Appendix

A Proofs

Proof of Lemma 1 (Variance of the Beta-Binomial posterior mean)

Under the Beta-Binomial structure, the posterior mean is $\mathbb{E}(\theta \mid s) = (\alpha + s)/(\alpha + \beta + n)$, which is linear in s. Therefore:

$$\mathrm{Var}_s[\mathbb{E}(\theta \mid s)] = \frac{\mathrm{Var}[s]}{(\alpha+\beta+n)^2} \tag{A1}$$

The marginal variance of s under the Beta-Binomial is $\mathrm{Var}[s] = n\mu_\theta(1-\mu_\theta)[1+(n-1)\sigma_\theta^2/(\mu_\theta(1-\mu_\theta))]$. Substituting and simplifying using the Beta moment condition $\sigma_\theta^2 = \alpha\beta/[(\alpha+\beta)^2(\alpha+\beta+1)]$:

$$\mathrm{Var}_s[\mathbb{E}(\theta \mid s)] = \frac{n\alpha\beta(\alpha+\beta+n)}{(\alpha+\beta)^2(\alpha+\beta+1)(\alpha+\beta+n)^2} = \sigma_\theta^2 \cdot \frac{n}{\alpha+\beta+n} \tag{A2}$$

Since $\omega_\theta \equiv (\alpha+\beta)/(\alpha+\beta+n)$, we have $n/(\alpha+\beta+n) = 1-\omega_\theta$, giving the result. □

Proof of Proposition 1 (Stage 2 optimal choices)

Step 1: Optimal response conditional on signal. After observing s, the parent maximizes $\mathbb{E}[u(x,\theta) \mid s]$ over $x > 0$:

$$\max_{x>0} \; \hat{\mu}_\gamma \, \mathbb{E}(\theta \mid s) \, \ln x - \frac{1}{2\delta}(\ln x - \ln x^R)^2 \tag{A3}$$

The first-order condition with respect to $\ln x$ yields:

$$\ln x^*(s) = \ln x^R + \delta\hat{\mu}_\gamma \, \mathbb{E}(\theta \mid s) \tag{A4}$$

The posterior mean under the Beta-Binomial structure is $\mathbb{E}(\theta \mid s) = \omega_\theta\mu_\theta + (1-\omega_\theta)s/n$, substituting which into (A4) gives part (ii) of the proposition.

Step 2: Expected indirect utility. Substituting (A4) into the payoff and defining $Q(s)$ as the resulting value conditional on s:

$$\begin{aligned} Q(s) &= \hat{\mu}_\gamma \, \mathbb{E}(\theta \mid s) \left(\ln x^R + \delta\hat{\mu}_\gamma \, \mathbb{E}(\theta \mid s)\right) - \frac{1}{2\delta}\left(\delta\hat{\mu}_\gamma \, \mathbb{E}(\theta \mid s)\right)^2 \\ &= \hat{\mu}_\gamma \, \mathbb{E}(\theta \mid s) \, \ln x^R + \frac{\delta}{2}\hat{\mu}_\gamma^2 \, [\mathbb{E}(\theta \mid s)]^2 \end{aligned} \tag{A5}$$

Taking expectations, applying the law of iterated expectations, and invoking Lemma 1:

$$\mathbb{E}_s[Q(s)] = \hat{\mu}_\gamma \, \mu_\theta \, \ln x^R + \frac{\delta}{2}\hat{\mu}_\gamma^2 \, \mu_\theta^2 + \frac{\delta}{2}\hat{\mu}_\gamma^2 \, \sigma_\theta^2(1-\omega_\theta) \tag{A6}$$

Step 3: Optimal information acquisition. Subtracting the information cost $\kappa(b,\rho) \cdot C(\omega_\theta) = -\kappa(b,\rho) \ln \omega_\theta$ from (A6), the Stage 2 objective is:

$$\max_{\omega_\theta \in (0,1]} \quad \hat{\mu}_\gamma \, \mu_\theta \, \ln x^R + \frac{\delta}{2} \hat{\mu}_\gamma^2 \, \mu_\theta^2 + \frac{\delta}{2} \hat{\mu}_\gamma^2 \, \sigma_\theta^2 (1 - \omega_\theta) + \kappa(b,\rho) \ln \omega_\theta \tag{A7}$$

The first-order condition with respect to ω_θ is:

$$-\frac{\delta}{2} \hat{\mu}_\gamma^2 \, \sigma_\theta^2 + \frac{\kappa(b,\rho)}{\omega_\theta} = 0 \tag{A8}$$

Solving gives the interior solution $\omega_\theta^* = 2\kappa(b,\rho)/(\delta \hat{\mu}_\gamma^2 \sigma_\theta^2)$. The second-order condition holds since $\partial^2/\partial \omega_\theta^2 = -\kappa(b,\rho)/\omega_\theta^2 < 0$. Feasibility requires $\omega_\theta^* \leqslant 1$, i.e. $\delta \hat{\mu}_\gamma^2 \sigma_\theta^2 \geqslant 2\kappa(b,\rho)$, equivalently $\hat{\mu}_\gamma \geqslant \bar{\mu}_\gamma$; otherwise the corner $\omega_\theta^* = 1$ is optimal. □

Proof of Proposition 2 (Comparative statics of ω_θ^*)

All derivatives are taken with respect to $\omega_\theta^* = 2\kappa(b,\rho)/(\delta \hat{\mu}_\gamma^2 \, \sigma_\theta^2)$ in the interior regime.

(i)

$$\frac{\partial \omega_\theta^*}{\partial \hat{\mu}_\gamma} = -\frac{4\kappa(b,\rho)}{\delta \hat{\mu}_\gamma^3 \, \sigma_\theta^2} < 0 \tag{A9}$$

The effect is proportional to $\hat{\mu}_\gamma^{-3}$, so it is strongest at low beliefs.

(ii)

$$\frac{\partial \omega_\theta^*}{\partial b} = -\frac{2\kappa_0}{\delta \hat{\mu}_\gamma^2 \, \sigma_\theta^2 \, b^2 (1 + \rho)} < 0 \tag{A10}$$

(iii)

$$\frac{\partial \omega_\theta^*}{\partial \rho} = -\frac{2\kappa_0}{\delta \hat{\mu}_\gamma^2 \, \sigma_\theta^2 \, b (1 + \rho)^2} < 0 \tag{A11}$$

□

Proof of Proposition 3 (Properties of G)

(i) Continuity and differentiability. At $\hat{\mu}_\gamma = \bar{\mu}_\gamma$, we have $\delta \bar{\mu}_\gamma^2 \sigma_\theta^2 = 2\kappa(b,\rho)$ by definition of $\bar{\mu}_\gamma$, so $\ln(\delta \bar{\mu}_\gamma^2 \sigma_\theta^2/(2\kappa(b,\rho))) = 0$. The interior expression gives:

$$G\big|_{\hat{\mu}_\gamma = \bar{\mu}_\gamma} = \psi \kappa(b,\rho)(1 + 0) = \psi \kappa(b,\rho)$$

The corner expression gives:

$$G\big|_{\hat{\mu}_\gamma = \bar{\mu}_\gamma} = \frac{\psi \delta}{2} \bar{\mu}_\gamma^2 \sigma_\theta^2 = \frac{\psi \delta}{2} \cdot \frac{2\kappa(b,\rho)}{\delta} = \psi \kappa(b,\rho)$$

Both expressions agree, so G is continuous. For differentiability, take $\partial G/\partial\hat{\mu}_\gamma$ from each side. In the interior:

$$\frac{\partial G}{\partial\hat{\mu}_\gamma} = \frac{2\psi\kappa(b,\rho)}{\hat{\mu}_\gamma}$$

evaluated at $\hat{\mu}_\gamma = \bar{\mu}_\gamma$ gives $2\psi\kappa(b,\rho)/\bar{\mu}_\gamma$. In the corner:

$$\frac{\partial G}{\partial\hat{\mu}_\gamma} = \psi\delta\hat{\mu}_\gamma\sigma_\theta^2$$

evaluated at $\hat{\mu}_\gamma = \bar{\mu}_\gamma$ gives $\psi\delta\bar{\mu}_\gamma\sigma_\theta^2 = \psi\delta\bar{\mu}_\gamma \cdot 2\kappa(b,\rho)/(\delta\bar{\mu}_\gamma^2) = 2\psi\kappa(b,\rho)/\bar{\mu}_\gamma$. Both one-sided derivatives agree, so G is differentiable at $\bar{\mu}_\gamma$.

(ii) Increasing in $\hat{\mu}_\gamma$. In the interior regime, $\partial G/\partial\hat{\mu}_\gamma = 2\psi\kappa(b,\rho)/\hat{\mu}_\gamma > 0$. In the corner regime, $\partial G/\partial\hat{\mu}_\gamma = \psi\delta\hat{\mu}_\gamma\sigma_\theta^2 > 0$. Both are positive in their respective regimes.

(iii) Decreasing in b and ρ. In the interior regime, $G = \psi\kappa(b,\rho)(1+\ln(\delta\hat{\mu}_\gamma^2\sigma_\theta^2/(2\kappa(b,\rho))))$. Since $\hat{\mu}_\gamma > \bar{\mu}_\gamma$ implies $\delta\hat{\mu}_\gamma^2\sigma_\theta^2 > 2\kappa(b,\rho)$, the log term is positive, so the factor $(1+\ln(\cdot)) > 1 > 0$. Since $\kappa(b,\rho) = \kappa_0/[b(1+\rho)]$ is decreasing in both b and ρ, G is decreasing in both. In the corner regime, $G = \frac{\psi\delta}{2}\hat{\mu}_\gamma^2\sigma_\theta^2$ does not depend on b or ρ.

(iv) Larger shortfall in the corner. At $\hat{\mu}_\gamma = \bar{\mu}_\gamma$, both regimes give $G = \psi\kappa(b,\rho)$ as shown in part (i). As $\hat{\mu}_\gamma$ falls below $\bar{\mu}_\gamma$, the corner shortfall $\frac{\delta}{2}\hat{\mu}_\gamma^2\sigma_\theta^2$ decreases at rate $\delta\hat{\mu}_\gamma\sigma_\theta^2 = 2\kappa(b,\rho)/\bar{\mu}_\gamma > 0$ evaluated at the threshold. The two are therefore tangent at $\bar{\mu}_\gamma$, consistent with differentiability in part (i). □

Proof of Proposition 4 (Optimal working belief)

We prove the proposition in three steps, corresponding to the three results established below.

Result 1: Corner regime solution ($\mu_{\gamma,0} \leqslant \bar{\mu}_{\gamma,0}$). In the corner regime, $\hat{\mu}_\gamma \leqslant \bar{\mu}_\gamma$, the Stage 2 value function and psychological cost reduce to:

$$V_2(\hat{\mu}_\gamma, b, \rho) = \hat{\mu}_\gamma\mu_\theta \ln x^R + \frac{\delta}{2}\hat{\mu}_\gamma^2\mu_\theta^2$$
$$G(\hat{\mu}_\gamma, b, \rho) = \frac{\psi\delta}{2}\hat{\mu}_\gamma^2\sigma_\theta^2$$

The Stage 1 objective in the corner regime is:

$$V_1^{\text{corner}}(\hat{\mu}_\gamma) = \hat{\mu}_\gamma\mu_\theta \ln x^R + \frac{\delta}{2}\hat{\mu}_\gamma^2\mu_\theta^2 - \frac{\lambda_\gamma}{2\sigma_{\gamma,0}^2}(\hat{\mu}_\gamma - \mu_{\gamma,0})^2 - \frac{\psi\delta}{2}\hat{\mu}_\gamma^2\sigma_\theta^2 \tag{A12}$$

The second derivative is:

$$\frac{\partial^2 V_1^{\text{corner}}}{\partial\hat{\mu}_\gamma^2} = \delta\mu_\theta^2 - \frac{\lambda_\gamma}{\sigma_{\gamma,0}^2} - \psi\delta\sigma_\theta^2 = -\left[A_\delta + (1+\psi)\delta\sigma_\theta^2\right] = -P < 0$$

under Assumption 1, so V_1^{corner} is strictly concave. The first-order condition is:

$$\mu_\theta \ln x^R + \delta\hat{\mu}_\gamma \mu_\theta^2 - \frac{\lambda_\gamma}{\sigma_{\gamma,0}^2}(\hat{\mu}_\gamma - \mu_{\gamma,0}) - \psi\delta\hat{\mu}_\gamma\sigma_\theta^2 = 0$$

Collecting terms in $\hat{\mu}_\gamma$ and using $B \equiv \mu_\theta \ln x^R + \lambda_\gamma \phi_0$:

$$B = \hat{\mu}_\gamma \left(\frac{\lambda_\gamma}{\sigma_{\gamma,0}^2} - \delta\mu_\theta^2 + \psi\delta\sigma_\theta^2 \right) = \hat{\mu}_\gamma P$$

The unique solution is $\hat{\mu}_\gamma^* = B/P$, which is well-defined since $P > 0$ under Assumption 1. □

Result 2: Interior regime solution ($\mu_{\gamma,0} > \bar{\mu}_{\gamma,0}$). In the interior regime, $\hat{\mu}_\gamma > \bar{\mu}_\gamma$, differentiating V_2 and G with respect to $\hat{\mu}_\gamma$:

$$\frac{dV_2}{d\hat{\mu}_\gamma} = \mu_\theta \ln x^R + \delta\hat{\mu}_\gamma(\mu_\theta^2 + \sigma_\theta^2) - \frac{2\kappa(b,\rho)}{\hat{\mu}_\gamma}$$
$$\frac{dG}{d\hat{\mu}_\gamma} = \frac{2\psi\kappa(b,\rho)}{\hat{\mu}_\gamma}$$

The first derivative of V_2 follows from direct differentiation of the indirect utility (17), verified by the envelope theorem: since ω_θ^* is an interior solution to the Stage 2 problem, $\partial V_2/\partial\omega_\theta\big|_{\omega_\theta^*} = 0$, so the indirect effect of $\hat{\mu}_\gamma$ through ω_θ^* vanishes. The Stage 1 first-order condition is:

$$\mu_\theta \ln x^R + \delta\hat{\mu}_\gamma(\mu_\theta^2 + \sigma_\theta^2) - \frac{2(1+\psi)\kappa(b,\rho)}{\hat{\mu}_\gamma} - \frac{\lambda_\gamma}{\sigma_{\gamma,0}^2}(\hat{\mu}_\gamma - \mu_{\gamma,0}) = 0$$

Multiplying through by $\hat{\mu}_\gamma > 0$ and rearranging using A_δ and B:

$$A_\delta \hat{\mu}_\gamma^2 - B\hat{\mu}_\gamma + 2(1+\psi)\kappa(b,\rho) = 0 \tag{A13}$$

The second derivative is:

$$\frac{d^2 V_1^{\text{interior}}}{d\hat{\mu}_\gamma^2} = -A_\delta + \frac{2(1+\psi)\kappa(b,\rho)}{\hat{\mu}_\gamma^2}$$

The two roots of (A13) are $\hat{\mu}_\gamma^\pm = (B \pm \sqrt{\Delta})/(2A_\delta)$ where $\Delta \equiv B^2 - 8A_\delta(1+\psi)\kappa(b,\rho)$. Substituting $2(1+\psi)\kappa(b,\rho) = B\hat{\mu}_\gamma^\pm - A_\delta\hat{\mu}_\gamma^{\pm 2}$ from (A13) into the second derivative:

$$\left.\frac{d^2 V_1^{\text{interior}}}{d\hat{\mu}_\gamma^2}\right|_{\hat{\mu}_\gamma^\pm} = \frac{B}{\hat{\mu}_\gamma^\pm} - 2A_\delta = 2A_\delta \cdot \frac{\mp\sqrt{\Delta}}{B \pm \sqrt{\Delta}}$$

Since $A_\delta > 0$ under Assumption 1 and $B > \sqrt{\Delta} > 0$ (the latter follows from $8A_\delta(1+\psi)\kappa > 0$):

- At $\hat{\mu}_\gamma^+$: second derivative < 0 — local maximum ✓
- At $\hat{\mu}_\gamma^-$: second derivative > 0 — local minimum ✓

The unique maximum is $\hat{\mu}_\gamma^* = (B + \sqrt{\Delta})/(2A_\delta)$. □

Result 3: Continuity and the regime threshold. We find $\bar{\mu}_{\gamma,0}$ by setting $\hat{\mu}_\gamma^{*,\text{corner}} = \hat{\mu}_\gamma^{*,\text{interior}}$:

$$\frac{B}{P} = \frac{B + \sqrt{\Delta}}{2A_\delta}$$

Rearranging:

$$BQ = P\sqrt{\Delta}$$

where $Q \equiv A_\delta - (1+\psi)\delta\sigma_\theta^2 > 0$ under Assumption 1. Squaring and substituting $\Delta = B^2 - 8A_\delta(1+\psi)\kappa(b,\rho)$:

$$B^2Q^2 = P^2\left(B^2 - 8A_\delta(1+\psi)\kappa(b,\rho)\right)$$

Using $P^2 - Q^2 = (P+Q)(P-Q) = 2A_\delta \cdot 2(1+\psi)\delta\sigma_\theta^2 = 4A_\delta(1+\psi)\delta\sigma_\theta^2$:

$$\delta\sigma_\theta^2 B^2 = 2\kappa(b,\rho)P^2$$

Since $\bar{\mu}_\gamma^2 = 2\kappa(b,\rho)/(\delta\sigma_\theta^2)$, this gives $B = \bar{\mu}_\gamma P$. Solving for ϕ_0 yields $\bar{\phi}$ and $\bar{\mu}_{\gamma,0}$ as stated in the proposition.

Substituting $B = \bar{\mu}_\gamma P$ into the corner solution gives $\hat{\mu}_\gamma^{*,\text{corner}} = \bar{\mu}_\gamma$. For the interior solution, using $8A_\delta(1+\psi)\kappa = 4A_\delta(1+\psi)\delta\sigma_\theta^2\bar{\mu}_\gamma^2$ and the identity $P^2 - 4A_\delta(1+\psi)\delta\sigma_\theta^2 = Q^2$:

$$\sqrt{\Delta}\big|_{B=\bar{\mu}_\gamma P} = \bar{\mu}_\gamma Q$$

Therefore:

$$\hat{\mu}_\gamma^{*,\text{interior}} = \frac{\bar{\mu}_\gamma P + \bar{\mu}_\gamma Q}{2A_\delta} = \frac{\bar{\mu}_\gamma(P+Q)}{2A_\delta} = \frac{2A_\delta\bar{\mu}_\gamma}{2A_\delta} = \bar{\mu}_\gamma$$

Both solutions agree at $\mu_{\gamma,0} = \bar{\mu}_{\gamma,0}$, so $\hat{\mu}_\gamma^*$ is continuous everywhere. Since Δ is strictly increasing in ϕ_0 and equals zero at $\phi_0 = \bar{\phi}$:

$$\Delta \geqslant 0 \iff \phi_0 \geqslant \bar{\phi} \iff \mu_{\gamma,0} \geqslant \bar{\mu}_{\gamma,0}$$

confirming that the interior solution is well-defined precisely when $\mu_{\gamma,0} \geqslant \bar{\mu}_{\gamma,0}$. □

Proof of Proposition 5 (Direction of distortion)

We prove the proposition in two steps, corresponding to Results 4 and 5 established below.

Result 4: Direction of distortion in the corner regime. In the corner regime, the optimal working belief is $\hat{\mu}_\gamma^* = B/P$ from Proposition 4. We find the threshold $\tilde{\mu}_{\gamma,0}^{\text{corner}}$ by setting $\hat{\mu}_\gamma^* = \mu_{\gamma,0}$:

$$\frac{\mu_\theta \ln x^R + \lambda_\gamma \mu_{\gamma,0}/\sigma_{\gamma,0}^2}{P} = \mu_{\gamma,0}$$

Multiplying through by P and collecting terms in $\mu_{\gamma,0}$:

$$\mu_\theta \ln x^R = \mu_{\gamma,0}\left(P - \frac{\lambda_\gamma}{\sigma^2_{\gamma,0}}\right)$$

Computing $P - \lambda_\gamma/\sigma^2_{\gamma,0}$ using $P = A_\delta + (1+\psi)\delta\sigma^2_\theta$ and $A_\delta = \lambda_\gamma/\sigma^2_{\gamma,0} - \delta(\mu^2_\theta + \sigma^2_\theta)$:

$$P - \frac{\lambda_\gamma}{\sigma^2_{\gamma,0}} = \delta\left(\psi\sigma^2_\theta - \mu^2_\theta\right) > 0$$

under Assumption 2. Solving for $\mu_{\gamma,0}$ gives $\tilde{\mu}^{\text{corner}}_{\gamma,0}$ as stated in the proposition.

We now show the slope of the corner solution with respect to $\mu_{\gamma,0}$ is less than one. Differentiating $\hat{\mu}^*_\gamma = B/P$:

$$\frac{\partial\hat{\mu}^*_\gamma}{\partial\mu_{\gamma,0}} = \frac{\lambda_\gamma/\sigma^2_{\gamma,0}}{P}$$

For this to be less than one we need $\lambda_\gamma/\sigma^2_{\gamma,0} < P$, which reduces to $\delta(\psi\sigma^2_\theta - \mu^2_\theta) > 0$, holding under Assumption 2. Therefore the corner solution is a strictly increasing function of $\mu_{\gamma,0}$ with slope less than one, crossing the 45-degree line exactly once at $\tilde{\mu}^{\text{corner}}_{\gamma,0}$:

- $\mu_{\gamma,0} < \tilde{\mu}^{\text{corner}}_{\gamma,0}$: $\hat{\mu}^*_\gamma > \mu_{\gamma,0}$ — upward distortion
- $\mu_{\gamma,0} = \tilde{\mu}^{\text{corner}}_{\gamma,0}$: $\hat{\mu}^*_\gamma = \mu_{\gamma,0}$ — no distortion
- $\tilde{\mu}^{\text{corner}}_{\gamma,0} < \mu_{\gamma,0} \leqslant \bar{\mu}_{\gamma,0}$: $\hat{\mu}^*_\gamma < \mu_{\gamma,0}$ — downward distortion

Since G is independent of b and ρ in the corner regime, $\tilde{\mu}^{\text{corner}}_{\gamma,0}$ is also independent of b and ρ. □

Result 5: Direction of distortion in the interior regime. We find $\tilde{\mu}^{\text{interior}}_{\gamma,0}$ by setting $\hat{\mu}^*_\gamma = \mu_{\gamma,0}$ in the interior solution and isolating the square root:

$$\sqrt{B^2 - 8A_\delta(1+\psi)\kappa(b,\rho)} = 2A_\delta\mu_{\gamma,0} - B$$

Squaring both sides, canceling B^2, and dividing by $4A_\delta > 0$:

$$-2(1+\psi)\kappa(b,\rho) = A_\delta\mu^2_{\gamma,0} - \mu_{\gamma,0}B$$

Substituting $B = \mu_\theta \ln x^R + \lambda_\gamma\mu_{\gamma,0}/\sigma^2_{\gamma,0}$ and using $A_\delta - \lambda_\gamma/\sigma^2_{\gamma,0} = -\delta(\mu^2_\theta + \sigma^2_\theta)$:

$$\delta(\mu^2_\theta + \sigma^2_\theta)\mu^2_{\gamma,0} + \mu_\theta \ln x^R \cdot \mu_{\gamma,0} - 2(1+\psi)\kappa(b,\rho) = 0 \quad \text{(A14)}$$

The discriminant $\Delta_\mu \equiv (\mu_\theta \ln x^R)^2 + 8(1+\psi)\kappa(b,\rho)\delta(\mu^2_\theta + \sigma^2_\theta) > 0$ and the product of the roots is $-2(1+\psi)\kappa/[\delta(\mu^2_\theta + \sigma^2_\theta)] < 0$, so one root is positive and one is negative. The unique positive root is $\tilde{\mu}^{\text{interior}}_{\gamma,0}$ as stated in the proposition.

Establishing $\tilde{\mu}_{\gamma,0}^{\textbf{interior}} > \bar{\mu}_{\gamma,0}$. Define $f(\mu_{\gamma,0}) \equiv \delta(\mu_\theta^2 + \sigma_\theta^2)\mu_{\gamma,0}^2 + \mu_\theta \ln x^R \cdot \mu_{\gamma,0} - 2(1+\psi)\kappa(b,\rho)$. Since the leading coefficient is positive, f is an upward-opening parabola with $f(\mu_{\gamma,0}) < 0$ for $\mu_{\gamma,0} \in (0, \tilde{\mu}_{\gamma,0}^{\text{interior}})$. It therefore suffices to show $f(\bar{\mu}_{\gamma,0}) < 0$.

From Proposition 4, the interior solution at $\mu_{\gamma,0} = \bar{\mu}_{\gamma,0}$ delivers $\hat{\mu}_\gamma^* = \bar{\mu}_\gamma$. Under Assumption 3, $\bar{\mu}_\gamma > \tilde{\mu}_{\gamma,0}^{\text{corner}} = \mu_\theta \ln x^R / [\delta(\psi\sigma_\theta^2 - \mu_\theta^2)]$, which is equivalent to:

$$\mu_\theta \ln x^R < \bar{\mu}_\gamma \delta(\psi\sigma_\theta^2 - \mu_\theta^2) = \bar{\mu}_\gamma \left(P - \frac{\lambda_\gamma}{\sigma_{\gamma,0}^2} \right)$$

Rearranging:

$$\bar{\mu}_\gamma < \frac{\sigma_{\gamma,0}^2}{\lambda_\gamma} \left[\bar{\mu}_\gamma P - \mu_\theta \ln x^R \right] = \bar{\mu}_{\gamma,0}$$

Therefore $\hat{\mu}_\gamma^*|_{\mu_{\gamma,0}=\bar{\mu}_{\gamma,0}} = \bar{\mu}_\gamma < \bar{\mu}_{\gamma,0}$, so the interior solution lies strictly below the 45-degree line at the regime boundary. This means $f(\bar{\mu}_{\gamma,0}) < 0$, confirming $\tilde{\mu}_{\gamma,0}^{\text{interior}} > \bar{\mu}_{\gamma,0}$.

Direction of distortion. Since f is an upward-opening parabola with positive root at $\tilde{\mu}_{\gamma,0}^{\text{interior}} > \bar{\mu}_{\gamma,0}$:

- $\bar{\mu}_{\gamma,0} < \mu_{\gamma,0} < \tilde{\mu}_{\gamma,0}^{\text{interior}}$: $f(\mu_{\gamma,0}) < 0$, so $\hat{\mu}_\gamma^* < \mu_{\gamma,0}$ — downward distortion
- $\mu_{\gamma,0} = \tilde{\mu}_{\gamma,0}^{\text{interior}}$: $\hat{\mu}_\gamma^* = \mu_{\gamma,0}$ — no distortion
- $\mu_{\gamma,0} > \tilde{\mu}_{\gamma,0}^{\text{interior}}$: $f(\mu_{\gamma,0}) > 0$, so $\hat{\mu}_\gamma^* > \mu_{\gamma,0}$ — upward distortion

Comparative statics of $\tilde{\mu}_{\gamma,0}^{\textbf{interior}}$. Differentiating (29) with respect to $\kappa(b,\rho)$:

$$\frac{\partial \tilde{\mu}_{\gamma,0}^{\text{interior}}}{\partial \kappa} = \frac{4(1+\psi)\delta(\mu_\theta^2 + \sigma_\theta^2)}{\delta(\mu_\theta^2 + \sigma_\theta^2)\sqrt{(\mu_\theta \ln x^R)^2 + 8(1+\psi)\kappa\delta(\mu_\theta^2 + \sigma_\theta^2)}} > 0$$

Higher attention costs expand the downward distortion region in the interior regime. Equivalently, higher bandwidth b or cue-detection expertise ρ reduce $\kappa(b,\rho)$ and shift $\tilde{\mu}_{\gamma,0}^{\text{interior}}$ leftward, expanding the upward distortion region. □

B Supplementary Theoretical Results

B.1 Properties of V_2

Proposition B1 (Properties of V_2)**.** *The function* $V_2(\hat{\mu}_\gamma, b, \rho)$ *is* C^1 *but not* C^2 *with respect to* $\hat{\mu}_\gamma$*. Specifically:*

(i) V_2 *is continuous at the threshold* $\hat{\mu}_\gamma = \bar{\mu}_\gamma$*.*

(ii) $\partial V_2/\partial\hat{\mu}_\gamma$ *is continuous at* $\bar{\mu}_\gamma$*.*

(iii) $\partial^2 V_2/\partial\hat{\mu}_\gamma^2$ *jumps by* $2\delta\sigma_\theta^2$ *at* $\hat{\mu}_\gamma = \bar{\mu}_\gamma$*, from* $\delta\mu_\theta^2$ *in the corner regime to* $\delta\mu_\theta^2 + 2\delta\sigma_\theta^2$ *in the interior regime.*

Proof. See below. □

Proof of Proposition B1 (Properties of V_2)

(i) Continuity. At $\hat{\mu}_\gamma = \bar{\mu}_\gamma$, $\omega_\theta^* = 1$, so $\delta\hat{\mu}_\gamma^2\sigma_\theta^2 - 2\kappa(b,\rho) = 0$ and $\ln(2\kappa(b,\rho)/(\delta\hat{\mu}_\gamma^2\sigma_\theta^2)) = 0$. Both regimes yield:

$$V_2\big|_{\hat{\mu}_\gamma=\bar{\mu}_\gamma} = \bar{\mu}_\gamma\,\mu_\theta\,\ln x^R + \frac{\delta}{2}\bar{\mu}_\gamma^2\,\mu_\theta^2 \tag{A15}$$

(ii) First derivative. By the envelope theorem, the indirect effect of $\hat{\mu}_\gamma$ through ω_θ^* vanishes in the interior regime. The direct derivative is:

$$\frac{\partial V_2}{\partial\hat{\mu}_\gamma} = \begin{cases} \mu_\theta\,\ln x^R + \delta\hat{\mu}_\gamma\,\mu_\theta^2 + \delta\hat{\mu}_\gamma\,\sigma_\theta^2(1-\omega_\theta^*), & \hat{\mu}_\gamma > \bar{\mu}_\gamma \\ \mu_\theta\,\ln x^R + \delta\hat{\mu}_\gamma\,\mu_\theta^2, & \hat{\mu}_\gamma < \bar{\mu}_\gamma \end{cases} \tag{A16}$$

At $\hat{\mu}_\gamma = \bar{\mu}_\gamma$, $\omega_\theta^* = 1$ and the extra term vanishes, so both one-sided limits equal $\mu_\theta\,\ln x^R + \delta\bar{\mu}_\gamma\,\mu_\theta^2$.

(iii) Second derivative. In the corner regime: $\partial^2 V_2^{\text{corner}}/\partial\hat{\mu}_\gamma^2 = \delta\mu_\theta^2$. In the interior regime, differentiating (A16) and using $d\omega_\theta^*/d\hat{\mu}_\gamma = -2\omega_\theta^*/\hat{\mu}_\gamma$:

$$\frac{\partial^2 V_2^{\text{interior}}}{\partial\hat{\mu}_\gamma^2} = \delta\mu_\theta^2 + \delta\sigma_\theta^2(1-\omega_\theta^*) + \delta\sigma_\theta^2\cdot 2\omega_\theta^* = \delta\mu_\theta^2 + \delta\sigma_\theta^2(1+\omega_\theta^*) \tag{A17}$$

Evaluating at $\hat{\mu}_\gamma = \bar{\mu}_\gamma$ (where $\omega_\theta^* = 1$), the corner gives $\delta\mu_\theta^2$ and the interior gives $\delta\mu_\theta^2 + 2\delta\sigma_\theta^2$. The second derivative jumps by $2\delta\sigma_\theta^2$, so V_2 is C^1 but not C^2. □

Proposition B2 (Complementarities of V_2)**.** *In the interior regime* $(\hat{\mu}_\gamma > \bar{\mu}_\gamma)$*, where* V_2 *is* C^2*:*

(i) $\partial^2 V_2/\partial\hat{\mu}_\gamma\,\partial b > 0$*: believed returns and bandwidth are complements.*

(ii) $\partial^2 V_2/\partial\hat{\mu}_\gamma\,\partial\rho > 0$*: believed returns and cue-detection expertise are complements.*

(iii) $\partial^2 V_2/\partial b\,\partial\rho = [\kappa(b,\rho)/b(1+\rho)](1+\ln\omega_\theta^*)$, *which is positive when* $\omega_\theta^* > 1/e$ *and negative when* $\omega_\theta^* < 1/e$.

Proof. See below. □

Proof of Proposition B2 (Complementarities of V_2)

All cross-partials are computed in the interior regime where V_2 is C^2. By the envelope theorem, indirect effects through ω_θ^* vanish.

(i) From (A16), $\partial V_2/\partial\hat{\mu}_\gamma = \mu_\theta \ln x^R + \delta\hat{\mu}_\gamma\mu_\theta^2 + \delta\hat{\mu}_\gamma\sigma_\theta^2(1-\omega_\theta^*)$. Since $\omega_\theta^* = 2\kappa(b,\rho)/(\delta\hat{\mu}_\gamma^2\sigma_\theta^2)$ is decreasing in b, a rise in b increases $(1-\omega_\theta^*)$, so $\partial^2 V_2/\partial\hat{\mu}_\gamma\,\partial b > 0$.

(ii) By the same argument with ρ in place of b, $\partial^2 V_2/\partial\hat{\mu}_\gamma\,\partial\rho > 0$.

(iii) Differentiating V_2 with respect to b and then ρ, using $\kappa(b,\rho) = \kappa_0/[b(1+\rho)]$:

$$\frac{\partial^2 V_2}{\partial b\,\partial\rho} = \frac{\kappa(b,\rho)}{b(1+\rho)}\,(1+\ln\omega_\theta^*) \tag{A18}$$

Since $\kappa(b,\rho)/[b(1+\rho)] > 0$, the sign is determined by $1+\ln\omega_\theta^*$, which is positive when $\omega_\theta^* > 1/e$ and negative when $\omega_\theta^* < 1/e$. □

Lemma B1 (Continuity and differentiability of V_1)**.** *Under Assumption 1:*

(i) ***Continuity of*** $\hat{\mu}_\gamma^*$. *The optimal working belief* $\hat{\mu}_\gamma^*$ *is continuous in* $\mu_{\gamma,0}$ *everywhere, including at the threshold* $\phi_0 = \bar{\phi}$*:*

$$\lim_{\phi_0\downarrow\bar{\phi}} \hat{\mu}_\gamma^{*,int} = \lim_{\phi_0\uparrow\bar{\phi}} \hat{\mu}_\gamma^{*,cor} = \bar{\mu}_\gamma(b,\rho) \tag{A19}$$

(ii) ***Continuity of*** V_1. *The Stage 1 value function* $V_1(\mu_{\gamma,0}, b, \rho)$ *is continuous in* $\mu_{\gamma,0}$ *everywhere.*

(iii) ***Differentiability of*** V_1. V_1 *is continuously differentiable in* $\mu_{\gamma,0}$*, with:*

$$\frac{dV_1}{d\mu_{\gamma,0}} = \frac{\lambda_\gamma}{\sigma_{\gamma,0}^2}\left(\hat{\mu}_\gamma^* - \mu_{\gamma,0}\right) \tag{A20}$$

The sign of this derivative equals the sign of the net distortion $\hat{\mu}_\gamma^* - \mu_{\gamma,0}$*: it is negative when* $\kappa(b,\rho) > \kappa^\dagger(\mu_{\gamma,0})$ *(downward distortion, low bandwidth) and non-negative when* $\kappa(b,\rho) \leqslant \kappa^\dagger(\mu_{\gamma,0})$ *(upward or zero distortion, high bandwidth), where* $\kappa^\dagger(\mu_{\gamma,0})$ *is defined in Proposition 4(iii).*

Proof. See below. □

Proposition B3 (The four escape channels)**.** *Under Assumptions 1, 2, and 3, the critical prior mean* $\bar{\mu}_{\gamma,0}(b,\rho,\ln x^R)$ *satisfies:*

(i) ***Belief channel.*** *$\bar{\mu}_{\gamma,0}$ is independent of $\mu_{\gamma,0}$. A parent escapes the trap if and only if $\mu_{\gamma,0} > \bar{\mu}_{\gamma,0}$. In the corner regime:*

$$\frac{\partial \hat{\mu}^*_\gamma}{\partial \mu_{\gamma,0}} = \frac{\lambda_\gamma/\sigma^2_{\gamma,0}}{P} \in (0,1)$$

(ii) ***Bandwidth channel.***

$$\frac{\partial \bar{\mu}_{\gamma,0}}{\partial b} = -\frac{\sigma^2_{\gamma,0} P \bar{\mu}_\gamma}{2b\lambda_\gamma} < 0$$

(iii) ***Cue-detection expertise channel.***

$$\frac{\partial \bar{\mu}_{\gamma,0}}{\partial \rho} = -\frac{\sigma^2_{\gamma,0} P \bar{\mu}_\gamma}{2(1+\rho)\lambda_\gamma} < 0$$

The relative effectiveness of channels (ii) and (iii) is:

$$\frac{\partial \bar{\mu}_{\gamma,0}/\partial b}{\partial \bar{\mu}_{\gamma,0}/\partial \rho} = \frac{1+\rho}{b}$$

(iv) ***Trap compression channel.***

$$\frac{\partial \bar{\mu}_{\gamma,0}}{\partial \ln x^R} = -\frac{\sigma^2_{\gamma,0}\mu_\theta}{\lambda_\gamma} < 0$$

In the corner regime:

$$\frac{\partial \hat{\mu}^{*,corner}_\gamma}{\partial \ln x^R} = \frac{\mu_\theta}{P} > 0$$

Proof. See below. □

Proof of Proposition B3 (The four escape channels)

We prove each part by direct differentiation of:

$$\bar{\mu}_{\gamma,0} = \frac{\sigma^2_{\gamma,0}}{\lambda_\gamma}\left[\bar{\mu}_\gamma(b,\rho)P - \mu_\theta \ln x^R\right]$$

where $\bar{\mu}_\gamma(b,\rho) = \sqrt{2\kappa(b,\rho)/(\delta\sigma^2_\theta)}$ and $\kappa(b,\rho) = \kappa_0/[b(1+\rho)]$.

Part (i): Belief channel. $\bar{\mu}_{\gamma,0}$ does not depend on $\mu_{\gamma,0}$ — it depends only on b, ρ, $\ln x^R$, and structural parameters. The corner solution is $\hat{\mu}^*_\gamma = B/P$ where $B = \mu_\theta \ln x^R + \lambda_\gamma \phi_0$ and $\phi_0 = \mu_{\gamma,0}/\sigma^2_{\gamma,0}$. Differentiating with respect to $\mu_{\gamma,0}$:

$$\frac{\partial \hat{\mu}^*_\gamma}{\partial \mu_{\gamma,0}} = \frac{\lambda_\gamma/\sigma^2_{\gamma,0}}{P}$$

This is strictly positive since $\lambda_\gamma > 0$ and $P > 0$ under Assumption 1. It is strictly less than one since $\lambda_\gamma/\sigma^2_{\gamma,0} < P$ reduces to $\delta(\psi\sigma^2_\theta - \mu^2_\theta) > 0$, which holds under Assumption 2. □

Part (ii): Bandwidth channel. By the chain rule:

$$\frac{\partial \bar{\mu}_{\gamma,0}}{\partial b} = \frac{\sigma^2_{\gamma,0}}{\lambda_\gamma} P \cdot \frac{\partial \bar{\mu}_\gamma}{\partial b}$$

Computing $\partial\bar{\mu}_\gamma/\partial b$. Since $\bar{\mu}_\gamma = \sqrt{2\kappa/(\delta\sigma^2_\theta)}$ and $\partial\kappa/\partial b = -\kappa/b$:

$$\frac{\partial \bar{\mu}_\gamma}{\partial b} = \frac{1}{\sqrt{2\kappa\delta\sigma^2_\theta}} \cdot \left(-\frac{\kappa}{b}\right) = -\frac{\bar{\mu}_\gamma}{2b}$$

Substituting:

$$\frac{\partial \bar{\mu}_{\gamma,0}}{\partial b} = \frac{\sigma^2_{\gamma,0}}{\lambda_\gamma} P \cdot \left(-\frac{\bar{\mu}_\gamma}{2b}\right) = -\frac{\sigma^2_{\gamma,0} P \bar{\mu}_\gamma}{2b\lambda_\gamma} < 0$$

since $\sigma^2_{\gamma,0}$, P, $\bar{\mu}_\gamma$, b, and λ_γ are all strictly positive. □

Part (iii): Cue-detection expertise channel. By the chain rule:

$$\frac{\partial \bar{\mu}_{\gamma,0}}{\partial \rho} = \frac{\sigma^2_{\gamma,0}}{\lambda_\gamma} P \cdot \frac{\partial \bar{\mu}_\gamma}{\partial \rho}$$

Computing $\partial\bar{\mu}_\gamma/\partial\rho$. Since $\partial\kappa/\partial\rho = -\kappa/(1+\rho)$:

$$\frac{\partial \bar{\mu}_\gamma}{\partial \rho} = \frac{1}{\sqrt{2\kappa\delta\sigma^2_\theta}} \cdot \left(-\frac{\kappa}{1+\rho}\right) = -\frac{\bar{\mu}_\gamma}{2(1+\rho)}$$

Substituting:

$$\frac{\partial \bar{\mu}_{\gamma,0}}{\partial \rho} = \frac{\sigma^2_{\gamma,0}}{\lambda_\gamma} P \cdot \left(-\frac{\bar{\mu}_\gamma}{2(1+\rho)}\right) = -\frac{\sigma^2_{\gamma,0} P \bar{\mu}_\gamma}{2(1+\rho)\lambda_\gamma} < 0$$

The relative effectiveness of channels (ii) and (iii) follows directly from dividing the two derivatives:

$$\frac{\partial \bar{\mu}_{\gamma,0}/\partial b}{\partial \bar{\mu}_{\gamma,0}/\partial \rho} = \frac{2(1+\rho)\lambda_\gamma}{2b\lambda_\gamma} = \frac{1+\rho}{b}$$

□

Part (iv): Trap compression channel. Differentiating $\bar{\mu}_{\gamma,0}$ directly with respect to $\ln x^R$, noting that $\bar{\mu}_\gamma$ and P do not depend on $\ln x^R$:

$$\frac{\partial \bar{\mu}_{\gamma,0}}{\partial \ln x^R} = \frac{\sigma^2_{\gamma,0}}{\lambda_\gamma} \cdot (-\mu_\theta) = -\frac{\sigma^2_{\gamma,0}\mu_\theta}{\lambda_\gamma} < 0$$

For the corner solution, differentiating $\hat{\mu}_{\gamma}^{*} = B/P$ with respect to $\ln x^{R}$, noting that $\partial B / \partial \ln x^{R} = \mu_{\theta}$:

$$\frac{\partial \hat{\mu}_{\gamma}^{*,\text{corner}}}{\partial \ln x^{R}} = \frac{\mu_{\theta}}{P} > 0$$

since $\mu_{\theta} > 0$ and $P > 0$ under Assumption 1. □

C Supplementary Information about the LENA Start Program

The LENA Start Program is a ten-week group-based parenting intervention designed to increase the quantity and quality of parent-child conversational interactions in the first three years of life. The program is delivered by a trained Family Liaison in weekly group sessions of approximately 90 minutes. Each session combines educational content, group discussion, coaching activities, and individualized feedback. Families receive age-appropriate children's books at each session and a monthly Developmental Snapshot report. The five components of the program — Education, Coaching, Feedback, Book Reading, and Reporting — are woven throughout the ten sessions rather than delivered sequentially. Below we describe each component in detail and then provide a session-by-session overview.

C.1 The Five Components

Education. The Education component delivers scientific content about the returns to responsive interaction through videos shown during group sessions and a parent handbook distributed at enrollment. Topics covered across the ten sessions include the neuroscience of early brain development, the role of conversational turns in language acquisition, the serve-and-return framework for responsive interaction, and the developmental consequences of adult word counts and conversational turns at different stages of early childhood. The content is designed to be accessible and engaging for parents with varying levels of educational attainment, and is delivered in both English and Spanish depending on the language group of the cohort. The passive delivery format — group sessions at a scheduled time, with a trained Family Liaison presenting the material — is a deliberate design choice. Unlike voluntary information acquisition, which requires the parent to seek out information about the returns to responsive parenting actively, the Education component delivers the signal without requiring the parent to overcome any informational search cost.

Coaching. The Coaching component teaches parents to recognize and respond to child communicative bids through two mechanisms. First, it provides structured practice in identifying child vocalizations, gestures, and bids for joint attention as cues that merit a contingent response. Second, it builds a repertoire of interaction activities that can be integrated into existing daily routines — narrating surroundings during cooking, naming objects during shopping trips, counting steps during walks, and reading at bedtime — so that responsive engagement does not require the parent to carve out dedicated interaction time but instead becomes embedded in activities the parent already performs. The 14 LENA Start talking tips, distributed across the ten sessions, provide concrete guidance for both mechanisms. Table C1 in Appendix C.3 classifies each talking tip within the model's channel taxonomy, identifying whether it primarily raises cue-detection expertise ρ or the reference investment x^R.

Feedback. The Feedback component provides parents with weekly personalized reports based on LENA recordings made in the home during the week preceding each session. Each report summarizes the parent's adult word count, conversational turns, and television or electronic media exposure for the recording day, and presents these measures relative to the parent's own prior recordings and relative to age-appropriate benchmarks. Reports are reviewed with the Family Liaison in a dialectical structure: the Family Liaison begins by praising achievements relative to the previous week, followed by collaborative identification of one or two areas for growth in the coming week. The feedback cycle is operationalized through a star system: parents receive stars for meeting weekly targets on conversational turns and adult word counts, and for keeping television and electronic media exposure below a threshold. Targets are set individually based on each parent's baseline recording and are adjusted upward as the parent improves across sessions. The quantitative, personalized, and weekly nature of the feedback is feasible because the LENA system's recording, processing, and report generation are fully automated: the parent wears the device at home, uploads the recording, and receives a report without a research team present. Figure C.4 in Appendix C.4 displays a representative LENA Start report, illustrating the weekly feedback structure, the quantitative targets, and the daily book reading tracker.

Book Reading. The Book Reading component provides age-appropriate children's books at each session and teaches dialogic reading techniques through in-session practice. Dialogic reading techniques include asking open-ended questions about the pictures, waiting for the child to point or vocalize before turning the page, expanding on the child's vocalizations, and following the child's gaze to determine which elements of the page the child is attending to. These techniques transform shared book reading from a one-directional adult activity into a conversational exchange in which the child's communicative bids structure the interaction. The books provided at each session serve a dual purpose: they provide a tangible practice opportunity between sessions, and because they are new at each session, they provide a novel stimulus more likely to elicit child vocalizations and bids for joint attention, giving the parent practice in detecting and responding to a range of child cues.

Reporting. The Reporting component administers a LENA Developmental Snapshot approximately monthly — at baseline, at the midpoint of the program, and at endline. The Snapshot is a parent-reported instrument that assesses the child's language development across several domains and provides a percentile score relative to age-matched peers in a normative sample. The Family Liaison reviews the Snapshot results with the parent individually, explaining what the percentile score means for the child's developmental trajectory and connecting it to the conversational turn counts recorded by the LENA device. The individualized nature of the Snapshot report is a deliberate design choice: a child in the 15th percentile for expressive language is a child whose conversational opportunities were not fully realized, and this evidence is harder to rationalize away than the general scientific content delivered in the Education component because it is specific to the parent's own child.

C.2 Session-by-Session Overview

Each of the ten weekly sessions weaves together the five program components — Education, Coaching, Feedback, Book Reading, and Reporting — in a structured format that builds progressively across the program. The session-by-session overview below identifies the primary educational topic, the coaching activities, and the feedback and reporting elements delivered in each session.

Session 1: **Introduction to LENA Start.** Parents explore why adult words and conversational turns matter for child development (*Education*), learn to operate the LENA recorder (*Coaching*), and complete their first LENA Snapshot evaluation (*Reporting*).

Session 2: **Understanding LENA Reports and Talking Tips.** Parents interpret their initial LENA report (*Feedback*), discover evidence-based talking tips to boost words and turns (*Coaching*), and practice applying these techniques in small groups.

Session 3: **Shared Reading.** Through interactive exercises, parents master dialogic shared reading strategies to increase verbal engagement (*Book Reading*), then apply and reflect on these techniques (*Coaching*).

Session 4: **Songs and Rhymes.** Parents learn to integrate songs, rhymes, and fingerplays into daily routines (*Coaching*), review multiple-recording LENA reports (*Feedback*), set personalized language goals, and earn stars for meeting targets.

Session 5: **More About Your Baby's Brain.** This session deepens parents' understanding of early brain development (*Education*), reviews group LENA reports (*Feedback*), and practices talking tips (*Coaching*).

Session 6: **Midpoint Reflection.** Families share successes and challenges, revisit their talking strategies (*Coaching*), analyze updated LENA data (*Feedback*), complete a midpoint Developmental Snapshot (*Reporting*), and celebrate achievements.

Session 7: **Math Talk — Movement.** Parents discover movement-based math talk activities (*Coaching*), practice new conversational prompts, review group reports (*Feedback*), and earn recognition for progress.

Session 8: **Language of Food.** Focusing on mealtime routines, parents learn language-rich narrations around food (*Coaching*), apply talking tips, and assess progress via group LENA reports and awards (*Feedback*).

Session 9: **Math Talk — Space.** Parents explore spatial language through everyday play (*Coaching*), practice targeted prompts, analyze group report trends (*Feedback*), and receive stars for goal attainment.

Session 10: **Graduation Day.** In a celebratory finale, families share reflections, review final LENA reports (*Feedback*), watch the "Growing Together" video (*Education*), complete a final Developmental Snapshot (*Reporting*), and receive diplomas and group photos.

C.3 The 14 Talking Tips and the Model's Channel Taxonomy

The LENA Start Program delivers 14 talking tips to parents across sessions, providing concrete, actionable guidance for increasing conversational engagement with their child. Although the curriculum presents the tips as a unified set, the model's channel taxonomy reveals that they operate through two distinct mechanisms.

Tips that train the parent to notice, interpret, and respond contingently to child communicative bids build cue-detection expertise, raising ρ and reducing the attention cost $\kappa(b,\rho)$. These are serve-and-return behaviors: the parent waits for the child's signal, reads it, and calibrates her response to it.

Tips that increase the parent's default level of verbal and physical engagement, regardless of whether the child has emitted a specific cue, raise the reference investment x^R. These behaviors improve developmental outcomes through the level channel even in the corner regime, where the parent is not yet reading child cues.

Table C1 classifies all 14 tips within this taxonomy. The majority of tips operate through ρ, consistent with the program's primary goal of building responsive, cue-contingent parenting. The remaining tips operate through x^R, raising the baseline of verbal and physical engagement. The two sets are complementary: ρ-raising tips make engagement more calibrated, while x^R-raising tips make engagement more frequent.

Table C1: The 14 Talking Tips: Channel Classification

Tip	Description	Channel
1	Talk about what you're doing and thinking	x^R
2	Comment on what they're doing or looking at	x^R
3	Name things that they're interested in	x^R
4	Get down to their level: face to face	x^R
5	Touch, hug, hold	x^R
6	Tune in and respond to what they look at, do, and say	ρ
7	Wait for their response	ρ
8	Imitate them, and add words	ρ
9	Make faces, use gestures	x^R
10	Take turns — don't do all the talking	ρ
11	Repeat and add to what they say and do	ρ
12	Follow their lead, do what interests them	ρ
13	Encourage them, be positive	x^R
14	Be silly! Relax and have fun!	x^R

Notes: ρ = cue-detection expertise: tips that train the parent to notice, interpret, and respond contingently to child cues, reducing the attention cost $\kappa(b,\rho)$. x^R = reference investment: tips that raise the parent's default level of engagement regardless of child cues.

C.4 The LENA Start Report

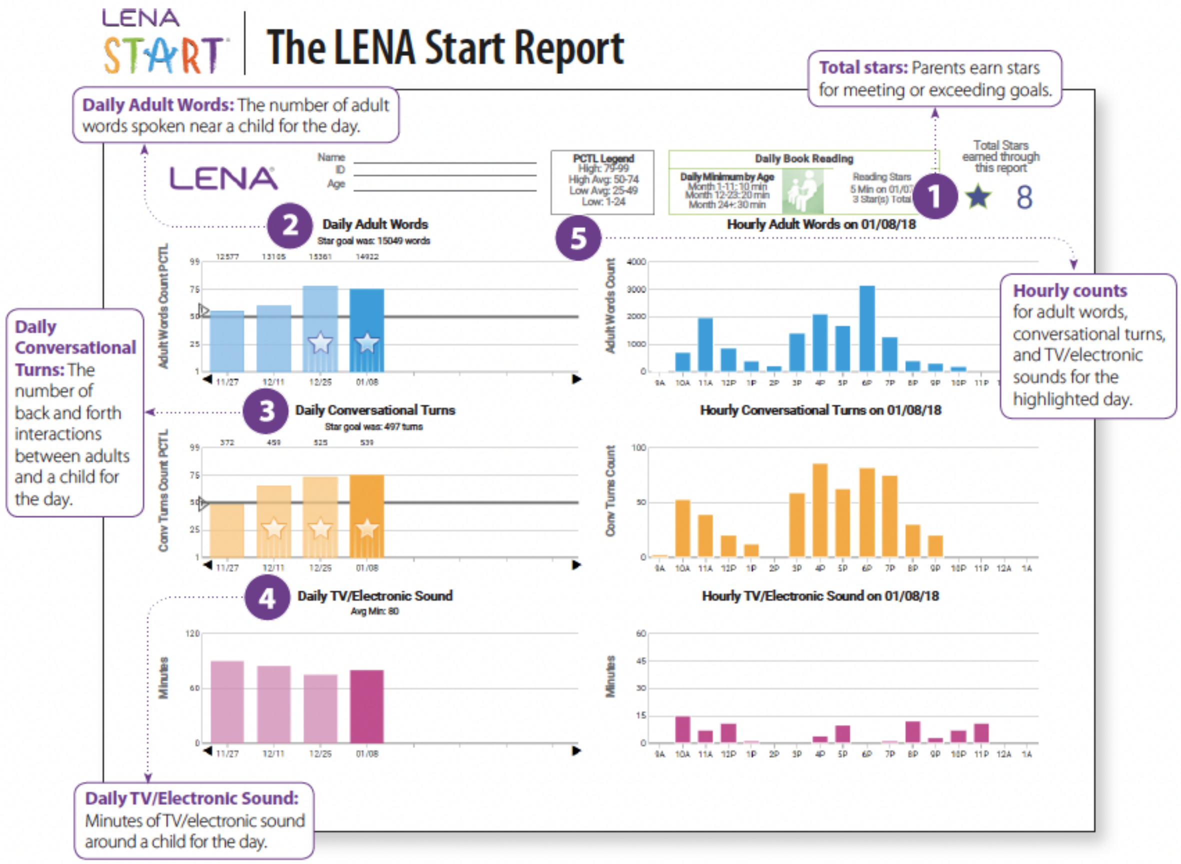

Figure C.4 displays a representative LENA Start Report, the weekly personalized feedback document that forms the core of the Feedback component described in Section 4. Each report is generated automatically from the child's LENA recording and reviewed with the Family Liaison at the following group session.

The report has four main components. At the top (1), the parent's total stars earned to date are displayed alongside a percentile legend that situates the family's conversational engagement relative to other LENA Start families. Below, the Daily Adult Words (2) and Daily Conversational Turns (3) panels display the parent's daily counts for the recording period, the star goal for each metric, and a percentile ranking. The Hourly panels show the within-day distribution of adult words, conversational turns, and television or electronic sound exposure (4) across the recording day, allowing the parent and Family Liaison to identify specific times of day when engagement is high or low and to target routine-building efforts accordingly. Finally, the Daily Book Reading (5) panel tracks the parent's self-reported reading minutes against the age-graded daily minimum (10 minutes for children under 12 months, 20 minutes for 12–23 months, and 30 minutes for children 24 months and older), and awards Reading Stars for meeting the target.

In the model's language, the LENA Start Report operates through two channels simultaneously. The objective, personalized data on conversational turns and adult words raises the accuracy cost of maintaining distorted beliefs about the returns to responsive interaction. A parent who can see her own hourly engagement profile, and then report on the child's language development monthly, can no longer vaguely believe that responsive parenting does not matter. The star system and quantitative targets raise $\ln x^R$, anchoring the parent's default engagement level to a higher reference point and activating the trap compression channel established in Section 3.3.

C.5 Derivations of Testable Implications

This appendix provides the full derivations of the four testable implications stated in Section 4.2. Each derivation identifies the model mechanism that generates the prediction and traces it through the model's parameters and channels.

C.5.1 TI-1: Large Effects on Parental Beliefs

The belief channel and the cue-detection expertise channel each push $\hat{\mu}_\gamma^*$ upward through distinct mechanisms, and their effects are cumulative.

The belief channel operates through the Education and Reporting components. Education delivers scientific content about the returns to responsive interaction passively, in a group setting, at a scheduled time. This shifts $\mu_{\gamma,0}$ toward the truth and collapses $\sigma^2_{\gamma,0}$. The tighter prior raises $A_\delta = \lambda_\gamma/\sigma^2_{\gamma,0} - \delta(\mu_\theta^2 + \sigma_\theta^2)$, increasing the penalty per unit of downward distortion of $\hat{\mu}_\gamma^*$. The parent who has received the Education component can no longer shade $\hat{\mu}_\gamma^*$ downward at negligible cost. Reporting reinforces this effect by providing output-side evidence about the consequences of the parent's engagement: a child in the 15th percentile for expressive language is a child whose conversational opportunities were not fully realized, and this signal about γ is harder to rationalize away than general scientific content.

The cue-detection expertise channel operates through the Coaching and Feedback components. Coaching raises ρ, which reduces $\kappa(b,\rho) = \kappa_0/[b(1+\rho)]$, which in turn reduces $G(\hat{\mu}_\gamma, b, \rho)$. A parent with greater cue-detection expertise achieves a Stage 2 value V_2 closer to the unconstrained benchmark V_2^*, because cue detection no longer requires the same deliberate cognitive effort. The aspiration-capacity gap narrows, making it less painful to hold accurate beliefs. Feedback reinforces this by building cue-detection awareness across recording sessions: the parent becomes attentive to her conversational behavior on recording days, and this attentiveness generalizes across sessions, further raising ρ.

A program that activates both channels simultaneously produces larger and more persistent belief effects than one that activates only one, because the two channels are complementary: the belief channel raises the material return to accurate beliefs, while the cue-detection expertise channel reduces their psychological cost. LENA Start activates both channels, and the model therefore predicts large effects on the belief factor score.

C.5.2 TI-2: Effects on Responsive Engagement

The model predicts that the intervention increases responsive engagement through two distinct margins.

The first margin operates for parents in the corner regime ($\hat{\mu}^*_\gamma \leqslant \bar{\mu}_\gamma$), who currently respond uniformly regardless of the child's signal. The intervention may move these parents into the interior regime through two routes. Through the belief channel, Education and Reporting raise $\mu_{\gamma,0}$, which pushes $\hat{\mu}^*_\gamma$ upward. If the resulting working belief crosses $\bar{\mu}_\gamma$, the parent activates cue detection for the first time. Through the cue-detection expertise channel, Coaching, Feedback, and Book Reading raise ρ, which lowers $\bar{\mu}_\gamma = \sqrt{2\kappa_0/[\delta\sigma^2_\theta b(1+\rho)]}$ and therefore lowers $\bar{\mu}_{\gamma,0}$, moving the threshold toward the parent rather than the parent toward the threshold. Either route, or both simultaneously, can move a corner regime parent into the interior regime and activate cue detection.

The second margin operates for parents already in the interior regime ($\hat{\mu}^*_\gamma > \bar{\mu}_\gamma$). For these parents, the intervention raises $\hat{\mu}^*_\gamma$, which increases engagement intensity through the optimal attention ratio:

$$\omega^*_\theta = \frac{2\kappa(b,\rho)}{\delta\hat{\mu}^{*2}_\gamma\sigma^2_\theta}$$

Since ω^*_θ is decreasing in $\hat{\mu}^*_\gamma$, stronger beliefs generate more attentive monitoring of child cues, increasing the weight placed on the observed signal s/n relative to the prior mean μ_θ in the optimal response. The parent reads child cues more attentively and responds more contingently.

Both margins predict positive treatment effects on conversational turns. The first margin predicts effects concentrated in child-initiated exchanges, since parents moving from the corner to the interior regime begin reading and responding to child cues for the first time. The second margin predicts effects across both child-initiated and adult-initiated exchanges, since more attentive monitoring raises the quality of all interactions. The concentration of treatment effects in child-initiated blocks documented in Section 6 is therefore consistent with the first margin being the dominant source of the treatment effect.

C.5.3 TI-3: Differential Effects on Behavioral Versus Material Components

TI-3 is stated and derived in full in Section 4.2 of the main paper and is not repeated here.

C.5.4 TI-4: Null Effects on Preferences and Self-Efficacy

The model's distortion mechanism operates entirely through the psychological cost $G = \psi(V^*_2 - V_2)$, which is the pain of knowing one cannot fully deliver on one's beliefs about the returns to responsive parenting. This is a belief-specific psychological cost, not a general preference parameter or a global measure of confidence. Three implications follow.

First, the intervention leaves preferences over child outcomes unchanged. The parent's utility function $u(x, \theta) = \hat{\mu}_\gamma \theta \ln x - \frac{1}{2\delta}(\ln x - \ln x^R)^2$ reflects her believed returns to responsive interaction, not her fundamental preferences over child development outcomes. The intervention changes $\hat{\mu}^*_\gamma$ but not the structure of the utility function itself. A parent who values child development before the intervention values it equally after; the intervention changes what she believes about the returns to responsive engagement, not what she cares about.

Second, the intervention does not affect general self-efficacy. The mechanism through which Coaching and Feedback raise $\hat{\mu}^*_\gamma$ is cue-detection expertise: the parent becomes better at reading and responding to child cues, which narrows the aspiration-capacity gap $V^*_2 - V_2$ and reduces G. This is domain-specific competence improvement, not a global change in the parent's confidence in her ability to manage her life or achieve her goals. The model therefore predicts no spillover to general self-efficacy measures.

Third, the intervention does not affect stable psychological traits such as altruism, reciprocity, trust, or patience. These traits enter neither the Stage 1 nor the Stage 2 problem and are not moved by any of the four channels. This distinguishes the model from accounts based on confidence or motivation, which would predict spillovers to general self-efficacy, and from accounts based on preferences, which would predict spillovers to measures of how much the parent values child outcomes.

The null predictions in TI-4 are not merely the absence of an effect — they are sharp predictions that follow from the model's structure. The contrast between large precise effects on TI-1, TI-2, and TI-3 and uniformly small imprecise effects on TI-4 outcomes is itself evidence in favor of the model's distortion mechanism over alternative accounts.

C.6 Fidelity of Implementation

A member of the Rice research team observed each session of the LENA Start Program to collect detailed data on session quality, liaison adherence, liaison competence, attendance, and participant engagement. After each session, both the observer and the liaison completed a fidelity-of-implementation survey. These measures allow us to assess consistency in implementation quality across different cohort-language groups.

C.6.1 Adherence

Adherence refers to the extent to which liaisons successfully completed key session-specific tasks. These tasks, recorded on a binary scale (1 for accomplished, 0 for not accomplished), varied by session. We computed adherence scores as the percentage of tasks completed per session.

Table C2 summarizes average adherence scores by cohort and language group. Initial adherence scores in Fall 2022 were lower compared to subsequent cohorts, but overall adherence remained consistently high thereafter, generally within the 90-100% range. Notably, English liaisons achieved perfect (100%) adherence scores in Spring 2024 and Fall 2024, whereas Spanish liaisons reached perfect adherence in Spring 2025.

Table C2: Average Adherence per Cohort

Language Group	Fall '22	Spring '23	Fall '23	Spring '24	Fall '24	Spring '25
English	89.36	94.90	97.16	100.00	100.00	98.51
Spanish	91.04	98.51	92.54	97.01	94.03	100.00
Combined	90.16	96.88	94.82	98.61	96.53	99.31

Notes: Table reports average adherence rates (percentage) for participants across language groups and implementation cohorts.

C.6.2 Overall Competence

Competence captures the liaison's overall preparedness, relationship-building, and ability to create an effective learning environment. We evaluated liaison competence through seven items rated on a 3-point scale: "Needs support" (0), "Emerging" (1), and "Developed" (2).

Table C3: Average Competence per Cohort

Language Group	Fall '22	Spring '23	Fall '23	Spring '24	Fall '24	Spring '25
English	1.56	1.73	1.62	1.43	1.43	1.33
Spanish	1.68	1.60	1.65	1.78	1.39	1.87
Combined	1.62	1.69	1.63	1.57	1.41	1.60

Notes: Table reports average competence scores for implementation across language groups and cohorts. Scores are based on [insert scale/brief description if applicable].

Table C3 summarizes average competence scores by cohort and language group. For English-speaking cohorts, Spring 2023 had the highest average competence (1.73), whereas Spring 2025 had the lowest (1.33). For Spanish-speaking cohorts, Spring 2025 exhibited the highest average competence (1.87), and Fall 2024 the lowest (1.39). Overall, competence scores hovered around a midpoint of 1.5, suggesting that liaison competence generally fell between the "Emerging" and "Developed" categories throughout the implementation period.

C.7 LENA Start Graduation Data

Table C4 reports graduation rates for treatment group participants across the six cohorts. The overall graduation rate was 75% across all cohorts, with cohort-level rates ranging from 69.23% in Cohort 2 to 86.21% in Cohort 1.[10]

[10] To graduate from the LENA Start program, parents must attend the first four core sessions and at least three of the remaining sessions (a minimum of seven out of ten classes total) and complete six valid weekly recordings with the device.

Table C4: LENA Program Enrollment and Graduation by Cohort

Cohort	Enrolled	Graduated	Did Not Graduate	% Graduated
Fall 2022	29	25	4	86.2%
Spring 2023	26	18	8	69.2%
Fall 2023	27	19	8	70.4%
Spring 2024	23	16	7	69.6%
Fall 2024	24	19	5	79.2%
Spring 2025	19	14	5	73.7%
Total	**148**	**111**	**37**	**75%**

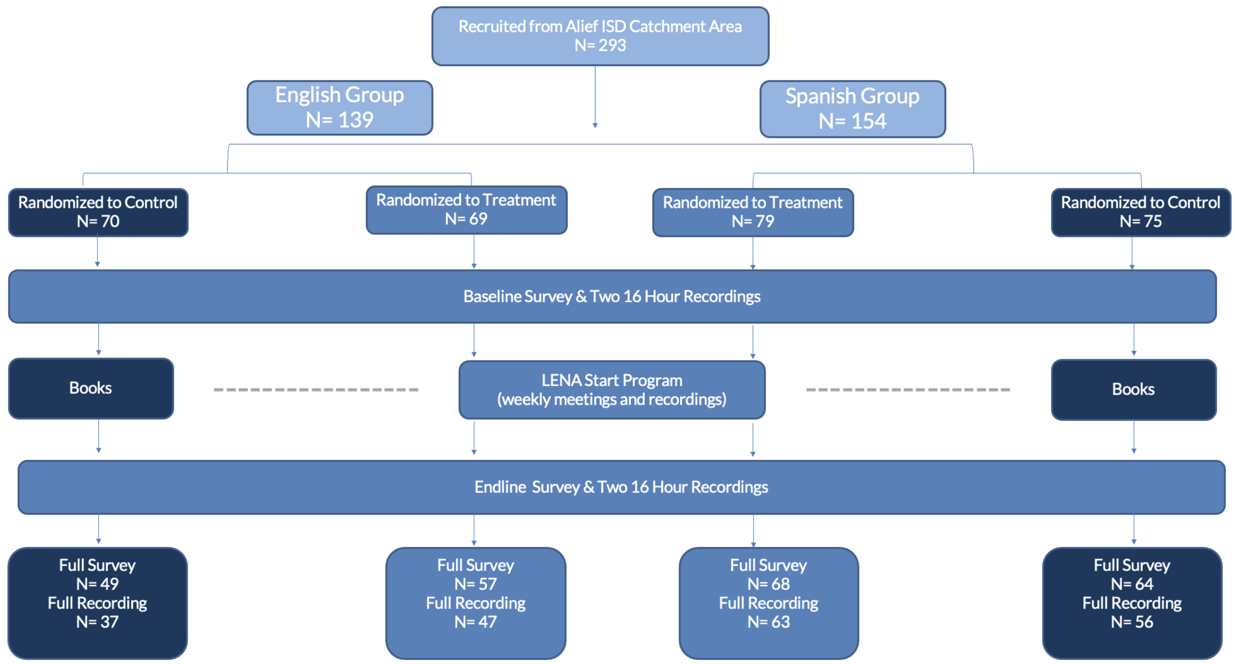

Figure D1: RCT Flow Diagram

D Supplementary Information about Experimental Design

D.1 Group Assignment

Figure D1 presents the flow diagram for the randomized controlled trial. A total of 293 families were recruited from the Alief ISD catchment area and stratified into two language groups: an English group (N = 139) and a Spanish group (N = 154). Within each language group, families were randomized to treatment or control. In the English group, 69 families were assigned to treatment and 70 to control; in the Spanish group, 79 families were assigned to treatment and 75 to control, for a total of 148 treatment and 145 control families. Not counted here are 10 participants we dropped from the study - 9 who had initially enrolled and then withdrew and 1 whose data had been compromised.

At baseline, all families completed a survey and provided LENA recordings. During the intervention period, treatment families participated in the ten-week LENA Start Program, which included weekly group sessions and weekly home recordings; control families received books only. At endline, all families were asked to complete a follow-up survey and provide two additional 16-hour recordings.

Attrition was higher for recordings than for surveys, and higher in the English group than in the Spanish group. In the English group, 57 of 69 treatment families (82.6%) and 49 of 70 control families (70%) completed the endline survey; 47 of 69 treatment families (68.1%) and 37 of 70 control families (52.9%) provided complete endline recordings. In the Spanish group, 68 of 79 treatment families (86.1%) and 64 of 75 control families (85.3%) completed the endline survey; 63

of 79 treatment families (79.7%) and 56 of 75 control families (74.7%) provided complete endline recordings. Across both language groups, recording attrition is lower in the treatment group, with the overall pattern analyzed in detail in Section G.

Table D1: Breakdown of Enrollment by Cohort and Language Group

Cohort	Status	English	Spanish	Total
Cohort 1	Enrolled	29	30	59
	Treatment	14	15	29
	Control	15	15	30
Cohort 2	Enrolled	29	24	53
	Treatment	14	12	26
	Control	15	12	27
Cohort 3	Enrolled	25	28	53
	Treatment	12	15	27
	Control	13	13	26
Cohort 4	Enrolled	21	23	44
	Treatment	11	12	23
	Control	10	11	21
Cohort 5	Enrolled	19	28	47
	Treatment	10	14	24
	Control	9	14	23
Cohort 6	Enrolled	16	21	37
	Treatment	8	11	19
	Control	8	10	18
All Cohorts	Enrolled	139	154	293
	Treatment	69	79	148
	Control	70	75	145

Notes: Table reports the number of children enrolled and their subsequent random assignment to treatment or control groups across all implementation cohorts.

Table D1 reports enrollment counts and treatment assignment by cohort and language group across the six implementation cohorts. Cohort sizes declined modestly over time, from 59 families in Fall 2022 to 37 families in Spring 2025, reflecting variation in recruitment across cohorts. The largest cohort was Fall 2022 ($N = 59$) and the smallest was Spring 2025 ($N = 37$). The ratio of English-speaking to Spanish-speaking families varied across cohorts, with Spanish-speaking families slightly more numerous in most cohorts and the gap widening in Fall 2024 (19 English, 28 Spanish) and Spring 2025 (16 English, 21 Spanish). Treatment and control group sizes are balanced within each cohort-language stratum, with differences of at most one family in any given stratum, confirming that the randomization protocol was implemented as designed.

D.2 Balance Tests

Tables D2, D3, and D4 report baseline means and standard deviations for treatment and control groups across all outcome and control variables, organized by testable implication. P-values are from tests of the difference in means between treatment and control groups, with standard errors clustered at the randomization cluster level.

Panel A reports balance on demographic control variables for the full analytical sample of 238 families. The sample is predominantly Hispanic (66–70%), with roughly equal shares of Non-Hispanic Black families (27–30%). Approximately half of primary family members have attended college, and approximately 91% are mothers of the enrolled child. About 70–74% of families have a co-parent present, and 56–59% are married. The mean child age is approximately 23 months and gender distributions are roughly equal across groups. All p-values in Panel A exceed 0.10, confirming that randomization achieved balance across demographic characteristics.

Panel B reports balance on TI-1 outcomes — parental knowledge, parental beliefs, and the belief factor score — for the full sample of 238 families. Baseline means are similar across treatment and control groups, with p-values of 0.139, 0.284, and 0.583 respectively, indicating no statistically significant pre-existing differences in parental beliefs or knowledge.

Panel C reports balance on TI-2 outcomes for the 203 families with valid baseline LENA recordings. Baseline conversational turn counts, adult word counts, and ADEX-decomposed measures of child-initiated, female adult-initiated, and male adult-initiated conversation blocks are all balanced across treatment and control groups, with all p-values exceeding 0.10.

Panel D reports balance on TI-3 outcomes — the StimQ Total Score and its four subscales — for the full sample of 238 families. All subscale scores are balanced across groups, with p-values ranging from 0.301 to 0.862.

Panel E reports balance on TI-4 placebo outcomes for the full sample, with observations ranging from 177 to 238 depending on the outcome. Two variables approach but do not reach conventional significance levels: Parenting Social Support ($p = 0.066$) and Trust ($p = 0.088$). All other placebo outcomes are well-balanced across groups. We return to the two near-significant variables when discussing TI-4 results in Section 6, and note that neither survives a Bonferroni correction for multiple testing across the eight placebo outcomes, which requires $p < 0.0125$ for significance at the 10% level.

Table D2: Balance Table - Part 1: Demographics and Parental Beliefs

Variable	Treated Mean	Treated SD	Control Mean	Control SD	p-value	N
Panel A: Demographic and Control Variables						
Child Male (Yes=1)	0.464	0.501	0.540	0.501	0.211	238
Child Age (in months)	23.328	3.996	23.265	4.363	0.905	238
College (Yes=1)	0.472	0.501	0.513	0.502	0.597	238
Mother (Yes=1)	0.912	0.284	0.920	0.272	0.769	238
Co-Parent (Yes=1)	0.744	0.438	0.708	0.457	0.550	238
Hispanic (Yes=1)	0.664	0.474	0.699	0.461	0.369	238
Non-Hispanic Black (Yes=1)	0.304	0.462	0.274	0.448	0.437	238
Married (Yes=1)	0.560	0.498	0.593	0.493	0.604	238
Cohort 1 (Yes=1)	0.192	0.395	0.177	0.383	0.124	238
Cohort 2 (Yes=1)	0.176	0.382	0.177	0.383	0.851	238
Cohort 3 (Yes=1)	0.184	0.389	0.195	0.398	0.280	238
Cohort 4 (Yes=1)	0.160	0.368	0.150	0.359	0.602	238
Cohort 5 (Yes=1)	0.152	0.360	0.168	0.376	0.141	238
Cohort 6 (Yes=1)	0.136	0.344	0.133	0.341	0.494	238
English (Yes=1)	0.456	0.500	0.434	0.498	0.117	238
Employed (PFM)	0.480	0.502	0.389	0.490	0.241	238
Household Size	5.024	1.298	5.168	1.487	0.449	238
Panel B: TI-1 Beliefs						
Parental Knowledge	0.590	0.129	0.571	0.123	0.139	238
Parental Belief	1.413	0.732	1.288	0.715	0.284	238
Belief Factor Score	-0.091	0.835	-0.170	0.945	0.583	238

Table D3: Balance Table - Part 2: LENA Metrics (TI-2)

Variable	Treated Mean	Treated SD	Control Mean	Control SD	p-value	N
Conversational Turns	369.809	211.471	348.742	188.708	0.435	203
Adult Word Count	10974.927	5821.242	10499.763	6501.857	0.683	203
Turns (Child Init)	151.118	93.691	138.161	76.689	0.181	203
Turns (Female Init)	118.655	73.217	109.355	72.863	0.393	203
Turns (Male Init)	37.136	40.023	36.559	34.973	0.922	203
Conv. Count (Child)	87.173	40.371	82.054	35.743	0.238	203
Conv. Count (Female)	63.382	30.123	58.860	29.788	0.380	203
Conv. Count (Male)	20.764	19.797	20.602	16.226	0.955	203
Duration (Child Init)	1874.137	1171.267	1762.077	1015.592	0.398	203
Duration (Female Init)	1666.006	1077.221	1544.422	1103.360	0.434	203
Duration (Male Init)	534.422	576.462	590.104	842.921	0.639	203

Table D4: Balance Table - Part 3: StimQ (TI-3) and Placebo Outcomes (TI-4)

Variable	Treated Mean	Treated SD	Control Mean	Control SD	p-value	N
Panel A: TI-3 StimQ Components						
StimQ Total Score	28.304	8.338	27.735	7.656	0.419	238
PVR	9.240	3.555	8.867	3.452	0.301	238
READ	10.536	3.423	10.478	3.246	0.862	238
PIDA	3.112	1.597	2.938	1.525	0.426	238
ALM	5.416	1.582	5.451	1.524	0.859	238
Panel B: TI-4 Placebo Outcomes						
Self Efficacy	-0.048	0.828	-0.077	0.790	0.789	238
Social Support	0.066	0.851	-0.112	0.821	0.066	238
Altruism	0.133	0.613	0.023	0.708	0.263	202
Pos. Reciprocity	0.044	0.779	0.017	0.662	0.711	219
Neg. Reciprocity	-0.009	0.775	0.136	0.746	0.249	177
Trust	-0.102	1.011	0.123	0.935	0.088	222
Patience	0.017	0.945	-0.096	1.186	0.385	230

E Supplementary Information about Measurement and Data

E.1 TI-1: Beliefs and Knowledge

Parental Knowledge. Parental knowledge is measured using the questionnaire developed by Suskind et al. (2016), which comprises five subscales: perceptions about how children learn to talk (9 items), the importance of reading to children (6 items), children's math learning (5 items), the relationship between language development and school readiness (7 items), and perceptions about television exposure and language development (3 items). Responses to each item were originally rated on a five-point Likert scale ranging from "strongly disagree" to "strongly agree" and were recoded as correct (1) or incorrect (0) according to the established guidelines by Suskind et al. (2016). The parental knowledge score is the average of these binary scores across all items.

Parental Beliefs. Parental beliefs were assessed using two scenario-based items (Cunha, Gerdes, Hu, and Nihtianova 2024). Scenario 1 described a home environment characterized by frequent parental engagement (talking and reading) and limited television exposure; Scenario 2 described minimal parental engagement and high television exposure. Respondents predicted future language development outcomes for each scenario using a five-point scale corresponding to Z-scores at the 10th, 25th, 50th, 75th, and 90th percentiles. The parental belief measure is the difference between the predicted outcomes (Z-scores) for Scenario 1 and Scenario 2, capturing the extent to which parents perceive the developmental impact of different home environments.

Belief Factor Score. To account for potential measurement error in both instruments, we followed Cunha, Elo, and Culhane (2022) and conducted a factor analysis on the items in the Parental Beliefs questionnaire and the Parental Knowledge questionnaire jointly. The resulting factor scores provide an error-adjusted measure of parental beliefs about the returns to responsive interaction and serve as the primary outcome for TI-1.

E.2 TI-2: Cue Detection and Responsive Engagement

LENA recordings. The LENA recording system documented each child's language environment over a typical day at baseline and endline (Gilkerson and Richards 2020). Families received specialized clothing with a vest pocket to securely hold the LENA Digital Language Processor, enabling automated capture of adult word counts, child vocalizations, and conversational turns between adults and children.

Conversational turns. A conversational turn is counted when a Key Child segment containing at least one speech-related vocalization and an Adult Male or Female segment containing at least one word occur in alternation, separated by no more than five seconds of silence or other non-speech sounds. The turn can be initiated by either party; once a segment has contributed to a turn it cannot initiate another. Segments from other children in the vicinity interrupt and preclude turns. The LENA system does not identify the directionality of adult speech and cannot distinguish child-directed from adult-directed utterances within a turn. A few parents provided recordings for two days; for those parents, one recording was randomly selected for analysis.

Recording period selection. Recorded data were analyzed in five-minute increments. Given families' varying routines, each family's best 12-hour consecutive period — defined as the period with the highest conversational engagement — was selected for primary analysis. Supplemental analyses of the best 8-hour and 10-hour periods are provided in Appendix F.

ADEX decomposition. The Advanced Data Extractor (ADEX) feature of the LENA system decomposes conversational turns by initiator: Key Child, Adult Female, or Adult Male. ADEX groups time-adjacent LENA segments into conversation blocks, defined as sequences containing at least one Key Child (CHN) or near-field adult speech segment (MAN or FAN) and separated from neighboring blocks by at least five seconds in which only silence, noise, television, or other non-live-speech segments occur. Within each block, ADEX identifies the initiator by scanning forward from the block's start time to the first qualifying live-speech segment: if that segment is CHN, the block is coded as child-initiated; if it is MAN or FAN, the block is coded as adult-initiated (male or female, respectively). Each block in the ADEX output is characterized by its initiator, its duration, and the number of conversational turns it contains.

E.3 TI-3: Behavioral Versus Material Components of the Home Environment

StimQ Toddler. The StimQ Toddler questionnaire (Cates et al. 2023) is a 39-point scale measuring cognitive stimulation in the homes of children aged 12 to 36 months. The total StimQ score is the sum of four subscale scores: Reading (READ), Parental Involvement in Developmental Advance (PIDA), Parental Verbal Responsivity (PVR), and Availability of Learning Materials (ALM).

Parental Verbal Responsivity (PVR). The PVR subscale asks about the nature of parent-child verbal interaction, including how often the caregiver talks and responds to the child during daily routines and play, engages in back-and-forth conversations, narrates surroundings, and joins in pretend play. PVR is the closest available survey-based proxy for the cue-contingent component of investment x: higher PVR reflects more contingent verbal engagement calibrated to the child's communicative bids, though the subscale also contains level items such as narrating surroundings regardless of the child's signal.

Reading (READ). The Reading subscale asks whether parents read books to the child, the names of children's books read in the recent past, the frequency of reading, and the types of books read. It captures the frequency, diversity, and quality of shared book reading, including dialogic techniques such as waiting for the child to point and following the child's gaze. READ maps primarily onto the reference investment x^R: it captures structured parental behaviors that improve developmental outcomes through the level channel.

Parental Involvement in Developmental Advance (PIDA). The PIDA subscale contains ten items about intentional teaching activities — counting with objects, naming colors and body parts, demonstrating play skills — that constitute deliberate engagement independent of whether the child has initiated a specific cue. PIDA maps onto the reference investment x^R through the level channel.

Availability of Learning Materials (ALM). The ALM subscale inventories toys and materials available to the child, divided into materials for symbolic play (dolls, puppets, housekeeping toys), art (crayons), fine motor skills (shape-sorter toys, stacking toys), and language (toy letters). ALM reflects the material component of the home environment and does not map onto either x or x^R as the model defines them. Since LENA Start provides no resources for parents to purchase learning materials, ALM serves as a within-instrument placebo for TI-3.

E.4 TI-4: Null Effects on Preferences and Self-Efficacy

Self-Efficacy and Social Support. Parental self-efficacy and social support are each measured using a small set of survey items developed by the LENA Foundation (Cunha, Gerdes, Hu, and Nihtianova 2024). Self-efficacy items ask parents to rate their confidence in specific parenting competencies on a five-point scale ranging from "least sure" to "very sure." Social support items ask parents to rate their agreement with statements about community and peer support on a five-point scale ranging from "strongly disagree" to "strongly agree." Latent trait scores for each construct were estimated using an Item Response Theory partial credit model, and predicted scores on the latent traits were used as outcome measures.

Global Preferences. Global preferences are measured using the validated instrument of Falk et al. (2018), which comprises nine items covering five preference domains. Patience is assessed by willingness to forgo immediate rewards. Positive reciprocity is assessed by two items: willingness to give in exchange for help and willingness to return a favor. Negative reciprocity is assessed by three items: willingness to take revenge, punish unfairness to oneself, and punish unfairness to others. Altruism is assessed by two items: donation decisions and willingness to give to good causes. Trust is assessed by beliefs about other people's intentions. Z-scores for each item were calculated and variables and their weights were constructed following Falk et al. (2018).

E.5 Control Variables

The survey collected detailed demographic information about the primary family member (PFM) and non-primary family member (NPFM, if applicable), including educational attainment, race and ethnicity, employment status, and household composition. Child demographic information includes age in months, gender, and race and ethnicity. These variables serve as control variables in all empirical specifications and are used to construct the inverse probability weights described in Appendix Section G.

F Sensitivity Analysis

Our main analyses utilize the 'best' 12-hour periods of LENA recording data. To evaluate the robustness of these findings, we repeated the analysis using shorter durations ('best' 10 and 8 hours).

Using the best 10-hour recordings expands our analytical sample to 210 participants, and further to 214 participants when using the best 8-hour recordings. Adult word count (AWC) treatment effects, initially weakly significant in the 12-hour model, lose significance and diminish in magnitude with shorter durations. Conversely, conversational turn treatment effects increase slightly in magnitude for the 10-hour analysis but decrease somewhat in the 8-hour analysis, dropping from the 1% to the 5% significance level.

Importantly, regardless of duration, results consistently indicate that the increase in conversational turns is primarily driven by children initiating interactions rather than adults. This pattern strongly suggests genuine shifts in parent-child conversational dynamics rather than merely performance effects related to device usage.

Table F1: Impact of LENA Start Program on Adult Words and Conversational Turns (10 Hours)

	Adult Words			Conversational Turns		
	(1)	(2)	(3)	(1)	(2)	(3)
Treatment	0.191	0.155	0.172	0.334***	0.315***	0.329***
	(0.122)	(0.126)	(0.126)	(0.104)	(0.100)	(0.103)
	[0.067]	[0.075]	[0.090]	[0.061]	[0.069]	[0.071]
	{0.118}	{0.222}	{0.175}	{0.002}	{0.003}	{0.002}
Start Time / Weekend	No	Yes	Yes	No	Yes	Yes
Child / PFM Controls	No	Yes	Yes	No	Yes	Yes
MCB Vocabulary Score	No	No	Yes	No	No	Yes

Notes: N = 210. Robust standard errors in (·), clustered standard errors in [·], and RI p-values in {·}. All specifications include baseline conversational turns, language group, and cohort dummies. Child controls: male, age. PFM controls: education, mother, co-parenting, ethnicity, and marital status. * p<0.10, ** p<0.05, *** p<0.01.

Table F2: Impact of LENA Start Program on Adult Words and Conversational Turns (Best 8 Hours)

	Adult Words			Conversational Turns		
	(1)	(2)	(3)	(1)	(2)	(3)
Treatment	0.163	0.145	0.159	0.270**	0.247**	0.266**
	(0.124)	(0.125)	(0.128)	(0.104)	(0.104)	(0.106)
	[0.088]	[0.099]	[0.113]	[0.064]	[0.071]	[0.078]
	{0.204}	{0.259}	{0.228}	{0.013}	{0.024}	{0.016}
Start Time / Weekend	No	Yes	Yes	No	Yes	Yes
Child / PFM Controls	No	Yes	Yes	No	Yes	Yes
MCB Vocabulary Score	No	No	Yes	No	No	Yes

Notes: N = 214. Robust standard errors in (·), clustered standard errors in [·], and RI p-values in {·}. All models include baseline outcome, language group, and cohort dummies. Child controls: male, age. PFM controls: education, mother, co-parenting, ethnicity, and marital status. * p<0.10, ** p<0.05, *** p<0.01.

Similar consistency is observed with the Advanced Data Extractor (ADEX) results. Although shorter recording durations slightly reduce treatment effect magnitudes and significance levels, findings consistently indicate meaningful increases in child-led conversational interactions.

Table F3: Impact of LENA Start Program on Conversation Block (10 Hours)

	(1) Initiated by Child	(2) Initiated by Female	(3) Initiated by Male
Conv. Block Count	0.279**	0.159	0.121
	(0.119)	(0.123)	(0.129)
	[0.098]	[0.113]	[0.126]
	{0.026}	{0.209}	{0.361}
Conv. Turn Count	0.303***	0.308***	0.096
	(0.106)	(0.113)	(0.128)
	[0.050]	[0.062]	[0.133]
	{0.006}	{0.008}	{0.449}
Block Duration	0.357***	0.291**	0.061
	(0.104)	(0.123)	(0.128)
	[0.062]	[0.090]	[0.127]
	{0.002}	{0.024}	{0.644}
Adult Word Count	0.310***	0.281**	-0.019
	(0.113)	(0.130)	(0.132)
	[0.070]	[0.093]	[0.113]
	{0.008}	{0.035}	{0.897}

Notes: Robust standard errors in $(\cdot)$, clustered standard errors in $[\cdot]$, and RI p-values in $\{\cdot\}$. All models include baseline dependent variable, cohort, language, child/parent demographics, recording duration, and start time. * $p<0.10$, ** $p<0.05$, *** $p<0.01$.

Table F4: Impact of LENA Start Program on Conversation Block (8 Hours)

	(1) Initiated by Child	(2) Initiated by Female	(3) Initiated by Male
Conv. Block Count	0.223*	0.068	0.104
	(0.117)	(0.119)	(0.136)
	[0.088]	[0.119]	[0.147]
	{0.084}	{0.589}	{0.431}
Conv. Turn Count	0.251**	0.245**	0.082
	(0.107)	(0.109)	(0.133)
	[0.066]	[0.071]	[0.144]
	{0.031}	{0.040}	{0.537}
Block Duration	0.323***	0.229*	0.079
	(0.105)	(0.118)	(0.137)
	[0.088]	[0.085]	[0.135]
	{0.004}	{0.084}	{0.548}
Adult Word Count	0.303***	0.235*	0.009
	(0.112)	(0.124)	(0.138)
	[0.093]	[0.088]	[0.108]
	{0.012}	{0.084}	{0.943}

Notes: Robust standard errors in $(\cdot)$, clustered standard errors in $[\cdot]$, and RI p-values in $\{\cdot\}$. All models include baseline dependent variable, cohort, language, child/parent demographics, recording duration, and start time. * $p<0.10$, ** $p<0.05$, *** $p<0.01$.

G Attrition Analysis

Attrition occurs along two dimensions: survey attrition, which reduces the sample from 293 to 238 families, and recording attrition, which reduces it from 293 to 203 families. Table G1 shows that attrition rates do not differ significantly between treatment and control groups for survey analysis at conventional levels ($p = 0.72$). For survey attrition, the treated group has an attrition rate of 15.5% and the control group has an attrition rate of 22.1%. Recording attrition differs more between treatment and control groups at a conventional level ($p = 0.037$), as the treated attrition rate is 25.7% and the control rate is 35.9%.

Table G2 examines whether attrition is selective on observable characteristics for the full sample. Two patterns emerge consistently across both attrition dimensions. First, English-speaking families are significantly more likely to attrit: attriting families are 15.5 percentage points more likely to be English-speaking for survey attrition ($p = 0.004$) and 19.7 percentage points more likely for recording attrition ($p = 0.002$). Second, families with higher baseline StimQ scores are significantly more likely to attrit: attriting families have StimQ scores that are 3.4 points higher for survey attrition ($p = 0.006$) and 3.4 points higher for recording attrition ($p = 0.001$). No other demographic characteristics predict attrition.

Table G3 shows that selective attrition on English language and baseline StimQ score is concentrated in the control group. In Panel A, control group attriters are 23.7 percentage points more likely to be English-speaking for recording attrition ($p = 0.006$) and 22.3 percentage points more likely for survey attrition ($p = 0.005$), and have StimQ scores that are 4.2 points higher for recording attrition ($p = 0.012$) and 5.2 points higher for survey attrition ($p = 0.006$). In Panel B, neither English language nor baseline StimQ score predicts attrition in the treatment group, and all differences are small and statistically insignificant at conventional levels.

Table G4 confirms that none of the demographic characteristics, baseline beliefs, responsive engagement measures, StimQ subscales, or placebo outcomes predict attrition once treatment status and cohort indicators are controlled for. The coefficient on treatment status is negative for both recording and survey attrition but is not statistically significant at conventional levels using clustered standard errors. The only variable that approaches significance is baseline StimQ score ($p = 0.082$ for recordings), consistent with the pattern identified in Table G2. To account for the selective nature of attrition on observable characteristics, all specifications include inverse probability weights constructed from the covariates in Table D2.

Table G1: T-Test Results: Sample Attrition Comparisons

Comparison	Diff.	SE	t-stat	p-value	Mean (T)	Mean (C)	N(T)	N(C)
293 to 203	-0.1019**	0.0430	-2.370	0.0371	0.2568	0.3586	148	145
293 to 238	-0.0653*	0.0328	-1.991	0.0720	0.1554	0.2207	148	145

Notes: Table reports t-test comparisons for different sample size thresholds. Diff. represents the difference in means between groups. Mean (T) and Mean (C) refer to Treatment and Control group means, respectively.

Table G2: Attrition Analysis for Full Sample

	Survey Attrition				Recording Attrition			
	Non-Attrit	Attrit	Difference	P-value	Non-Attrit	Attrit	Difference	P-value
Child Age (Months)	23.298	23.600	-0.302	0.549	23.281	23.522	-0.241	0.641
Child Male	0.500	0.527	-0.027	0.756	0.488	0.544	-0.057	0.442
College Educated	0.492	0.491	0.001	0.994	0.498	0.478	0.020	0.773
Mother	0.916	0.909	0.007	0.821	0.901	0.944	-0.043	0.129
Co-Parent in Home	0.727	0.673	0.054	0.463	0.729	0.689	0.040	0.547
Hispanic	0.681	0.600	0.081	0.239	0.704	0.578	0.127**	0.045
Non-Hispanic Black	0.290	0.273	0.017	0.717	0.266	0.333	-0.067	0.248
Married	0.576	0.618	-0.043	0.380	0.581	0.589	-0.008	0.898
StimQ Total Score	28.034	31.473	-3.439***	0.006	27.650	31.000	-3.350***	0.001
English-Speaking	0.445	0.600	-0.155***	0.004	0.414	0.611	-0.197***	0.002
Observations	238	55			203	90		

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table G3: Attrition Analysis by Treatment Status

	Survey Attrition				Recording Attrition			
	Non-Attrit	Attrit	Difference	P-value	Non-Attrit	Attrit	Difference	P-value
Panel A: Control Group								
Child Age (Months)	23.265	24.375	-1.110	0.168	23.086	24.269	-1.183	0.143
Child Male	0.540	0.500	0.040	0.744	0.495	0.596	-0.102	0.300
College Educated	0.513	0.531	-0.018	0.865	0.538	0.481	0.057	0.437
Mother	0.920	0.906	0.014	0.768	0.903	0.942	-0.039	0.334
Co-Parent in Home	0.708	0.719	-0.011	0.907	0.710	0.712	-0.002	0.983
Hispanic	0.699	0.531	0.168*	0.098	0.731	0.538	0.193**	0.029
Non-Hispanic Black	0.274	0.281	-0.007	0.921	0.237	0.346	-0.110*	0.095
Married	0.593	0.625	-0.032	0.567	0.602	0.596	0.006	0.925
StimQ Total Score	27.735	32.938	-5.203***	0.006	27.376	31.577	-4.201**	0.012
English-Speaking	0.434	0.656	-0.223***	0.005	0.398	0.635	-0.237***	0.006
Observations	32	113			93	52		
Panel B: Treatment Group								
Child Age (Months)	23.328	22.522	0.806	0.254	23.445	22.500	0.945	0.258
Child Male	0.464	0.565	-0.101	0.397	0.482	0.474	0.008	0.940
College Educated	0.472	0.435	0.037	0.730	0.464	0.474	-0.010	0.913
Mother	0.912	0.913	-0.001	0.988	0.900	0.947	-0.047	0.379
Co-Parent in Home	0.744	0.609	0.135	0.194	0.745	0.658	0.088	0.358
Hispanic	0.664	0.696	-0.032	0.602	0.682	0.632	0.050	0.415
Non-Hispanic Black	0.304	0.261	0.043	0.503	0.291	0.316	-0.025	0.740
Married	0.560	0.609	-0.049	0.592	0.564	0.579	-0.015	0.872
StimQ Total Score	28.304	29.435	-1.131	0.484	27.882	30.211	-2.329	0.129
English-Speaking	0.456	0.522	-0.066	0.340	0.427	0.579	-0.152*	0.062
Observations	125	23			110	38		

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table G4: Attrition Analysis

	Recordings	Survey		Recordings	Survey
Treatment	-0.265*	-0.197	Married	0.056	0.060
	(0.160)	(0.177)		(0.175)	(0.195)
	[0.129]	[0.118]		[0.158]	[0.171]
Child Age	0.001	-0.006	StimQ Score	0.021*	0.020
	(0.020)	(0.023)		(0.012)	(0.014)
	[0.020]	[0.019]		[0.012]	[0.014]
Child Male	0.168	0.081	English	0.332	0.232
	(0.162)	(0.179)		(0.262)	(0.287)
	[0.172]	[0.213]		[0.261]	[0.276]
College	-0.255	-0.205	Cohort 2	0.246	-0.067
	(0.173)	(0.192)		(0.266)	(0.293)
	[0.173]	[0.184]		[0.108]	[0.148]
Mother	0.466	-0.038	Cohort 3	-0.133	-0.218
	(0.324)	(0.333)		(0.273)	(0.297)
	[0.246]	[0.237]		[0.130]	[0.129]
Co-Parent	0.006	-0.205	Cohort 4	0.021	-0.246
	(0.197)	(0.215)		(0.288)	(0.318)
	[0.194]	[0.222]		[0.120]	[0.130]
Hispanic	-0.603	-0.685	Cohort 5	0.225	-0.063
	(0.419)	(0.432)		(0.276)	(0.298)
	[0.176]	[0.240]		[0.150]	[0.112]
Black	-0.558	-0.917**	Cohort 6	-0.112	-0.491
	(0.391)	(0.404)		(0.303)	(0.345)
	[0.281]	[0.200]		[0.144]	[0.183]

Standard errors in parentheses, clustered SE in brackets.
Stars based on most conservative SE: * $p<0.10$, ** $p<0.05$, *** $p<0.01$

G.1 Inverse Probability Weighting

We construct inverse probability weights to account for selective attrition along two dimensions: from 293 to 238 observations for survey-based outcomes, and from 293 to 203 observations for LENA-based outcomes. For both dimensions, the baseline probit model predicts observation at endline using primary family member demographics, child demographics, and cohort-language group indicators (Model 1). Model 2 adds baseline StimQ scores to the probit. This approach maximizes the information used to predict attrition while maintaining a common sample across models.

G.2 Attrition Balancing

We implement an attrition balancing procedure as an alternative robustness check to inverse probability weighting. We first estimate a probit model on the full sample to predict the probability of attrition, controlling for demographic variables, baseline StimQ scores, language group, and cohort indicators. We then match each treated participant to a control participant within the same language-cohort group based on their predicted attrition probabilities. Matched pairs are retained or dropped according to the following rule: pairs in which both participants are observed at endline are retained; pairs in which both attrit are dropped; and pairs in which exactly one participant attrits are dropped. This procedure produces a balanced sample in which attrition rates are equalized across treatment and control groups by construction, providing a check on whether selective attrition drives the main results.

G.3 Lee Bounds

We follow Lee (2009) to construct nonparametric bounds on the ITT effects that are robust to selective attrition. The Lee bounds approach exploits random assignment to construct worst-case bounds on the treatment effect without imposing assumptions on the selection mechanism. The identifying assumption is monotonicity: treatment assignment affects sample selection in only one direction — that is, treatment either weakly increases or weakly decreases the probability of being observed at endline relative to control. Under this assumption, the true ITT lies within an identified interval $[\beta^L, \beta^U]$, where the bounds are obtained by trimming the outcome distribution of the group with lower attrition by the excess share of observed units.

Formally, let p^T and p^C denote the endline observation rates for the treatment and control groups respectively. For recording outcomes, $p^T > p^C$ in our sample, so we trim the treatment group by the proportion $(p^T - p^C)/p^T$. The lower bound β^L is obtained by dropping the $(p^T - p^C)/p^T$ fraction of treatment observations with the *highest* outcome values and recomputing the ITT; the upper bound β^U is obtained by dropping the same fraction with the *lowest* outcome values. Confidence intervals for the bounds are constructed using the bootstrap procedure described in Lee (2009).

For survey outcomes, attrition is approximately balanced across treatment and control groups, so the Lee bounds are narrow and point estimates fall comfortably within the identified interval. For recording outcomes, the control group exhibits weakly higher attrition, so the trimming proportion is small and the bounds remain informative. All significant ITT estimates — conversational turns, StimQ subscales, parental knowledge, parental beliefs, and the belief factor score — fall within their respective Lee bounds, confirming that selective attrition is unlikely to account for the observed treatment effects.

G.4 Results

G.4.1 TI-1: Beliefs and Knowledge

Table G5 reports the robustness of the TI-1 belief results to three approaches to selective attrition: inverse probability weighting (IPW), attrition balancing, and Lee bounds.

The baseline estimates in column (1) are robust across all specifications. Parental knowledge increases by 49.6% of a standard deviation in the baseline specification and remains virtually unchanged under IPW (49.8%) and attrition balancing (48.7%), with all estimates significant at the 1% level. The Lee bounds in columns (4) and (5) confirm robustness: the lower bound of 45.3% and upper bound of 75.4% are both statistically significant, indicating that the knowledge effect is positive and significant regardless of the direction of selective attrition.

Table G5: Attrition Robustness: TI-1 Beliefs

Dependent Var.	(1) No Correction	(2) IPW	(3) Balancing	(4) Lee Lower	(5) Lee Upper
Parental Knowledge	0.496***	0.498***	0.487***	0.453***	0.754***
	(0.097)	(0.095)	(0.106)	(0.153)	(0.152)
	[0.086]	[0.083]	[0.100]		
Parental Beliefs	0.694***	0.673***	0.736***	0.166	0.797***
	(0.118)	(0.117)	(0.127)	(0.150)	(0.155)
	[0.126]	[0.131]	[0.130]		
Factor Score	0.866***	0.856***	0.872***	0.702***	1.008***
	(0.112)	(0.111)	(0.127)	(0.143)	(0.151)
	[0.098]	[0.091]	[0.122]		

Notes: N = 238. Robust standard errors in (·) and clustered standard errors in [·]. Column (1) presents baseline results. Column (2) uses Inverse Probability Weighting (IPW). Column (3) utilizes balancing within cohort-language groups. Columns (4) and (5) report Lee Bounds. Probit models for IPW/Balancing control for cohorts, language groups, and StimQ scores. * p<0.10, ** p<0.05, *** p<0.01.

Parental beliefs increase by 69.4% of a standard deviation at baseline, with estimates of 67.3% under IPW and 73.6% under balancing, all significant at the 1% level. The Lee lower bound of 16.6% is not statistically significant, indicating that under the worst-case trimming assumption, the parental beliefs effect could be attenuated to near zero. The upper bound of 79.7% is significant, so the effect is bounded above by a large positive value. We note that the lower bound scenario requires extreme selection on unobservables and is inconsistent with the pattern of selective attrition documented in Appendix G, which is concentrated in the control group on English language and baseline StimQ rather than on outcomes.

The belief factor score is the most robust outcome. The baseline estimate of 86.6% of a standard deviation is stable under IPW (85.6%) and balancing (87.2%), and the Lee lower bound of 70.2% is statistically significant, confirming that even under the most conservative trimming assumption, the factor score effect remains large and precisely estimated. The upper bound of 100.8% indicates the effect could be as large as a full standard deviation under the opposite trimming scenario.

Taken together, the attrition robustness analysis strongly supports TI-1. The belief factor score is significant under all five specifications, including the Lee lower bound, providing the strongest available evidence that selective attrition does not account for the observed belief effects.

G.4.2 TI-2: Cue Detection and Responsive Engagement

Table G6 reports the robustness of the primary TI-2 outcome — conversational turns — to selective attrition. The baseline estimate of 30.0% of a standard deviation is stable under IPW (30.5%, significant at the 1% level) and remains positive and significant under attrition balancing (28.0%, significant at the 5% level), confirming that the conversational turns effect is not driven by differential attrition on observable characteristics.

Table G6: Attrition Robustness: TI-2 Responsive Engagement

Dependent Var.	(1) No Correction	(2) IPW	(3) Balancing	(4) Lee Lower	(5) Lee Upper
Conversational Turns	0.300***	0.305***	0.280**	0.083	0.617***
	(0.099)	(0.102)	(0.123)	(0.187)	(0.178)
	[0.068]	[0.068]	[0.093]		

Notes: N = 203. Robust standard errors in (·) and clustered standard errors in [·]. Column (1) presents baseline results. Column (2) uses Inverse Probability Weighting (IPW). Column (3) utilizes balancing within cohort-language groups. Columns (4) and (5) report Lee Bounds. Probit models for IPW/Balancing control for cohorts, language groups, and StimQ. * p<0.10, ** p<0.05, *** p<0.01.

The Lee bounds reveal that the conversational turns effect is sensitive to the direction of worst-case selection on unobservables. The lower bound of 8.3% is not statistically significant, indicating that under the most conservative trimming assumption the effect could be attenuated below conventional significance thresholds. The upper bound of 61.7% is large and precisely estimated. As with the parental beliefs lower bound, we note that the worst-case trimming scenario is inconsistent with the pattern of attrition documented in Appendix G: recording attrition is higher in the control group, which means the Lee bounds trim the treatment group, and the selective attrition on observable characteristics is concentrated in the control group on English language and baseline StimQ rather than on conversational turns outcomes. The IPW and balancing estimates, which account for selection on observables, are both positive and significant, and the covariate-conditioned Lee bounds would be tighter than those reported here.

Table G7 reports the robustness of the child-initiated ADEX outcomes to selective attrition. The pattern across all four outcomes is consistent: baseline estimates are positive and significant, IPW estimates are positive and equally significant, attrition balancing estimates are positive and significant, and Lee lower bounds are uniformly insignificant while upper bounds are large and precisely estimated.

Conversational block count increases by 26.2% of a standard deviation at baseline and 25.5% under IPW, both significant at the 5% level. Under attrition balancing the estimate increases to 30.1% and maintains significance. Conversational turn count increases by 30.8% at baseline, 31.2% under IPW, and 33.0% under balancing, with the first two significant at the 1% level and the third at the 5% level. Block duration increases by 34.7% at baseline, 36.4% under IPW, and 32.8% under balancing, all significant at the 1% or 5% level. Adult word count in child-initiated blocks increases by 33.9% at baseline, 35.0% under IPW (significant at the 1% level), and 36.9% under balancing (significant at the 1% level).

Table G7: Robustness to Attrition: Child-Initiated Conversation Metrics

Dependent Var.	(1) No Correction	(2) IPW	(3) Balancing	(4) Lee Lower	(5) Lee Upper
Conv. Block Count	0.262**	0.255**	0.301**	0.011	0.577***
	(0.123)	(0.128)	(0.145)	(0.177)	(0.173)
	[0.102]	[0.122]	[0.109]		
Conv. Turn Count	0.308***	0.312***	0.330**	0.052	0.608***
	(0.107)	(0.109)	(0.129)	(0.181)	(0.172)
	[0.065]	[0.063]	[0.089]		
Block Duration	0.347***	0.364***	0.328**	0.072	0.620***
	(0.107)	(0.107)	(0.129)	(0.167)	(0.172)
	[0.079]	[0.079]	[0.085]		
Adult Word Count	0.339**	0.350***	0.369***	0.104	0.650***
	(0.110)	(0.113)	(0.136)	(0.167)	(0.177)
	[0.081]	[0.073]	[0.105]		

Notes: Robust standard errors in (·) and clustered standard errors in [·]. All specifications include full controls for demographics, cohort, and language. IPW and Balancing models include baseline StimQ. * p<0.10, ** p<0.05, *** p<0.01.

The Lee lower bounds are uniformly statistically insignificant across all four outcomes, ranging from 1.1% to 10.4% of a standard deviation. As with the conversational turns result, this reflects the mechanical consequence of trimming the treatment group under worst-case selection on unobservables, which is inconsistent with the pattern of attrition documented in Appendix G. The IPW and balancing estimates, which account for selection on the observable characteristics that drive attrition in this sample, are consistently positive and significant, providing stronger evidence against an attrition-driven explanation of the child-initiated block effects. The concentration of treatment effects across all four child-initiated metrics — block count, turn count, duration, and adult words — is robust to the two observables-based corrections and is therefore unlikely to reflect differential attrition on unobservables.

Table G8 reports the robustness of the female adult-initiated ADEX outcomes to selective attrition. The pattern differs from the child-initiated results in one important respect: conversational block count is not significant in the baseline, IPW, or balancing specifications, while the remaining three outcomes — turn count, block duration, and adult word count — are consistently positive and significant across all observables-based corrections.

Conversational block count shows a baseline estimate of 10.5% of a standard deviation that is not statistically significant under any specification, including IPW (9.0%) and balancing (17.5%). The Lee upper bound of 42.3% is significant at the 5% level, suggesting that under favorable selection assumptions a positive effect could emerge, but the central estimates provide no evidence of a significant effect on the number of female adult-initiated blocks.

Conversational turn count increases by 26.7% of a standard deviation at baseline and 27.4% under IPW, both significant at the 5% level, rising to 31.8% under balancing, also significant

Table G8: Robustness to Attrition: Female Adult-Initiated Conversation Metrics

Dependent Var.	(1) No Correction	(2) IPW	(3) Balancing	(4) Lee Lower	(5) Lee Upper
Conv. Block Count	0.105	0.090	0.175	-0.092	0.423**
	(0.121)	(0.122)	(0.144)	(0.175)	(0.178)
	[0.114]	[0.119]	[0.127]		
Conv. Turn Count	0.267**	0.274**	0.318**	0.037	0.543***
	(0.108)	(0.109)	(0.132)	(0.176)	(0.173)
	[0.077]	[0.076]	[0.102]		
Block Duration	0.271**	0.281**	0.305**	0.016	0.528***
	(0.120)	(0.121)	(0.150)	(0.169)	(0.176)
	[0.097]	[0.098]	[0.108]		
Adult Word Count	0.274**	0.283**	0.322**	-0.011	0.508***
	(0.126)	(0.127)	(0.159)	(0.172)	(0.174)
	[0.099]	[0.101]	[0.117]		

Notes: Robust standard errors in (·) and clustered standard errors in [·]. All specifications include full controls. IPW and Balancing models include baseline StimQ. * p<0.10, ** p<0.05, *** p<0.01.

at the 5% level. Block duration increases by 27.1% at baseline, 28.1% under IPW, and 30.5% under balancing, all significant at the 5% level. Adult word count in female adult-initiated blocks increases by 27.4% at baseline, 28.3% under IPW, and 32.2% under balancing, all significant at the 5% level. The stability of these three estimates across the observables-based corrections is notable: balancing produces slightly larger estimates than the baseline, suggesting that differential attrition on observables if anything attenuates the baseline estimates rather than inflating them.

The Lee lower bounds are uniformly insignificant across all four outcomes, and two are negative, reflecting the mechanical consequence of worst-case trimming. As with the child-initiated outcomes, the IPW and balancing estimates provide the more relevant robustness check given the pattern of attrition documented in Appendix G. The three significant female adult-initiated outcomes — turn count, duration, and adult words — are robust to both observables-based corrections, consistent with the interpretation that the intervention raised the quality and intensity of female adult engagement beyond what is captured by block count alone.

Table G9 reports the robustness of the male adult-initiated ADEX outcomes to selective attrition. Consistent with the main results in Section 6.2.2, all four outcomes are statistically insignificant across all observables-based specifications, confirming that the null result for male adult-initiated blocks is not an artifact of differential attrition.

Conversational block count, turn count, and block duration show positive point estimates ranging from 12.3% to 24.4% of a standard deviation across the baseline, IPW, and balancing specifications, but none reaches conventional significance levels under robust or clustered standard errors. Adult word count in male adult-initiated blocks is the smallest and least precisely estimated effect, ranging from 3.3% in the IPW specification to 13.3% under balancing, and is insignificant under all specifications.

Table G9: Robustness to Attrition: Male Adult-Initiated Conversation Metrics

Dependent Var.	(1) No Correction	(2) IPW	(3) Balancing	(4) Lee Lower	(5) Lee Upper
Conv. Block Count	0.160	0.158	0.244	-0.152	0.379**
	(0.134)	(0.153)	(0.158)	(0.161)	(0.174)
	[0.114]	[0.122]	[0.176]		
Conv. Turn Count	0.170	0.175	0.183	-0.151	0.369**
	(0.129)	(0.146)	(0.155)	(0.170)	(0.169)
	[0.112]	[0.111]	[0.168]		
Block Duration	0.135	0.123	0.196	-0.194	0.305*
	(0.138)	(0.161)	(0.170)	(0.169)	(0.171)
	[0.113]	[0.114]	[0.158]		
Adult Word Count	0.048	0.033	0.133	-0.270	0.192
	(0.143)	(0.166)	(0.176)	(0.165)	(0.167)
	[0.096]	[0.093]	[0.146]		

Notes: Robust standard errors in (·) and clustered standard errors in [·]. All specifications include full controls. IPW and Balancing models include baseline StimQ. * p<0.10, ** p<0.05, *** p<0.01.

The Lee lower bounds are negative for all four outcomes, ranging from −15.2% to −27.0% of a standard deviation, and are uniformly insignificant. The upper bounds are positive and significant for block count (37.9%, $p < 0.05$), turn count (36.9%, $p < 0.05$), and block duration (30.5%, $p < 0.10$), but not for adult word count. The significant upper bounds indicate that under favorable selection assumptions a positive effect on male adult-initiated engagement could emerge, but the central estimates and lower bounds provide no evidence that such an effect is present in the data.

The uniform insignificance of the male adult-initiated outcomes across all five specifications is consistent with the model's predictions and with the program design: no fathers were enrolled in LENA Start, so the belief and cue-detection expertise channels were not directly activated for male adults. The positive but imprecise point estimates are consistent with a small positive spillover from treated mothers to other household members, but the data are underpowered to detect such effects given the low baseline levels of male adult engagement documented in Table D2.

G.4.3 TI-3: Behavioral Versus Material Components of the Home Environment

Table G10 reports the robustness of the TI-3 StimQ results to selective attrition. The pattern is strikingly consistent: the three behavioral subscales — READ, PVR, and PIDA — are significant across all five specifications, while the materials subscale ALM becomes marginally significant under IPW (10) and significant at the 10% level with balancing.

The StimQ total score increases by 58.3% of a standard deviation at baseline and is stable under IPW (57.9%) and balancing (66.4%), all significant at the 1% level. The Lee lower and upper bounds of 56.6% and 78.0% are both statistically significant, confirming that the total StimQ effect is robustly positive regardless of the direction of selection on unobservables.

Table G10: Attrition Robustness: TI-3 StimQ Components

Dependent Var.	(1) No Correction	(2) IPW	(3) Balancing	(4) Lee Lower	(5) Lee Upper
StimQ Total Score	0.583***	0.579***	0.664***	0.566***	0.780***
	(0.093)	(0.092)	(0.104)	(0.143)	(0.150)
	[0.090]	[0.088]	[0.115]		
READ	0.708***	0.708***	0.778***	0.662***	0.859***
	(0.100)	(0.101)	(0.108)	(0.143)	(0.152)
	[0.122]	[0.122]	[0.153]		
PIDA	0.282***	0.274***	0.339***	0.244*	0.462***
	(0.104)	(0.100)	(0.115)	(0.141)	(0.157)
	[0.102]	[0.099]	[0.103]		
PVR	0.456***	0.451***	0.451***	0.288*	0.671***
	(0.105)	(0.102)	(0.119)	(0.150)	(0.165)
	[0.089]	[0.087]	[0.104]		
ALM	0.183	0.183*	0.235**	0.111	0.491***
	(0.111)	(0.108)	(0.119)	(0.141)	(0.154)
	[0.104]	[0.101]	[0.103]		

Notes: $N = 239$. Robust standard errors in $(\cdot)$ and clustered standard errors in $[\cdot]$. Column (1) presents baseline results. Column (2) uses Inverse Probability Weighting (IPW). Column (3) utilizes balancing within cohort-language groups. Columns (4) and (5) report Lee Bounds. Probit models for IPW/Balancing control for cohorts, language groups, and baseline StimQ scores. READ: Reading; PIDA: Parental Involvement in Developmental Advance; PVR: Parental Verbal Responsivity; ALM: Availability of Learning Materials. * p<0.10, ** p<0.05, *** p<0.01.

Among the behavioral subscales, READ is the most robust outcome. The baseline estimate of 70.8% of a standard deviation is virtually unchanged under IPW (70.8%) and rises to 77.8% under balancing, all significant at the 1% level. The Lee lower bound of 66.2% is significant at the 1% level, indicating that even under the most conservative trimming assumption the READ effect remains large and precisely estimated. PVR increases by 45.6% at baseline, 45.1% under IPW, and 45.1% under balancing, all significant at the 1% level. The Lee lower bound of 28.8% is marginally significant at the 10% level. PIDA increases by 28.2% at baseline and is stable under IPW (27.4%) and balancing (33.9%), all significant at the 1% level. The Lee lower bound of 24.44% is marginally significant ($p < 0.10$), confirming that the PIDA effect survives even under conservative trimming.

The ALM subscale shows a baseline estimate of 18.3% of a standard deviation that is not statistically significant, consistent with the null prediction for the materials component under TI-3. ALM becomes marginally significant under IPW (18.3%, $p < 0.10$) and significant under balancing (23.5%, $p < 0.05$). The Lee lower bound of 11.1% is not significant. The ALM results under IPW and balancing are sensitive to the specification of the probit model used to construct the weights and should therefore be interpreted with caution. The within-instrument falsification result is broadly preserved: the behavioral subscales respond consistently and significantly across all corrections, while the ALM result is fragile and does not survive all specifications.

G.4.4 TI-4: Null Effects on Preferences and Self-Efficacy

Table G11 reports the robustness of the TI-4 placebo outcomes to selective attrition. The overall pattern is consistent with TI-4: placebo outcomes do not show a robust positive effect across all specifications, and the results that approach significance in some columns are fragile across the attrition corrections. The only exception is Negative Reciprocity, which already suffers from low response rates in the survey, making any interpretations uncertain.

Parenting efficacy shows a marginally significant baseline estimate of 17.5% of a standard deviation ($p < 0.10$ under robust standard errors) that maintains marginal significance under IPW (17.3%) and balancing (20.9%). The Lee lower bound of 14.3% is not significant. The result does not show strong significance or magnitude across the alternative specifications, thus we cannot reject the null prediction for self-efficacy under TI-4.

Table G11: Attrition Robustness: TI-4 Placebo Outcomes

Dependent Var.	(1) No Correction	(2) IPW	(3) Balancing	(4) Lee Lower	(5) Lee Upper
Parenting Efficacy	0.175*	0.173*	0.209*	0.143	0.359**
	(0.096)	(0.095)	(0.110)	(0.152)	(0.163)
	[0.091]	[0.091]	[0.096]		
Parenting Support	0.168	0.150	0.171	0.246	0.446***
	(0.101)	(0.103)	(0.111)	(0.163)	(0.162)
	[0.100]	[0.095]	[0.134]		
Altruism	0.186	0.164	0.312**	0.123	0.526***
	(0.125)	(0.123)	(0.144)	(0.163)	(0.150)
	[0.139]	[0.129]	[0.119]		
Trust	0.143	0.143	0.100	0.024	0.307*
	(0.117)	(0.119)	(0.134)	(0.173)	(0.186)
	[0.116]	[0.121]	[0.114]		
Pos. Reciprocity	0.093	0.083	0.068	-0.001	0.197
	(0.134)	(0.131)	(0.165)	(0.155)	(0.181)
	[0.169]	[0.165]	[0.181]		
Neg. Reciprocity	-0.276**	-0.268**	-0.284*	-0.604***	0.029
	(0.126)	(0.123)	(0.168)	(0.201)	(0.203)
	[0.084]	[0.095]	[0.129]		
Patience	0.004	0.007	0.100	-0.015	0.373**
	(0.129)	(0.128)	(0.143)	(0.147)	(0.175)
	[0.125]	[0.130]	[0.151]		

Notes: Robust standard errors in $(\cdot)$ and clustered standard errors in $[\cdot]$. Column (1) utilizes Model 3 specifications. Lee bounds are estimated without covariates. N varies by outcome: Parenting Efficacy/Support (239), Altruism (202), Trust (223), Pos. Reciprocity (220), Neg. Reciprocity (178), Patience (231). * p<0.10, ** p<0.05, *** p<0.01.

Parenting social support shows a baseline estimate of 16.8%, IPW estimate of 15.0%, and balancing estimate of 17.1%, all of which are insignificant. The Lee lower bound of 24.6% is not significant. Altruism is insignificant at baseline (18.6%) and under IPW (16.4%), but becomes significant under balancing (31.2%, $p < 0.05$). The Lee lower bound of 12.3% is not significant. This isolated result under one specification, for an outcome with a relatively small sample ($N = 202$) and large standard errors, does not constitute evidence against TI-4.

Trust, positive reciprocity, and patience show no significant effects under any specification. Trust and positive reciprocity are uniformly small and insignificant. Patience is essentially zero at baseline and remains so across all corrections.

Negative reciprocity is the most anomalous result in the table. The baseline estimate of −27.6% is significant at the 5% level, dropping to −26.8% under IPW, and growing under balancing (−28.4%) (both significant at the 5% level). The Lee lower bound of −60.4% is significant at the 1% level, indicating that under worst-case trimming the negative effect is large. A negative effect on negative reciprocity — reduced willingness to punish unfairness — is difficult to interpret as a model-consistent finding, since the model predicts no effects on global preferences in either direction. The result has the smallest sample ($N = 178$), and is the only outcome for which the Lee lower bound is more extreme than the baseline estimate. We therefore treat this as a spurious result driven by the high rate of missing data for this outcome rather than as evidence of a genuine program effect on preferences.

H Heterogeneity Analysis

To investigate treatment effect heterogeneity, we estimate Conditional Average Treatment Effects (CATEs) for conversational turns using Generalized Random Forests (Athey, Tibshirani, and Wager 2019). The causal forest is trained on the full set of covariates from the most comprehensive model specification, including demographic characteristics of the child and primary family member, cohort and language indicators, and recording variables.

H.1 Subgroup Analysis

Figures H1 and H2 report CATEs for high and low realizations of the training covariates and covariates of interest, respectively. Across all subgroups, the confidence intervals for high and low subgroups overlap substantially, providing no evidence of heterogeneous treatment effects.

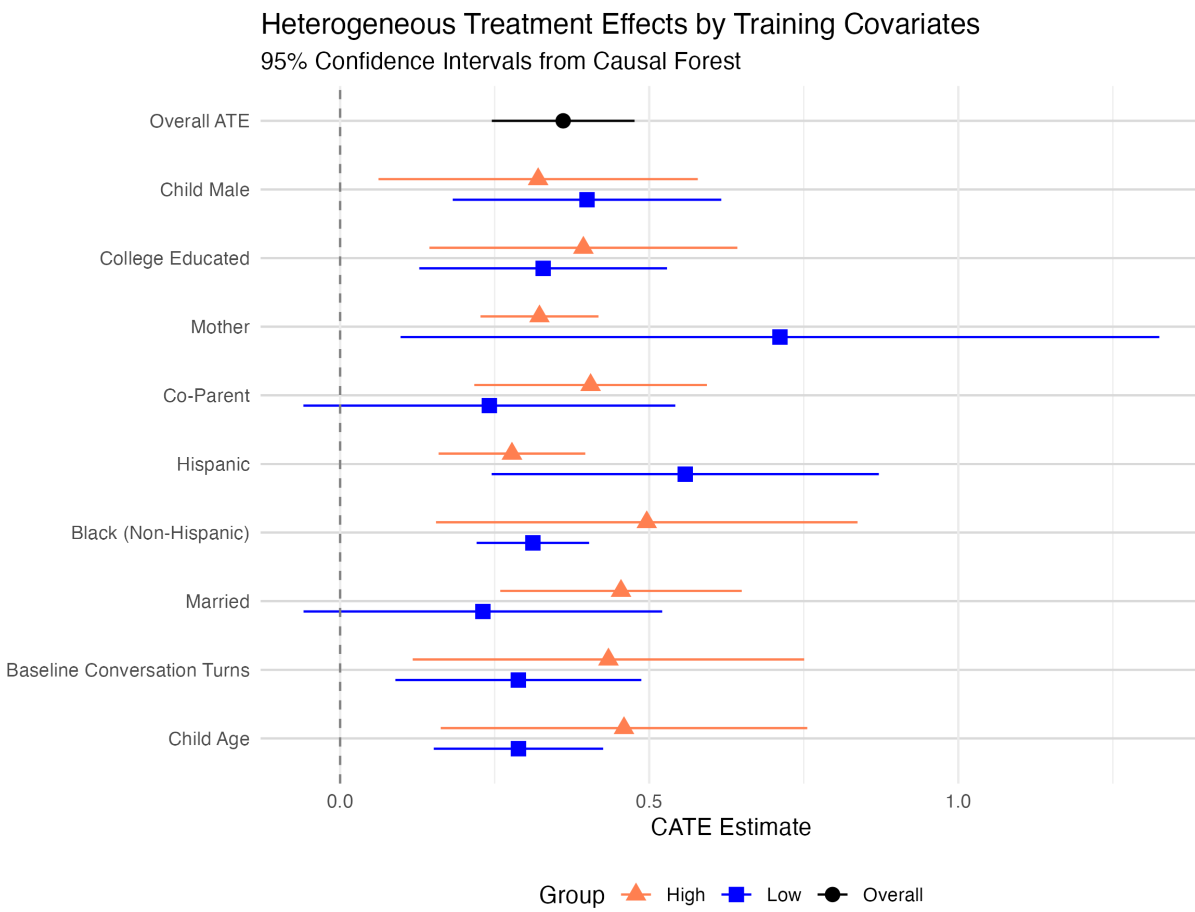

Figure H1: Heterogeneity: Training Covariates

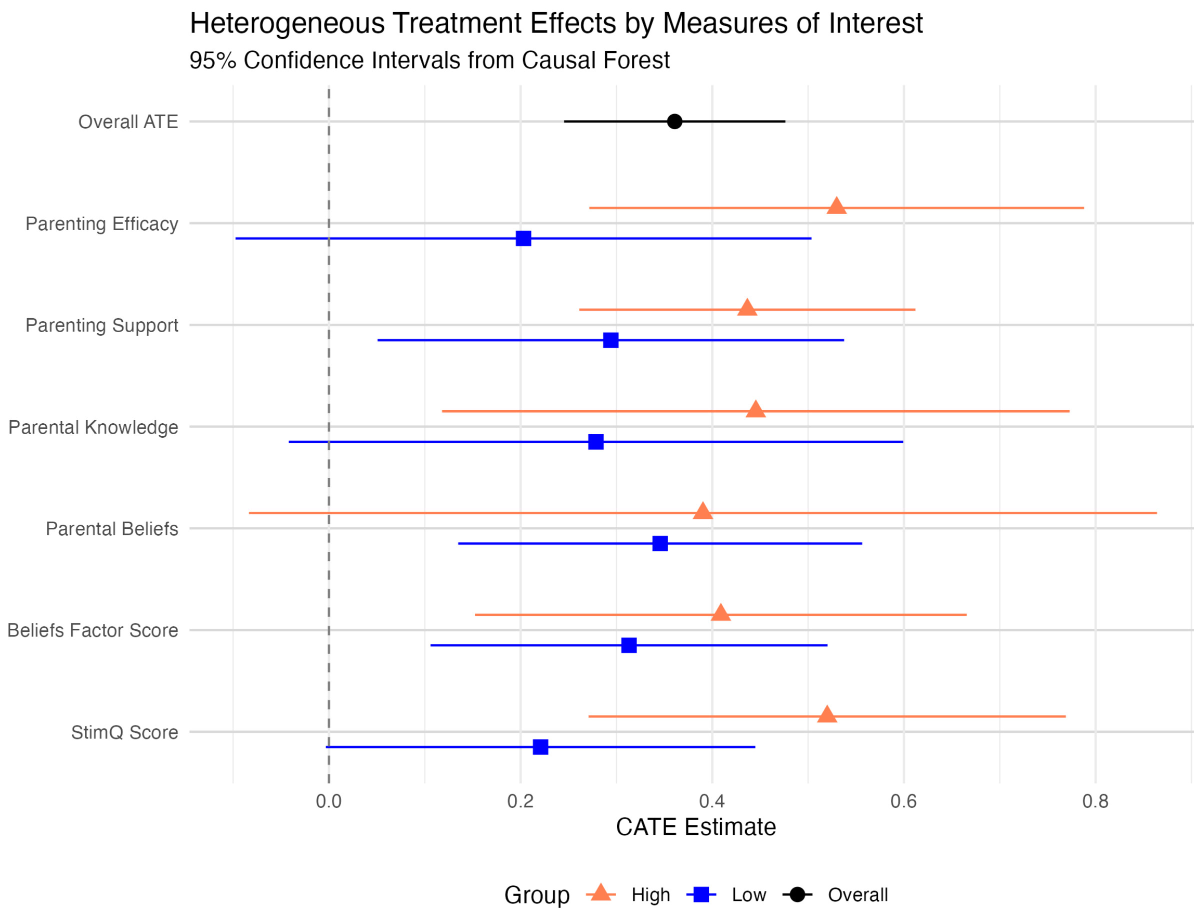

Figure H2: Heterogeneity: Covariates of Interest

H.2 Continuous Variables

Extending the subgroup analysis, we investigate CATEs against continuous measures of baseline variables of interest. Since the causal forests produce individual CATE estimates, we are able to plot them against baseline measures. Figure H3 are generated using local polynomial regression to fit the data, and the band represents the standard error of the fit at each point. The variables show a fairly horizontal CATE across values of the baseline variable. We therefore find no evidence of heterogeneous treatment effects even from this more granular analysis.

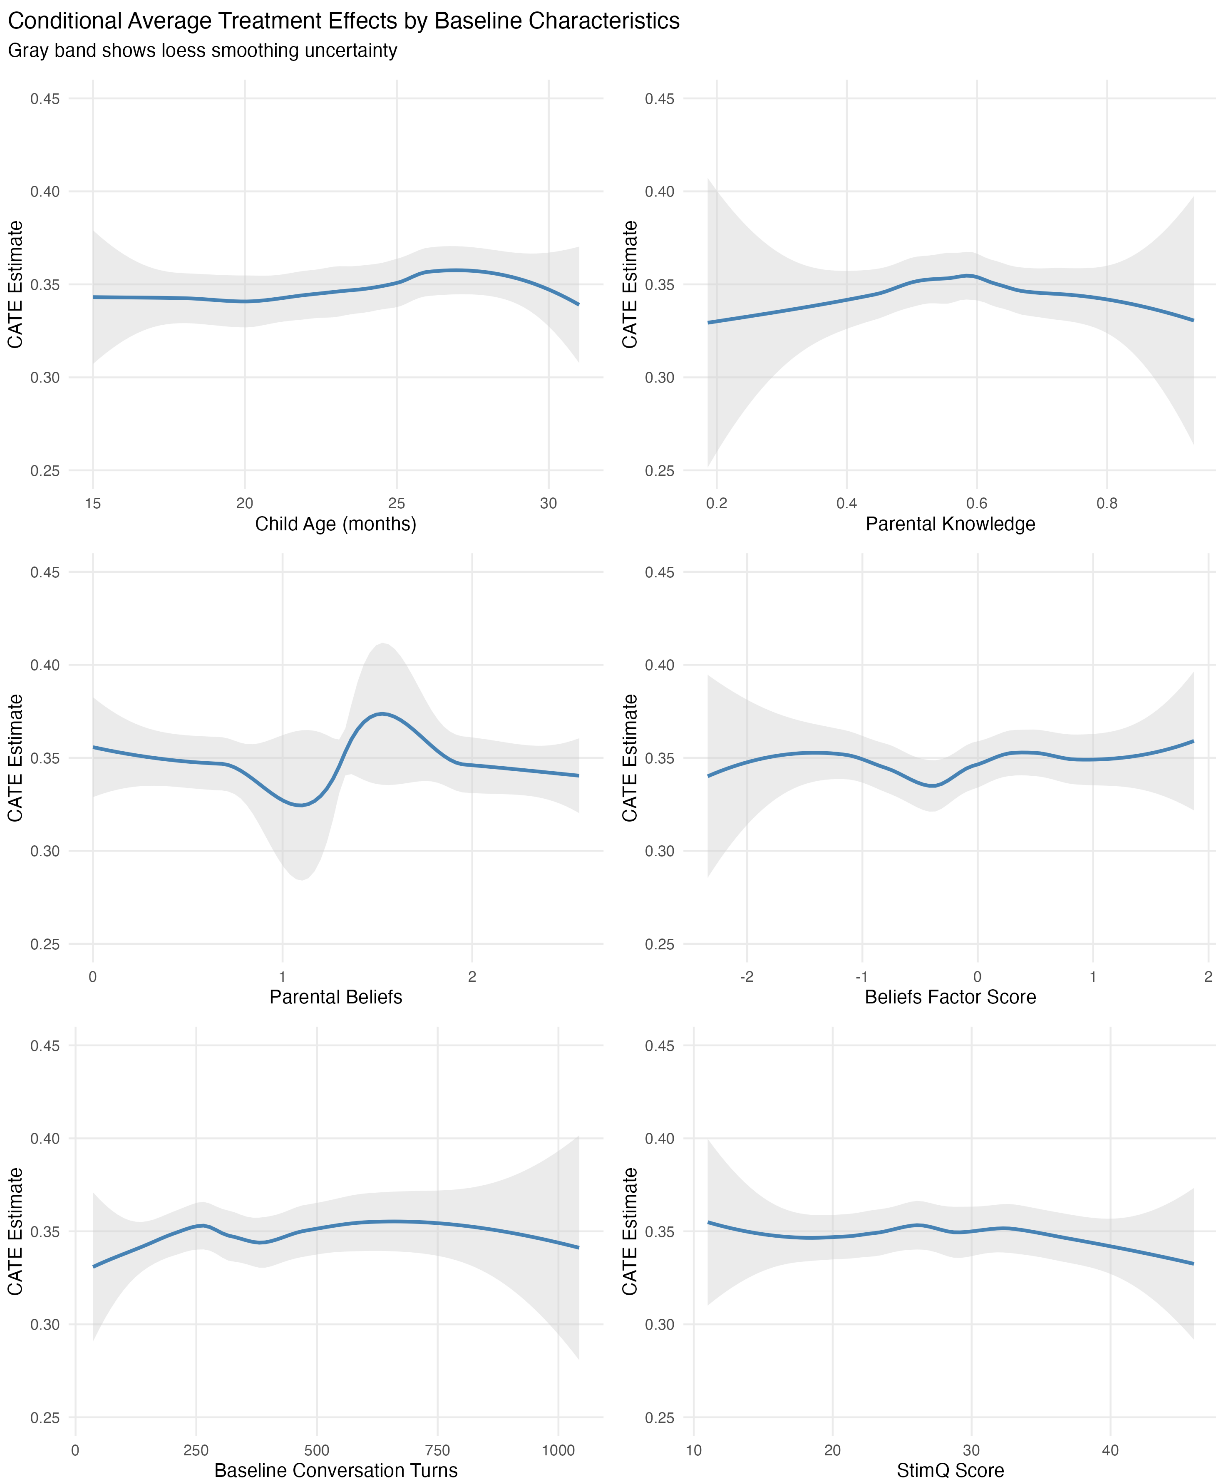

Figure H3: CATEs Against Continuous Baseline Variables