The cost of free. Why it matters.

Generative AI, everyday emotional support, and the economics of access

Prepared for Kenya Jackson-Saulters, Founder, My Postpartum Pal
October 29, 2026 | U.S. focus | Discussion paper

Executive summary

Postpartum support should be free or low cost because even a modest recurring fee can compete with childcare, household expenses, and reduced earnings. Keeping that barrier low requires a sustainable funding model. Free access is a commitment with an operating cost.

Generative AI may make a limited layer of responsive emotional support available at a low marginal delivery cost. The credible case is for helping adults reflect, name feelings, use grounding practices, and find a next step alongside human relationships and care. Lower message-generation costs do not establish clinical effectiveness, safety, or a lower total cost of care.

Sword Health's acquisition of Headspace provides a timely industry example. Headspace reports that the deal closed in fall 2026. Its existing AI companion, Ebb, is positioned as nonclinical support. The company says the acquisition brings no immediate changes to AI use or member pricing, and human expertise remains essential. The transaction is a strategic signal, not evidence that replacing clinicians works. [1-3]

Three findings	Implication
Human availability has a real cost	Keeping one support role staffed 24/7 across multiple shifts requires 8,760 hours/year. At assumed $30-$50 hourly wages, total team wages are $262,800-$438,000; this is not one salary.
AI inference can be inexpensive	A transparent text-only example below costs $0.03 per conversation in model usage. This is a scenario, not MPP's actual bill.
Maternal funding is not a single open pool	Published nonprofit spending and restricted government awards illustrate different funding structures, purposes, and eligibility.

Recommendation: Fund a bounded, independently evaluated support pilot with paid human oversight, protected referral pathways, and clear reporting. Assess helpfulness and appropriate connection to care, rather than treating conversation volume as success.

Scope: This brief is an analysis, not a clinical guideline, procurement quote, systematic review, or audited MPP cost model. No clinical outcomes or cost savings have been established for MPP by the sources reviewed.

What the Headspace example can tell us

The acquisition raises a practical question: can digital availability complement human expertise at a scale that more people can access? Sword's announced plans connect AI capabilities with Headspace's provider network. Headspace's current FAQs preserve an important boundary: there are no immediate acquisition-related changes to AI use, and human care remains part of the experience. Future benefits should be measured rather than assumed. [1,2]

Ebb offers conversational reflection and personalized content recommendations. Headspace explicitly distinguishes it from clinical mental-health services and states that a person does not monitor chats in real time. This is a useful category comparison for a nonclinical support product; it does not establish that another product has equivalent testing or protections. [3]

Promising findings, specific limits

A Headspace-affiliated 2026 study examined usage by 393,969 members, alongside a 482-person survey and a 15-person diary study. For Ebb's newer version, 93.5% of submitted conversation ratings were positive; 50.8% of users completed at least two sessions within seven days. These are feedback and engagement measures. They do not mean 93.5% of all users improved, and they do not establish a treatment effect. [4]

The diary and survey work suggested interest in in-the-moment support alongside human care. Selection effects, company affiliation, and the small qualitative sample limit how broadly those findings should be generalized. The study is useful evidence about use and acceptability, not a substitute for an independent outcomes evaluation. [4]

Dartmouth's Therabot research provides a different kind of evidence: a randomized trial with 106 participants receiving access and 104 waitlist controls found symptom improvements. The system was purpose-built with mental-health experts, and the study team reviewed conversations and could intervene. It was not a head-to-head trial against a therapist and was not a postpartum-specific validation of MPP. [5]

The postpartum application remains a hypothesis

A mother might use a bounded chat to organize what she wants to tell a provider, identify a manageable next action, or work through a reflective prompt. Whether this reliably improves her experience, increases appropriate help-seeking, or reduces burden requires direct evaluation. An empathetic tone is not evidence of understanding, accurate assessment, or safe handling of every disclosure.

Defensible claim: Purpose-built AI warrants evaluation as an additional source of everyday support.

Unsupported claim: An inexpensive chatbot can replace maternal mental-health care or has already proven equivalent outcomes.

Availability is part of affordability

My Postpartum Pal's founder reports that the app is busiest between 11 p.m. and midnight. This is a product-level observation, not a population estimate. The reporting period, time zone, activity definition, and user counts were not supplied; it should not be interpreted as proof that every user is seeking help at local midnight or that human alternatives are unavailable.

Even with those limits, this observation raises a useful design question: can support be available during the hours people actually use it? A service can have a low price and still be difficult to access if its schedule does not fit a mother's available time. Responsive digital tools may help address that timing constraint, while existing human services remain essential.

Why the price paid by mothers matters

For a mother managing childcare costs, unpaid leave, or reduced working hours, an additional subscription can be difficult to sustain. The $99.99 monthly coaching benchmark in this brief is one published offering, not a market average or a claim that all human support costs that amount. Free human resources also exist. The goal is to preserve meaningful choice without making ability to pay a condition of basic support. [6,11]

Free trials and paid premium tiers can support a business, but they raise a continuity question: what help remains available when a trial ends or a family cannot renew? This brief does not quantify how common those models are across postpartum apps. For MPP, the proposed commitment is ongoing access to core supportive chat, with any optional paid features clearly separated from that access.

Who carries the cost of keeping access open?

Organizations can combine grants, sponsorship, institutional purchasing, donations, and appropriately designed technology to sustain services. Increasing funding and improving delivery efficiency are compatible strategies. AI may lower the marginal expense of some reflective interactions, but it cannot be assumed to replace the judgment or availability of a human professional.

A fundable proposition is to underwrite dependable access for a defined group of mothers, including the hours they use it, while paying for evaluation, oversight, maintenance, and referral pathways. That supports the user-facing promise of free access without pretending the underlying service is free to operate.

Core argument: Mothers should be able to access support at a price they can afford and a time they can use it. Sustaining that access requires funding; carefully evaluated AI may help that funding support more everyday interactions.

What does support actually cost?

Prices below are public benchmarks or explicit calculations, not equivalent services. Retail coaching fees include a business model; staffing scenarios count labor only; model prices count generated text only. Comparing these as a single percentage saving would be misleading.

Benchmark	Price / calculated cost	What it includes
Headspace human coaching [6]	$99.99/month	Three scheduled 30-minute text sessions. Promotional first-month pricing is excluded.
One support role staffed 24/7 across shifts	$262,800-$438,000/year in team wages	24 hours x 365 days x assumed $30-$50/hour. Multiple workers; one on duty at a time.
500 users: limited human support	$16,250-$27,083/month	500 x 15 minutes/week x 52/12 x assumed $30-$50/hour. Direct contact time only.
Claude Haiku 4.5 API [7]	$1 per million input tokens; $5 per million output tokens	Published text-processing rates. A price benchmark, not a clinical suitability recommendation.

At $30/hour, one full-time employee working 40 hours/week for 52 weeks earns $62,400; four earn $249,600 and cover 160 of the 168 hours in a week. Four people alone cannot cover 24/7. Annual coverage requires 8,760 / 2,080 = approximately 4.21 full-time equivalents before leave, breaks, training, and absences. This is a coverage calculation, not an exact headcount or one person's salary. The 52-week convention is 364 days; the annual estimate uses 365 days.

The $30-$50 hourly wages are illustrative assumptions, not verified pay rates for a specified credential. Employer payroll taxes, benefits, supervision, recruiting, and relief coverage are excluded. This models paid nonclinical support availability, not a therapy service or guaranteed response capacity. Concurrent demand may require more staff; scheduled or shared services may cost less.

A reproducible AI conversation scenario

Assume one text conversation uses **20,000 billed input tokens** across all requests, including repeatedly sent instructions and history, and **2,000 billed output tokens**. At the published rates above: (20,000/1,000,000 x $1) + (2,000/1,000,000 x $5) = **$0.03 per conversation**. A token is a unit of processed text, not a full message. [7; calculation]

Monthly scenario	Model usage only
500 active users x 10 conversations = 5,000 conversations	$150/month
Same users and frequency, with 5x the assumed token volume	$750/month

These calculations exclude hosting, database and platform fees, extra safety-model calls, voice, retries, monitoring, security, support, development, and taxes. They assume no caching or batch discounts. Actual consumption depends on conversation length and architecture. The model is a published legacy offering; price and availability must be rechecked before procurement. [7]

MPP-specific gap: The app's current model, Base44 billing, integration markups, token logs, and operating invoices were not audited. These figures must not be presented as MPP's actual cost or a guaranteed delivery price.

Free access depends on funded capacity

Maternal-support organizations pay for coordination, skilled staff, volunteer training, language access, outreach, technology, and continuity. A family-facing price of zero does not eliminate these inputs. Existing free human services are part of the support landscape; AI should not be justified by claiming no human help exists.

Published example	What the number means
Postpartum Support International, 2024 Form 990 [8]	$10,112,772 in total program-service expenses; $3,528,734 attributed to operating the National Maternal Mental Health Hotline. These are reported expenditures, not available grants or a per-user price.
HHS budget document [9]	$11 million FY2024 final funding for the maternal mental-health/substance-use screening and treatment program; $11 million requested for FY2026. A request is not an enacted appropriation.
HRSA-26-102 listing [10]	$1.29 million listed program funding and two expected awards; a $500,000 award maximum and matching requirement are also listed. These are opportunity figures, not confirmed disbursements.

The 2026 HRSA opportunity supports provider teleconsultation, training, and care coordination. Eligibility is limited to states, federally recognized tribes, and tribal organizations. Its application window has closed. It should not be described as a grant a standalone app company can directly claim. The listing's headline total and maximum-award fields should not be forced into a per-award calculation without the full award terms. [10]

PSI's 2024 impact report describes reliance on funders to sustain and expand services, and reports 1,829 English-language support groups and 10,011 HelpLine calls, texts, and online submissions. Those activity counts represent different services and cannot be divided into its overall budget to produce a reliable cost per mother. [11]

Why funding can still be difficult

These examples support an eligibility and sustainability argument, not a blanket claim that all organizations are underfunded. Funds may be tied to a program, applicant type, period, matching requirement, or set of deliverables. A large national award is not necessarily flexible money for local staffing or an emerging product. Annual fundraising must also support continuity after a pilot ends.

For MPP, a more plausible route to explore is a paid partnership with an eligible or established organization, employer, health system, or sponsor. Any subcontract, procurement, privacy, and grant-allowability terms would need review. No source here establishes MPP's eligibility, a funding commitment, or a national total for all maternal-support spending.

Funding principle: Preserve human services and finance responsible digital capacity alongside them. Low marginal message costs can create room in a budget; they do not remove the need for that budget.

Pay for access, accountability, and learning

A practical next step is a time-limited pilot for adults seeking everyday postpartum support, with an explicit scope and a funded evaluation. MPP's founder describes generative AI as confined to chat. This brief treats that as founder-reported positioning, not an independent technical audit. Public claims should match the deployed system and its actual data practices.

Illustrative 90-day pilot budget

Planning example for 500 active users and 10 text conversations per user monthly. Rates and work allocations below are assumptions for discussion, not quotations or a claim that this budget is sufficient for every risk profile.

Line item	Assumption	90-day total
Model usage	5,000 conversations/month x $0.03 x 3	$450
Platform, hosting, monitoring	$300/month x 3	$900
Qualified safety/content review	40 hours x $150	$6,000
Engineering, security, privacy work	40 hours x $100	$4,000
Partner onboarding and user support	60 hours x $40	$2,400
Evaluation and reporting	30 hours x $75	$2,250
Subtotal		$16,000
Contingency	20% of subtotal	$3,200
Illustrative funding target		$19,200

This totals $38.40 per participating user for 90 days if all 500 participate. If only 250 participate and the same budget is spent, it is $76.80. Model usage is 2.8% of the pre-contingency budget. The proposed review hours are not live chat coverage. Product development beyond these hours, legal work, insurance, integrations, recruitment incentives, and staffed clinical care may require additional funding.

What a sponsor should receive

Clear boundaries: Plain-language AI disclosure; no diagnosis or treatment claims; accurate notice about whether anyone monitors chats; accessible human-resource options; data minimization and retention controls. Postpartum risk scenarios should be assessed by qualified experts before and during a pilot.

Meaningful evaluation: Track opt-in helpfulness, ability to identify a next step, referral follow-through where consent permits, harmful or inappropriate responses, and performance across relevant groups. Use aggregate reporting without exposing private chats. Engagement alone is insufficient.

A continuation decision: Agree on escalation, adverse-event review, and pause criteria. Report costs alongside safety results. Claims of fewer clinical visits or lower spending require direct evidence.

Evidence notes and references

Public sources reviewed October 29, 2026. Figures retain their original reporting periods. Calculations use U.S. dollars and 52 weeks/year. This is a targeted review of primary organizational publications, official funding listings, published research, and vendor prices; it is not exhaustive. Published financial statements and grant listings describe different scopes and must not be added together as a sector total.

[1] Headspace. Headspace and Sword: A New Era for Mental Health. Acquisition closed fall 2026.
https://www.headspace.com/articles/headspace-and-sword-a-new-era-for-mental-health

[2] Headspace Help Center. Will Headspace use more AI?; subscription/pricing FAQ.
https://help.headspace.com/hc/en-us/sections/55913396898715-Sword-Acquisition

[3] Headspace. Meet Ebb: intended use, human-monitoring limits, and privacy information.
https://www.headspace.com/ai-mental-health-companion

[4] Callahan et al. Real-World Use of a Mental Health AI Companion: Multiple Methods Study. JMIR Formative Research, 2026;10:e86904. Headspace-affiliated authors.
https://formative.jmir.org/2026/1/e86904/

[5] Dartmouth. First Therapy Chatbot Trial Yields Mental Health Benefits, March 27, 2025. Institutional account of the NEJM AI Therabot trial.
https://ai.dartmouth.edu/news/2025/03/first-therapy-chatbot-trial-yields-mental-health-benefits

[6] Headspace. Mental health coaching enrollment: three 30-minute sessions, $99.99/month.
https://help.headspace.com/hc/en-us/articles/25083715994523-How-Do-I-Enroll-in-Headspace-Mental-Health-Coaching

[7] Anthropic. Claude Haiku 4.5 model documentation: $1 input / $5 output per million tokens; legacy model.
https://platform.claude.com/docs/en/models/haiku-4-5/overview

[8] Postpartum Support International. 2024 Form 990, Part III, lines 4b and 4e. EIN 77-0196208; filed November 17, 2025; IRS object 202503219349326165. Public filing index linked for retrieval.
https://990finder.com/770196208/

[9] HHS. FY2026 Administration for a Healthy America Congressional Justification: maternal mental-health screening and treatment funding history and request.
https://www.hhs.gov/sites/default/files/fy-2026-aha-cj.pdf

[10] HRSA / Simpler.Grants.gov. HRSA-26-102, posted June 17, 2026. Opportunity funding, eligibility, and matching fields; archived.
https://simpler.grants.gov/es/opportunity/64847f01-ba43-43b1-962f-5e1dc248327d

[11] Postpartum Support International. 2024 Impact Report: service activity and funding dependence.
https://postpartum.net/about-psi/2024-impact-report/

Interpretive conclusion: a lower unit cost can support a case for testing access, but an acquisition, positive ratings, and low token prices do not collectively prove clinical safety, efficacy, or cost-effectiveness. Those remain product-specific questions.