SplitHomework — Productization Build Plan

From a one-page front-end demo to a shippable "earn your feed by doing the work" study tool — built for the phone-shaped attention span, especially ADHD

Working document · drafted May 2026 · owner: Product · status: pre-build scoping · companion to the Product Thesis (docs/SplitHomework-Product-Thesis)

TL;DR

The demo proves the loop in a browser tab. To be a *product*, three things have to get real: (1) the **"homework"** side — real assignments or a credible practice ladder, not just generated MCQs; (2) the **gate** — a soft in-app lock is easy to route around (open YouTube in another tab), so a real version needs a browser-extension or device-level block; (3) the **evidence** — "does this actually get more homework done for ADHD students" is the only question that matters and the demo can't answer it. Pick **one customer** first (consumer parents of ADHD kids is the cleanest wedge), build the gate + a real homework source for that customer, run a small structured pilot, and only then widen. Realistic effort to a real v1: **~3–4 months, 2–3 people**; to a school/SpEd-ready, evidence-backed version: **~9–15 months** and at least one partnership.

1 THE STRATEGIC FORK — PICK THE CUSTOMER FIRST

The same loop sells to four buyers who want different products. You can't serve all four in v1; the build, the gate design, and the content source all branch on this choice.

CUSTOMER	WHAT THEY'D PAY FOR	WHY IT'S HARD	VERDICT
Parents of an ADHD student (consumer)	"My kid will actually start their homework." Set-and-forget on the family laptop; a parent view of "did the work get done"; willing to pay a subscription for something that works.	Reaching them cheaply; the kid has to not hate it; the gate has to be hard enough that the kid can't trivially bypass it.	**v1 wedge.** Clear pain, clear payer, no procurement cycle.
ADHD coaches / tutors / EF specialists	A structured tool to assign between sessions; visibility into follow-through; a shared vocabulary ("activation energy", contingency).	Small market; they want client-management features, not just the loop; long trust-building sale.	Strong **design partners** for v1; a real channel by v2 — not the v1 wedge.
Schools / SpEd departments	A documentable accommodation they can attach to a 504/IEP; aggregate reporting; evidence it works.	Procurement, data-privacy (FERPA, state laws), district pilots, and a (rightful) demand for efficacy data you won't have for a year+.	**Phase 3.** Real, but gated on evidence and a very different go-to-market.
Students themselves (self-buy; college / AuDHD adults)	"I want to make myself study." Self-imposed contingency, like Forest or website blockers, but with the reward built in.	People who'd self-buy a focus tool are a crowded, low-willingness-to-pay segment; high churn; the homework content matters less, so you compete on the gate alone.	Adjacent **upsell** later (a "self mode"); not the wedge.

The pick: consumer, parents of ADHD students, US, web first. Cleanest pain → payer line, no procurement, and it forces you to solve the two things that actually make this a product (a real gate, a real homework source) rather than polishing the demo. Everything in §3 is sequenced for that customer.

2 GAP ANALYSIS — WHAT'S ACTUALLY MISSING TODAY

Nothing on the market is squarely "earn an educational feed by doing the work, with both halves in one place." Adjacent tools each own a piece:

WHAT EXISTS	WHAT IT DOES	WHAT IT DOESN'T DO (THE GAP SPLITHOMEWORK SITS IN)
Forest, Flora, Pomodoro timers	Time-boxing + a soft gamified penalty for leaving the app.	No *reward* for finishing — just a stick; no homework; no contingent access to anything; nothing built for ADHD specifically.
Website blockers — Cold Turkey, Freedom, LeechBlock, SelfControl	Hard block of distracting sites for a set window.	Pure abstinence; no positive contingency; nothing earned back through work; brittle against a determined user; no homework loop.
"Earn screen time" parental controls — Screen Time, Family Link, Bark, Qustodio	Parent-set time budgets; some let chores/tasks add minutes.	The "task" is unverified (a parent toggle); the reward is the whole phone, not a metered educational feed in-context; not a study tool; the kid experiences it as surveillance.
Practice / homework apps — Khan, IXL, DeltaMath, Quizlet	Real curriculum-aligned practice and (for some) teacher-assigned work.	No contingent-reward layer; no "feed" reward at all; built for coverage, not for getting an avoidant student to *start*.
Body-doubling / focus rooms — Focusmate, Study Together	Social accountability via co-working.	Needs a scheduled human; no reward mechanic; no homework content; not always a fit for kids.
Commitment devices — Beeminder, StickK	Self-imposed stakes for hitting goals.	Stakes (a stick), abstract goals, delayed — the opposite of immediate-salient-positive; adult, niche.

The white space, as a sentence: **a positive, immediate, metered reward — an educational feed — earned in-context by doing actual schoolwork, with a gate hard enough to matter, designed around how ADHD attention responds to incentives.** Pieces of this exist; the combination doesn't. That combination is the moat — and the part you build yourself (§7).

> **THE POSITIONING SENTENCE**
>
> Not a focus timer (no reward), not a website blocker (pure abstinence), not a screen-time-budget app (unverified task, no study loop), not a practice app (no carrot to start). SplitHomework is the **earned-reward layer wrapped around real schoolwork** — and it should be pitched on the ADHD use case and the homework loop, not the crowded word "focus."

3 PHASED ROADMAP

Phase 0 — the demo WHERE WE ARE

done · ~1–2 weeks of build

BUILT	Split-screen web app; left = runtime-generated math MCQs (4 levels) with worked solutions; right = a YouTube panel locked until you earn time (+1:00 right / +0:30 wrong, 10-min cap, ticks down while watching, re-locks at 0); 16 educational categories with verified playlists, a "More ▾" picker, a paste-your-own flow; a stats grid (answered / accuracy / ACT-style "where you'd stand" / time earned); state in localStorage; deploys static to Vercel.
PROVES	The loop is legible and fun in 30 seconds: do problem → bank a minute → watch → re-lock → back to work. Good enough to put in front of parents, coaches, and a few ADHD students for reactions.
CAN'T PROVE	Whether it gets more homework done; whether the gate survives a motivated teenager; anything about real assignments.

EXIT: 5–10 qualitative sessions with target families / coaches. If the reaction is "this is just YouTube next to homework" → the framing or the gate is wrong; fix before Phase 1.

Phase 1 — a real v1 for one customer BUILD

~3–4 months · 2–3 people

HOMEWORK SIDE	Move past "generated MCQs only." Two options, in order of effort: (a) **self-entered work** — the student/parent loads a worksheet, problem set, or reading and checks each item done (+ a quick self-check); (b) **integrations** — pull assignments from Google Classroom / Canvas, or embed a partner practice ladder (Khan, DeltaMath) so completion is verified. v1 ships (a) + at least one of (b).
THE GATE	A **browser extension** (Chrome / Edge first) that blocks a configurable distraction list (YouTube, TikTok, Instagram, Reddit, games) and only opens metered windows the work-loop unlocks. The embedded in-app panel stays as the "nice" path; the extension is what makes the lock real. Parent sets the list + rates; a parent passcode to change them.
ACCOUNTS & SYNC	Light auth (email / Google), cross-device state, a simple **parent view**: problems done, time earned, sessions, this week vs. last. No grades, no leaderboards — consistent with the thesis.
ANTI-GAMING	Wrong answers already pay less and force a worked solution; add accuracy/streak-decayed payouts (guess-spamming drops the per-answer payout toward zero), per-day caps, and "active work" detection so idle time doesn't bank.
TELEMETRY	Instrument the falsifiable bets from the thesis: homework items completed/session, time-to-start, time-to-resume, session-end reason (cap hit vs. closed), retention. Privacy-first; minimal PII; parent-visible data only.

EXIT: a structured 4–6 week pilot with ~20–40 families (ideally recruited via ADHD coaches). Read: does homework throughput rise vs. their baseline? Do sessions end on "out of time" not "quit"? Would they keep paying? If yes on all three → Phase 2. If "fun but homework didn't move" → the thesis (as stated) is wrong; rethink before spending more.

3 PHASED ROADMAP (CONTINUED)

Phase 2 — make it stick & widen the funnel SCALE

~3–5 months after Phase 1

REACH Mobile (PWA or a thin native shell) + iOS / Android Screen-Time-style integration so the gate works on the device the kid actually uses; a coach / tutor console (assign sets, see follow-through across clients) — the v2 channel.

CONTENT DEPTH More homework subjects and a real practice ladder beyond math (reading, writing prompts, science, languages) so the left side is useful for more than one class; richer curated feed categories and per-student feed allow / deny lists.

REWARD TUNING Personalizable rates and caps (some kids need 30s bursts, some 3-min); an optional "earn toward a non-screen reward" mode for families who don't want screen time as the reinforcer at all.

EXIT: healthy retention & word-of-mouth in the consumer + coach channels; a clean data story on homework throughput — the credibility needed to walk into a school.

Phase 3 — the school / SpEd lane EXPAND

~9–15 months out · partnership-gated

WHAT Position it as a documentable executive-function support a school can attach to a 504 / IEP: a managed-device build, aggregate (not per-kid-public) reporting, FERPA / state-privacy compliance, district admin controls, LMS roster sync.

EVIDENCE Partner with an education researcher on a small efficacy study (ideally randomized or waitlist-control) on homework completion / task initiation. This is the unlock for the lane and for honest marketing.

RISK GATE Do not enter this lane on consumer momentum alone — without the study you'll over-claim, and one bad district reference travels. The study is the price of admission.

4 TEAM & BUDGET (ORDER-OF-MAGNITUDE)

US contractor-rate ballpark; roughly halve with offshore or equity-heavy founding hires. Numbers get you *through each phase's exit*, not to steady-state.

PHASE	PEOPLE	DURATION	ROUGH COST	WHAT THE MONEY BUYS
0 (done)	1 (you + AI)	~1–2 wks	~$0	The demo. Done.
1 — real v1	1–2 full-stack devs; PT designer; PT content / curriculum person; you on product / pilot	~3–4 mo	~$120k–220k	Browser-extension gate; accounts + sync + parent view; self-entered work + 1 LMS / practice integration; anti-gaming; telemetry; the recruited pilot.
2 — stick & widen	2–3 devs; designer; content; PT growth; PT support	~3–5 mo	~$200k–400k	Mobile + device-level gate; coach console; more subjects; reward tuning; the consumer + coach go-to-market.
3 — school / SpEd	+ enterprise eng for compliance & admin; a researcher (grant or contract); PT sales / partnerships	~6–9 mo	~$300k–600k+	Managed-device build, FERPA / privacy, district pilots, the efficacy study.

~3–4 mo
TO A REAL, SELLABLE V1

~9–15 mo
TO SCHOOL/SPED-READY, EVIDENCE-BACKED

2–3
PEOPLE FOR V1

1
NON-NEGOTIABLE LATER HIRE: A RESEARCHER

5 TOP RISKS & MITIGATIONS

RISK	MITIGATION
The gate is too soft. The kid just opens YouTube in another tab / app and the premise collapses.	The in-app panel is the carrot; the **browser extension / device-level block is the product** from v1 on — parent-set, passcode-locked. Be honest in marketing: it's a strong scaffold, not an unbreakable cage — and pair it with the positive loop so the kid *wants* the legit path.
The reward becomes the distraction. Sessions turn into "watched 10 min, did 2 problems."	Small caps, tick-down while watching, mandatory return to work to re-earn; per-day total caps; telemetry that flags "feed-dominant" sessions; if it shows up at scale, tighten the economy or make the default cap stricter.
Gaming. Guess-spam easy problems to bank time without learning.	Wrong answers pay less + force the worked solution (already in the demo); add accuracy/streak-decayed payouts, harder-problem bonuses, and "active work" gating — make the cheapest path to banked time actually be doing the work.
"Screen time as a reward" is contested. Some parents / clinicians won't reinforce schoolwork with YouTube.	Lead with the design discipline (educational-only feed, metered, capped, anti-engagement); offer a **non-screen reward mode**; publish the contingency-management rationale *and its limits*. Don't oversell; let skeptical families pick the alternate reinforcer.
YouTube / platform dependency & minors. Embed terms, video availability, a feed of third-party videos shown to kids, COPPA if under-13.	Curated playlists from stable channels (done); design the for-minors version carefully (no general feed, tight allow-lists, no comments, age gating, parental control); keep the feed source swappable so a platform change isn't fatal; get counsel before any under-13 push.
No efficacy data for a long time — and the claim ("gets ADHD kids to do homework") is exactly what people demand proof for.	Phrase v1 marketing as a behavioral scaffold built on established principles, with the pilot's *own* throughput numbers as the proof point; commission the real study before Phase 3 and before any clinical-sounding claim.
Crowded "focus app" perception. Gets lumped in with timers / blockers and dismissed.	The wedge is the differentiator: not abstinence, not a timer — *earned, in-context, metered educational reward tied to real schoolwork, for ADHD*. Lead with the ADHD use case and the homework loop.

6 NEXT 90 DAYS — CONCRETE ACTIONS

- **Weeks 1–2:** 8–12 discovery conversations — ADHD coaches / EF tutors, parents of ADHD teens, a couple of SpEd teachers. Validate the wedge (consumer parents), the pain ("starting homework"), and willingness to pay. Recruit 2–3 coaches as design partners.
- **Weeks 2–4:** tighten the demo into a credible pilot build — the anti-gaming basics (accuracy-decayed payouts, idle detection), a minimal parent-view page, a "load your own worksheet and check items off" mode so it works for real homework, basic telemetry. Still no backend if you can avoid it; localStorage + a thin analytics endpoint is fine for a pilot.
- **Weeks 3–6:** prototype the **browser-extension gate** (block list + metered-window unlock wired to the work loop). The riskiest, most differentiating piece — de-risk it early.
- **Weeks 5–10:** run a small structured pilot — ~15–25 families via the coach partners, 3–4 weeks each, with a pre-pilot baseline (how much homework gets done now) and a weekly check-in. Measure throughput, session-end reason, "would you keep using it."
- **Weeks 10–13:** write up the pilot honestly. Go / pivot / no-go on Phase 1 funding. If go: scope the v1 build and decide build-vs-partner for the homework side (§7).

7 BUILD IT YOURSELF VS. BUILD IT WITH A PARTNER

SplitHomework is two halves: the **gate + reward economy** (the novel part — build this yourself, it's the moat) and the **homework content** (the commodity part — partner here, don't rebuild a curriculum). Three partnership shapes for the content half:

PARTNER TYPE	WHAT YOU GET	TRADE-OFF
LMS — Google Classroom, Canvas	Real teacher-assigned work flows in; "completed" is meaningful; instant relevance — "do the thing your teacher actually assigned, then earn time."	API + auth complexity; you're a thin layer on their roster; for the consumer wedge, only works if the kid's school uses one and the parent can connect it.
Practice content — Khan Academy, DeltaMath, IXL, CK-12	A real difficulty ladder across subjects without building one; verified completion; credibility by association.	Licensing / partnership terms; some are school-sold and slow to deal with consumer apps; you're shaping their content into your loop.
Parental-control / screen-time vendor — for distribution, not content	They own the "manage my kid's screen time" relationship and the device-level gating tech; SplitHomework is the *positive* module they lack.	You become a feature in their suite; channel power sits with them; pick one whose brand isn't pure surveillance.

The specific cases: Google Classroom — best *content-side* first integration: huge install base, decent API, and "earn time by finishing what Ms. K assigned" is the most concrete pitch possible (build in Phase 1). **Khan Academy** — best *ladder* partner if self-entered work proves too loose: free, broad, credible, K-12+; explore a content / embed arrangement so the "homework" can be a real Khan set with verified mastery. **A parental-control vendor (a Bark / Qustodio / Canopy-class company)** — the fastest *distribution* if the consumer wedge proves out but CAC is brutal: they have the parents and the gating tech; SplitHomework supplies the missing carrot (a Phase 2 conversation, from a position of having pilot data).

8 APPENDIX — WHAT THE DEMO ALREADY IS, AND WHAT IT ISN'T

CAPABILITY	IN THE DEMO?	WHAT "REAL" NEEDS
Core loop (work → metered earned feed → re-lock)	YES	The thing the product rests on; legible and tuned.
Math problems + worked solutions, 4 levels	YES	Generated, correct-by-construction; a placeholder for "real homework" — fine for a demo, not the product.
16 feed categories, verified playlists, picker, paste-your-own	YES	Swappable; the for-minors version needs tighter allow-lists.
Stats grid (answered / accuracy / ACT-style standing / time earned)	YES	Ported from MathClash; nice-to-have, not load-bearing.
Real homework source (self-entered work or LMS / practice integration)	NO	Phase 1, item 1. Without this it's a toy.
A gate that survives a motivated teenager (browser extension / device-level)	NO	Phase 1, the riskiest piece. The in-app lock is bypassable by design.
Accounts, cross-device sync, parent view	NO	localStorage only today; Phase 1.
Anti-gaming beyond "wrong pays less + show solution"	PARTIAL	Accuracy-decay, idle detection, daily caps — Phase 1.
Telemetry for the falsifiable bets	NO	Phase 1 — until this exists we can't answer the only question that matters.
Mobile / device-level reach; school compliance, admin, reporting; efficacy study	NO	Phases 2–3 — partnership- and research-gated; do not skip the study.

YES in the demo PARTIAL started NO not yet — phase noted

BOTTOM LINE

The demo is the right size for what it is: proof that the loop is fun and legible. Turning it into a product is mostly the unglamorous half — a real homework source, a gate that holds, accounts, instrumentation — pointed at **one** customer (consumer parents of ADHD kids), validated by a small honest pilot before any big spend. Build the novel part (the gate + economy) yourself; partner for the commodity part (the curriculum). And don't claim it works until the pilot — and later the study — say it does.