H., Linardon, J., Mangelsdorf, S. N., Alvarez-Jimenez, M., Valentine, L., Eisner, E., Reynolds, K. A., Torous, J., Myin-Germeys, I., Bucci, S., Mohr, D. C., Reininghaus, U., & Nicholas, J. (2026). Public availability of evidence-based mental health apps: A systematic review and meta-analysis of self-guided apps. *npj Digital Medicine*. <https://doi.org/10.1038/s41746-026-03014-5>
- Bowen, S., & Marlatt, A. (2009). Surfing the urge: Brief mindfulness-based intervention for college student smokers. *Psychology of Addictive Behaviors*, 23(4), 666–671. <https://doi.org/10.1037/a0017127>
- Bulley, A., Miloyan, B., Pepper, G. V., Gullo, M. J., Henry, J. D., & Suddendorf, T. (2019). Cuing both positive and negative episodic foresight reduces delay discounting but does not affect risk-taking. *Quarterly Journal of Experimental Psychology*, 72(8), 1998–2017. <https://doi.org/10.1177/1747021818819777>
- Carter, B. L., & Tiffany, S. T. (1999). Meta-analysis of cue-reactivity in addiction research. *Addiction*, 94(3), 327–340. <https://doi.org/10.1046/j.1360-0443.1999.9433273.x>
- Dearing, R. L., Stuewig, J., & Tangney, J. P. (2005). On the importance of distinguishing shame from guilt: Relations to problematic alcohol and drug use. *Addictive Behaviors*, 30(7), 1392–1404. <https://doi.org/10.1016/j.addbeh.2005.02.002>

Doebel, S., Michaelson, L. E., & Munakata, Y. (2020). Good things come to those who wait: Delaying gratification likely does matter for later achievement (a commentary on Watts, Duncan, & Quan, 2018). *Psychological Science*, 31(1), 97–99.
<https://doi.org/10.1177/0956797619839045>

Ersner-Hershfield, H., Garton, M. T., Ballard, K., Samanez-Larkin, G. R., & Knutson, B. (2009). Don't stop thinking about tomorrow: Individual differences in future self-continuity account for saving. *Judgment and Decision Making*, 4(4), 280–286.

Federal Trade Commission. (2023, March 2). *FTC to ban BetterHelp from revealing consumers' data, including sensitive mental health information, to Facebook and others for targeted advertising* [Press release]. <https://www.ftc.gov/news-events/news/press-releases/2023/03/ftc-ban-betterhelp-revealing-consumers-data-including-sensitive-mental-health-information-facebook>

Ferster, C. B., & Skinner, B. F. (1957). *Schedules of reinforcement*. Appleton-Century-Crofts.

Galla, B. M., & Duckworth, A. L. (2015). More than resisting temptation: Beneficial habits mediate the relationship between self-control and positive life outcomes. *Journal of Personality and Social Psychology*, 109(3), 508–525.
<https://doi.org/10.1037/pspp0000026>

Ghebreyesus, T. A., Mpemba, C., & Murthy, V. (2025, July 14). Loneliness and isolation: The hidden threat to global health we can no longer ignore. World Health Organization. <https://www.who.int/news-room/commentaries/detail/loneliness-and-isolation-the-hidden-threat-to-global-health-we-can-no-longer-ignore>

Gollwitzer, P. M. (1999). Implementation intentions: Strong effects of simple plans. *American Psychologist*, 54(7), 493–503. <https://doi.org/10.1037/0003-066X.54.7.493>

Gollwitzer, P. M., & Sheeran, P. (2006). Implementation intentions and goal achievement: A meta-analysis of effects and processes. *Advances in Experimental Social Psychology*, 38, 69–119. [https://doi.org/10.1016/S0065-2601\(06\)38002-1](https://doi.org/10.1016/S0065-2601(06)38002-1)

Hagger, M. S., Chatzisarantis, N. L. D., Alberts, H., Anggono, C. O., Batailler, C., Birt, A. R., Brand, R., Brandt, M. J., Brewer, G., Bruyneel, S., Calvillo, D. P., Campbell, W. K., Cannon, P. R., Carlucci, M., Carruth, N. P., Cheung, T., Crowell, A., De Ridder, D. T. D., Dewitte, S., ... Zwieneberg, M. (2016). A multilab preregistered replication of the ego-depletion effect. *Perspectives on Psychological Science*, 11(4), 546–573.
<https://doi.org/10.1177/1745691616652873>

- Harkin, B., Webb, T. L., Chang, B. P. I., Prestwich, A., Conner, M., Kellar, I., Benn, Y., & Sheeran, P. (2016). Does monitoring goal progress promote goal attainment? A meta-analysis of the experimental evidence. *Psychological Bulletin*, 142(2), 198–229. <https://doi.org/10.1037/bul0000025>
- Hofmann, W., Baumeister, R. F., Förster, G., & Vohs, K. D. (2012). Everyday temptations: An experience sampling study of desire, conflict, and self-control. *Journal of Personality and Social Psychology*, 102(6), 1318–1335. <https://doi.org/10.1037/a0026545>
- Hofmann, W., Friese, M., & Strack, F. (2009). Impulse and self-control from a dual-systems perspective. *Perspectives on Psychological Science*, 4(2), 162–176. <https://doi.org/10.1111/j.1745-6924.2009.01116.x>
- Hofmann, W., Vohs, K. D., & Baumeister, R. F. (2012). What people desire, feel conflicted about, and try to resist in everyday life. *Psychological Science*, 23(6), 582–588. <https://doi.org/10.1177/0956797612437426>
- Hua, Y., Siddals, S., Ma, Z., Galatzer-Levy, I., Xia, W., Hau, C., Na, H., Flathers, M., Linardon, J., Ayubcha, C., & Torous, J. (2025). Charting the evolution of artificial intelligence mental health chatbots from rule-based systems to large language models: A systematic review. *World Psychiatry*, 24(3), 383–394. <https://doi.org/10.1002/wps.21352>
- Huckvale, K., Torous, J., & Larsen, M. E. (2019). Assessment of the data sharing and privacy practices of smartphone apps for depression and smoking cessation. *JAMA Network Open*, 2(4), e192542. <https://doi.org/10.1001/jamanetworkopen.2019.2542>
- Inzlicht, M., & Schmeichel, B. J. (2012). What is ego depletion? Toward a mechanistic revision of the resource model of self-control. *Perspectives on Psychological Science*, 7(5), 450–463. <https://doi.org/10.1177/1745691612454134>
- Kakoschke, N., Henry, B. A., Cowley, M. A., & Lee, K. (2024). Tackling cravings in medical weight management: An update on pathophysiology and an integrated approach to treatment. *Nutrients*, 16(19), 3238. <https://doi.org/10.3390/nu16193238>
- Kjærvik, S. L., & Bushman, B. J. (2024). A meta-analytic review of anger management activities that increase or decrease arousal: What fuels or douses rage? *Clinical Psychology Review*, 109, 102414. <https://doi.org/10.1016/j.cpr.2024.102414>

- Kross, E., Ayduk, O., & Mischel, W. (2005). When asking “why” does not hurt: Distinguishing rumination from reflective processing of negative emotions. *Psychological Science*, 16(9), 709–715. <https://doi.org/10.1111/j.1467-9280.2005.01600.x>
- Larsen, M. E., Huckvale, K., Nicholas, J., Torous, J., Birrell, L., Li, E., & Reda, B. (2019). Using science to sell apps: Evaluation of mental health app store quality claims. *npj Digital Medicine*, 2, 18. <https://doi.org/10.1038/s41746-019-0093-1>
- Linardon, J., Fuller-Tyszkiewicz, M., Firth, J., Goldberg, S. B., Anderson, C., McClure, Z., & Torous, J. (2024). Systematic review and meta-analysis of adverse events in clinical trials of mental health apps. *npj Digital Medicine*, 7, 363. <https://doi.org/10.1038/s41746-024-01388-y>
- Linardon, J., Torous, J., Firth, J., Cuijpers, P., Messer, M., & Fuller-Tyszkiewicz, M. (2024). Current evidence on the efficacy of mental health smartphone apps for symptoms of depression and anxiety: A meta-analysis of 176 randomized controlled trials. *World Psychiatry*, 23(1), 139–149. <https://doi.org/10.1002/wps.21183>
- Luoma, J. B., Chwyl, C., & Kaplan, J. (2019). Substance use and shame: A systematic and meta-analytic review. *Clinical Psychology Review*, 70, 1–12. <https://doi.org/10.1016/j.cpr.2019.03.002>
- Marshall, J. M., Dunstan, D. A., & Bartik, W. (2019). The digital psychiatrist: In search of evidence-based apps for anxiety and depression. *Frontiers in Psychiatry*, 10, 831. <https://doi.org/10.3389/fpsy.2019.00831>
- Marshall, J. M., Dunstan, D. A., & Bartik, W. (2020). Apps with maps: Anxiety and depression mobile apps with evidence-based frameworks: Systematic search of major app stores. *JMIR Mental Health*, 7(6), e16525. <https://doi.org/10.2196/16525>
- Michie, S., Abraham, C., Whittington, C., McAteer, J., & Gupta, S. (2009). Effective techniques in healthy eating and physical activity interventions: A meta-regression. *Health Psychology*, 28(6), 690–701. <https://doi.org/10.1037/a0016136>
- Milyavskaya, M., & Inzlicht, M. (2017). What’s so great about self-control? Examining the importance of effortful self-control and temptation in predicting real-life depletion and goal attainment. *Social Psychological and Personality Science*, 8(6), 603–611. <https://doi.org/10.1177/1948550616679237>
- Nordby, K., Løkken, R. A., & Pfuhl, G. (2019). Playing a video game is more than mere procrastination. *BMC Psychology*, 7, 33. <https://doi.org/10.1186/s40359-019-0309-9>

Office of the Surgeon General. (2023). Our epidemic of loneliness and isolation: The U.S. Surgeon General's advisory on the healing effects of social connection and community [Advisory]. U.S. Department of Health and Human Services. <https://www.hhs.gov/sites/default/files/surgeon-general-social-connection-advisory.pdf>

Peck, S., Preston, E., Smith, K. B., & Madden, G. J. (2022). Reducing impulsive choice: VIII. Effects of delay-exposure training in female rats. *Behavioural Processes*, 197, Article 104622. <https://doi.org/10.1016/j.beproc.2022.104622>

Pfund, R. A., Forman, D. P., Whalen, S. K., Zech, J. M., Ginley, M. K., Peter, S. C., McAfee, N. W., & Whelan, J. P. (2023). Effect of cognitive-behavioral techniques for problem gambling and gambling disorder: A systematic review and meta-analysis. *Addiction*, 118(9), 1661–1674. <https://doi.org/10.1111/add.16221>

Pfund, R. A., Ginley, M. K., Kim, H. S., Boness, C. L., Horn, T. L., & Whelan, J. P. (2023). Cognitive-behavioral treatment for gambling harm: Umbrella review and meta-analysis. *Clinical Psychology Review*, 105, 102336. <https://doi.org/10.1016/j.cpr.2023.102336>

Price, K. L., Saladin, M. E., Baker, N. L., Tolliver, B. K., DeSantis, S. M., McRae-Clark, A. L., & Brady, K. T. (2010). Extinction of drug cue reactivity in methamphetamine-dependent individuals. *Behaviour Research and Therapy*, 48(9), 860–865. <https://doi.org/10.1016/j.brat.2010.05.010>

Randles, D., & Tracy, J. L. (2013). Nonverbal displays of shame predict relapse and declining health in recovering alcoholics. *Clinical Psychological Science*, 1(2), 149–155. <https://doi.org/10.1177/2167702612470645>

Rickard, N. S., Kurt, P., & Meade, T. (2022). Systematic assessment of the quality and integrity of popular mental health smartphone apps using the American Psychiatric Association's app evaluation model. *Frontiers in Digital Health*, 4, 1003181. <https://doi.org/10.3389/fdgth.2022.1003181>

Scholten, H., Scheres, A., de Water, E., Graf, U., Granic, I., & Luijten, M. (2019). Behavioral trainings and manipulations to reduce delay discounting: A systematic review. *Psychonomic Bulletin & Review*, 26(6), 1803–1849. <https://doi.org/10.3758/s13423-019-01629-2>

Sheeran, P., Maki, A., Montanaro, E., Avishai-Yitshak, A., Bryan, A., Klein, W. M. P., Miles, E., & Rothman, A. J. (2016). The impact of changing attitudes, norms, and self-efficacy on health-related intentions and behavior: A meta-analysis. *Health Psychology*, 35(11), 1178–1188. <https://doi.org/10.1037/hea0000387>

Skinner, B. F. (1951). How to teach animals. *Scientific American*, 185(6), 26–29.
<https://doi.org/10.1038/scientificamerican1251-26>

Skinner, B. F. (1953). *Science and human behavior*. Macmillan.

Smith, T. R., Panfil, K., Bailey, C., & Kirkpatrick, K. (2019). Cognitive and behavioral training interventions to promote self-control. *Journal of Experimental Psychology: Animal Learning and Cognition*, 45(3), 259–279.
<https://doi.org/10.1037/xan0000208>

Sperber, J. F., Vandell, D. L., Duncan, G. J., & Watts, T. W. (2024). Delay of gratification and adult outcomes: The Marshmallow Test does not reliably predict adult functioning. *Child Development*, 95(6), 2015–2029.
<https://doi.org/10.1111/cdev.14129>

Stone, A. A., Shiffman, S., Schwartz, J. E., Broderick, J. E., & Hufford, M. R. (2002). Patient non-compliance with paper diaries. *BMJ*, 324(7347), 1193–1194.
<https://doi.org/10.1136/bmj.324.7347.1193>

Strack, F., & Deutsch, R. (2004). Reflective and impulsive determinants of social behavior. *Personality and Social Psychology Review*, 8(3), 220–247.
https://doi.org/10.1207/s15327957pspr0803_1

Tangney, J. P., Baumeister, R. F., & Boone, A. L. (2004). High self-control predicts good adjustment, less pathology, better grades, and interpersonal success. *Journal of Personality*, 72(2), 271–324. <https://doi.org/10.1111/j.0022-3506.2004.00263.x>

Tangney, J. P., Mashek, D., & Stuewig, J. (2005). Shame, guilt, and embarrassment: Will the real emotion please stand up? *Psychological Inquiry*, 16(1), 44–48.
<https://www.jstor.org/stable/20447260>

Tangney, J. P., Stuewig, J., & Hafez, L. (2011). Shame, guilt, and remorse: Implications for offender populations. *Journal of Forensic Psychiatry & Psychology*, 22(5), 706–723.
<https://doi.org/10.1080/14789949.2011.617541>

Trope, Y., & Liberman, N. (2003). Temporal construal. *Psychological Review*, 110(3), 403–421. <https://doi.org/10.1037/0033-295X.110.3.403>

Wang, D., Hagger, M. S., & Chatzisarantis, N. L. D. (2020). Ironic effects of thought suppression: A meta-analysis. *Perspectives on Psychological Science*, 15(3), 778–793.
<https://doi.org/10.1177/1745691619898795>

Wasil, A. R., Gillespie, S., Patel, R., Petre, A., Venturo-Conerly, K. E., Shingleton, R. M., Weisz, J. R., & DeRubeis, R. J. (2020). Reassessing evidence-based content in popular smartphone apps for depression and anxiety: Developing and applying user-adjusted analyses. *Journal of Consulting and Clinical Psychology*, 88(11), 983–993.

<https://doi.org/10.1037/ccp0000604>

Wasil, A. R., Venturo-Conerly, K. E., Shingleton, R. M., & Weisz, J. R. (2019). A review of popular smartphone apps for depression and anxiety: Assessing the inclusion of evidence-based content. *Behaviour Research and Therapy*, 123, 103498.

<https://doi.org/10.1016/j.brat.2019.103498>

Watts, T. W., Duncan, G. J., & Quan, H. (2018). Revisiting the marshmallow test: A conceptual replication investigating links between early delay of gratification and later outcomes. *Psychological Science*, 29(7), 1159–1177.

<https://doi.org/10.1177/0956797618761661>

Wegner, D. M. (1994). Ironic processes of mental control. *Psychological Review*, 101(1), 34–52. <https://doi.org/10.1037/0033-295X.101.1.34>

Wolpe, J. (1958). *Psychotherapy by reciprocal inhibition*. Stanford University Press.

Yang, F., Wei, J., Zhao, X., & An, R. (2025). Artificial intelligence-based mobile phone apps for child mental health: Comprehensive review and content analysis. *JMIR mHealth and uHealth*, 13, e58597. <https://doi.org/10.2196/58597>

Ye, J.-Y., Ding, Q.-Y., Cui, J.-F., Liu, Z., Jia, L.-X., Qin, X.-J., Xu, H., & Wang, Y. (2022). A meta-analysis of the effects of episodic future thinking on delay discounting. *Quarterly Journal of Experimental Psychology*, 75(10), 1876–1891.

<https://doi.org/10.1177/17470218211066282>

Imprint and Contact

Information pursuant to § 5 DDG and § 18 (2) MStV.

Onur Yavuz
c/o MDC Management #4163
Welserstraße 3
87463 Dietmannsried
Germany

Email: support@epistemist.de

Responsible for content pursuant to § 18 (2) MStV: Onur Yavuz, at the address above.

Professional designation. Psychologist (MSc). Psychologist (MSc) is a protected academic and professional title in Switzerland (PsyG) and across the EU, including Germany (§ 5 UWG), and denotes a European professional qualification rather than a United States, UK or Canadian license. The designation is a graduation title, earned on completion of the degree, and does not imply a university faculty appointment or ongoing academic affiliation.

Liability for content. As a service provider I am responsible under § 7 (1) DDG for my own content on these pages under general law. This document has been prepared with care and to the best of my knowledge, and no guarantee is given as to the accuracy, completeness or currency of the information in it. Content may be revised without notice. This document is versioned V. 1.3 in line with the app release family it describes. Liability for third-party content under §§ 8 to 10 DDG arises only from the point at which a specific infringement becomes known to me, and any such content will be removed on becoming aware of it.

Liability for links. This document and the associated website may reference external third-party sources. I have no influence over their content and accept no liability for it. Responsibility for the content of linked pages rests with their respective provider or operator.

Copyright. The content and works created by the operator on these pages are subject to German copyright law. Reproduction, modification, distribution or exploitation of any kind beyond the limits of copyright requires the written consent of the respective author.

Notice and Intellectual Property. This document is published for transparency, so that practitioners, researchers, and users can examine the reasoning behind BRAKED. The system architecture detailed in this document (including its templates, ladders, coaching delivery mechanism, and underlying behavioral intervention logic) constitutes proprietary intellectual property; copying or replicating this design is prohibited. The coaching material referenced in Section 7 is a separate, living body of work, updated by a psychologist at EPISTEMIST on at least a quarterly basis and likewise proprietary intellectual property. This document and its contents may not be used to train artificial intelligence models.

BRAKED

by EPISTEMIST LABS

A BEHAVIORAL MICROTRAINING INSTRUMENT
FOR URGE AND IMPULSE CONTROL

support@epistemist.de

