Fixed Budget is No Harder Than Fixed Confidence

in Best-Arm Identification, up to logarithmic factors

Kapilan Balagopalan[*,1] Yinan Li[*,1] Yao Zhao[1] Tuan Ngo Nguyen[1]
Anton Daitche[2] Houssam Nassif[2] Kwang-Sung Jun[3]
[1]University of Arizona [2]Meta Inc. [3]POSTECH CSE/GSAI *Equal contribution

Scan for paper

1. Problem: two protocols for BAI

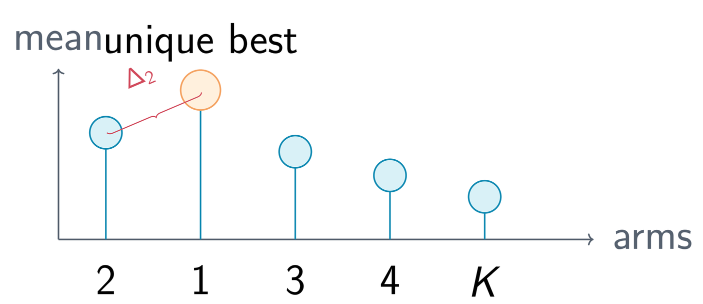

Given K arms with unknown means μ_i, identify the unique best arm. The performance question is: how many samples are needed to make $\Pr(\widehat{J} \neq 1)$ small?

Fixed confidence (FC)	Fixed budget (FB)
Input: failure rate δ	Input: budget B
Sample until stopping rule certifies answer	Must stop after at most B pulls
Guarantee: $\Pr(\widehat{J} \neq 1) \leq \delta$	Error decays as B grows
Sample complexity: stopping time τ	Sample complexity: budget for target error

2. Research question

Is fixed budget harder/easier than fixed confidence for generic BAI with a unique best arm?

Background. In ordinary K-armed BAI, optimal FC and FB sample complexities match up to logarithmic factors. For generic, potentially structured BAI, the answer is largely unknown. **Answer.** FC2FB constructs an FB algorithm from a strong FC algorithm with only logarithmic overhead. Thus, **FB is no harder than FC up to logarithmic factors** in sample complexity.

3. Contributions

1. A black-box reduction FC2FB from fixed confidence to fixed budget.
2. A weak-to-strong conversion FCW2S that broadens the usable FC algorithm inputs.
3. Improved FB guarantees in heterogeneous-noise, linear, unimodal, and cascading BAI.
4. Experiments showing FC2FB(PE-KHN) can outperform SH and SHVar.

4. Input guarantee needed

FC2FB needs a strong FC input. A weak FC algorithm can first be amplified by FCW2S.

Weak FC
For **some fixed** $\delta_0 \in (0,1)$,
$$\Pr(\widehat{J} \neq 1, \tau < \infty) \leq \delta_0,$$
$$\Pr(\tau > f(\delta_0)) \leq \delta_0.$$

Strong FC
For **every** $\delta \in (0,1)$,
$$\Pr(\widehat{J} \neq 1, \tau < \infty) \leq \delta,$$
$$\Pr\left(\tau > H_{\mathrm{FC}} \log\frac{1}{\delta} + C\right) \leq \delta.$$

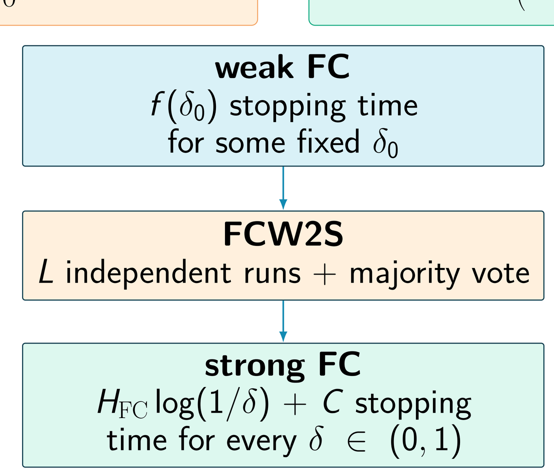

Note: FC2FB does *not* need to know H_{FC} or C. FCW2S does *not* need to know $f(\delta_0)$.

5. The reduction: FC2FB

Set
$$R = \lfloor \log_2(B/Q) \rfloor, \qquad B' = \lfloor B/R \rfloor, \qquad L_r = 2^{R-r}.$$
At stage r, run the base FC algorithm $\mathcal{A}(\delta_0^{L_r})$ with budget cap B'. If it self-terminates, return its recommendation; otherwise force-terminate and move to the next stage.

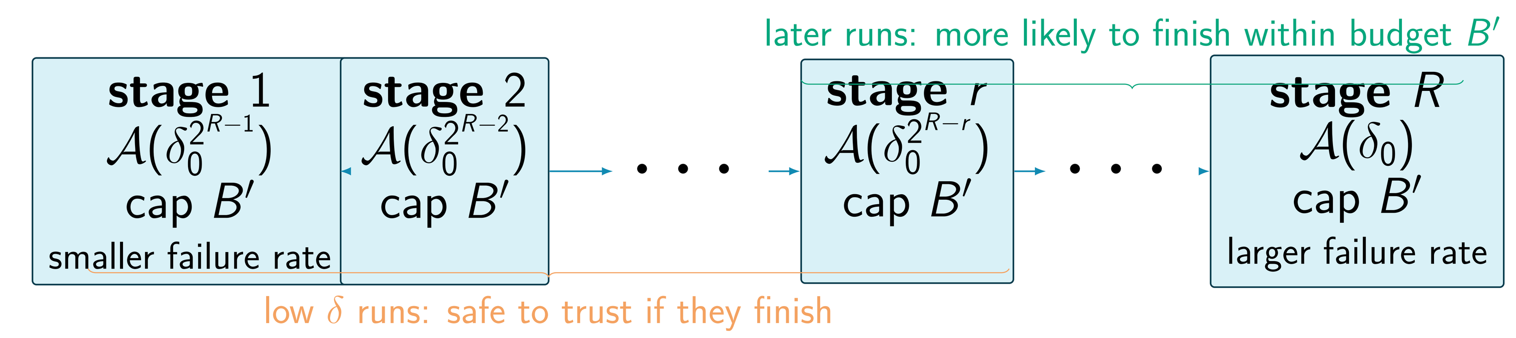

Try doubly exponentially increasing failure rates. Trust first self-terminated run.

6. Main theorem

Theorem, simplified. If $\mathcal{A}$ is strong FC with high-probability stopping time
$$\tau \lesssim H_{\mathrm{FC}} \log(1/\delta) + C,$$
then $\mathrm{FC2FB}(\mathcal{A})$ satisfies an error bound of the form
$$\Pr(\widehat{J} \neq 1) \leq 3 \exp\left(-\frac{B}{4Q/\log(1/\delta_0) + 4H_{\mathrm{FC}} \log_2(B/Q)}\right).$$
Equivalently, its sample complexity is
$$\boxed{B = \widetilde{O}\left(H_{\mathrm{FC}} \log\frac{1}{\delta} + C\right)}$$
without knowing H_{FC} or C.

A more explicit sample complexity bound, suppressing only constants,
$$B = O\left(Q \log\frac{1}{\delta} + H_{\mathrm{FC}} \log\frac{1}{\delta} \cdot \log\frac{H_{\mathrm{FC}} \log(1/\delta)}{Q} + C\right).$$

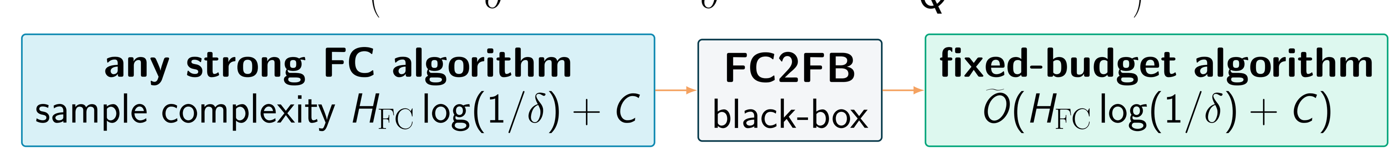

7. Proof intuition

1. **Early finishes are accurate.** A run with tiny input $\delta_0^{L_r}$ is unlikely to output the wrong arm.
2. **A threshold stage.** Define
$$r^* := \min\{r : B' \geq T^*_{\delta_0^{L_r}}\}.$$
r^* is *analysis-only*: it marks the first stage whose budget is large enough for the base FC stopping-time guarantee.
3. **Failure event decomposition.**
$$\Pr(\widehat{J} \neq 1) \leq \Pr(\text{wrong self-terminated run before } r^*) + \Pr(\text{no self-termination by } r^*).$$
The first term is small because early stages use tiny failure rates. The second is small because stage r^* has enough budget to self-terminate with high probability.

8. Practical parameter choices

Use a constant base failure rate such as $\delta_0 = 1/e$. Choose Q as a minimal meaningful per-stage budget: for unstructured BAI, $Q \approx K$; for linear bandits, $Q \approx d$.

Black-box reduction: H_{FC} and C stay unknown.

9. Applications: our bound vs. existing FB

All entries are sample complexities for achieving $\Pr(\widehat{J} \neq 1) \leq \delta$, ignoring logarithmic factors unless shown.

Setting	Existing fixed-budget bound	Our bound via FC2FB
Known heterogeneous noise	SHVar [Lalitha+23]: $\widetilde{O}\left(\frac{\sigma_1^2}{\Delta_2^2} + \sum_{i=2}^K \frac{\sigma_i^2}{\Delta_2^2}\right)$ SH [Karnin+13]: $\widetilde{O}\left(\frac{\sigma_1^2}{\Delta_2^2} + \sum_{i=2}^K \frac{\sigma_1^2+\sigma_i^2}{\Delta_i^2}\right)$	$\widetilde{O}\left(\frac{\sigma_1^2}{\Delta_2^2} + \sum_{i=2}^K \frac{\sigma_i^2}{\Delta_i^2}\right)$
Unknown variances	VBR [Faella+20]: $\widetilde{O}\left(K \max_{i\geq 2} \frac{\sigma_1^2+\sigma_i^2}{\Delta_i^2}\right)$ SHAdaVar [Lalitha+23]: $\widetilde{O}\left(K \frac{\sigma_{\max}^2}{\Delta_2^2}\right)$	[VD-BESTARMID: Lu+21] $\widetilde{O}\left(\sum_{i=1}^K \left(\frac{\sigma_i^2}{\Delta_i^2} + \frac{1}{\Delta_i}\right)\right)$
Linear bandits	[Katz-Samuels+20] $\widetilde{O}(\gamma^* \log(1/\delta) + \rho^* \log(1/\delta) + d)$	[FC Peace: Katz-Samuels+20] $\widetilde{O}(\gamma^* + \rho^* \log(1/\delta) + d)$
Unimodal bandits	[Ghosh+24] $\widetilde{O}(\Delta^{-2} \log(1/\delta))$, $\Delta := \min_i \lvert\mu_i - \mu_{i-1}\rvert$	[UniTT: Poiani+24] $\widetilde{O}((T_\mu(\delta_1) + K) \log(1/\delta))$ uses local complexity near the optimum, not the global smallest adjacent gap
Cascading bandits	None	[CASCADEBAI: Zhong+20] $\widetilde{O}(A_{\mathrm{cas}} \log(1/\delta) + C_{\mathrm{cas}})$ gives first FB guarantee from existing FC guarantees

10. Experiments

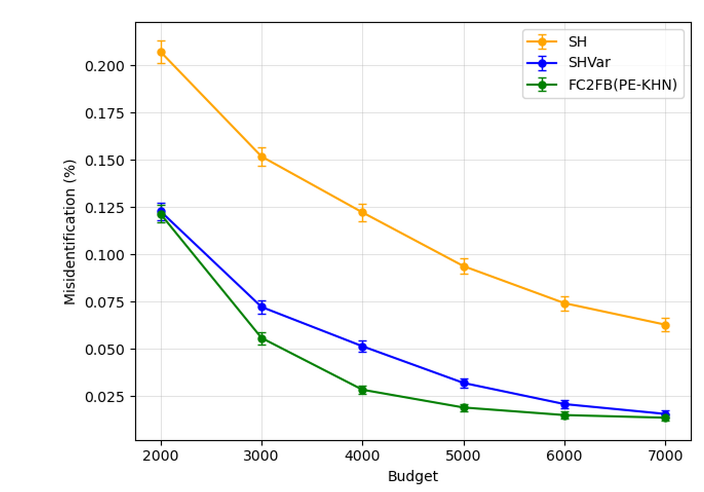

Error vs. budget, $K = 32$.

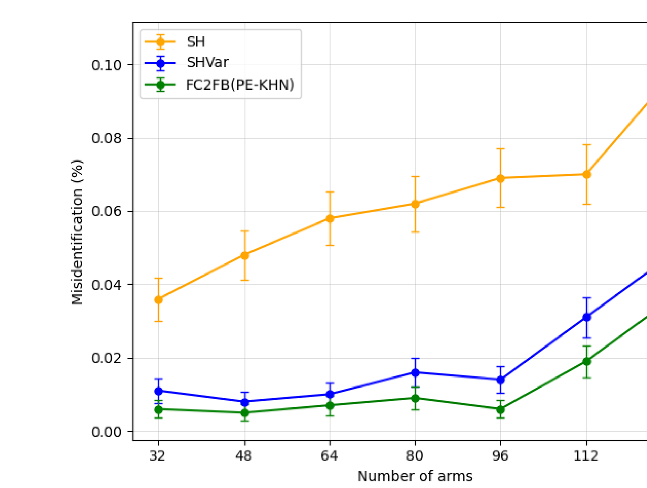

Error vs. number of arms, $B = 6000$.

Gaussian BAI with heterogeneous variances. FC2FB(PE-KHN) stays competitive across budgets and as K grows.

11. Future research problems

The reduction gives a general FC $\Rightarrow$ FB transfer, but several questions remain open.

1. Beyond high-probability identification
Can similar FC–FB relationships be proved for other performance measures, such as **simple regret**?

2. No unique best arm: ϵ-good-arm guarantees
When the best arm is not unique, can a meta-algorithm convert an FC algorithm with an ε-good-arm guarantee into an FB algorithm with a guarantee for **all** ε **simultaneously**?

3. More practical versions with samples shared across runs
FC2FB runs independent copies of the base FC algorithm. Can we design variants that **share samples across runs** and improve small-budget performance?

Main takeaway
Any good FC BAI algorithm $\xrightarrow{\text{FC2FB}}$ an almost-as-good FB BAI algorithm.

References: Audibert et al. (2010); Karnin et al. (2013); Katz-Samuels et al. (2020); Faella et al. (2020); Lalitha et al. (2023); Lu et al. (2021); Ghosh et al. (2024); Poiani et al. (2024); Zhong et al. (2020).