Cosmic Bandits: Exploration versus Exploitation in Cosmological Surveys

Ely D. Kovetz
University of Texas at Austin

ITC Seminar, Dec. 10th, 2013

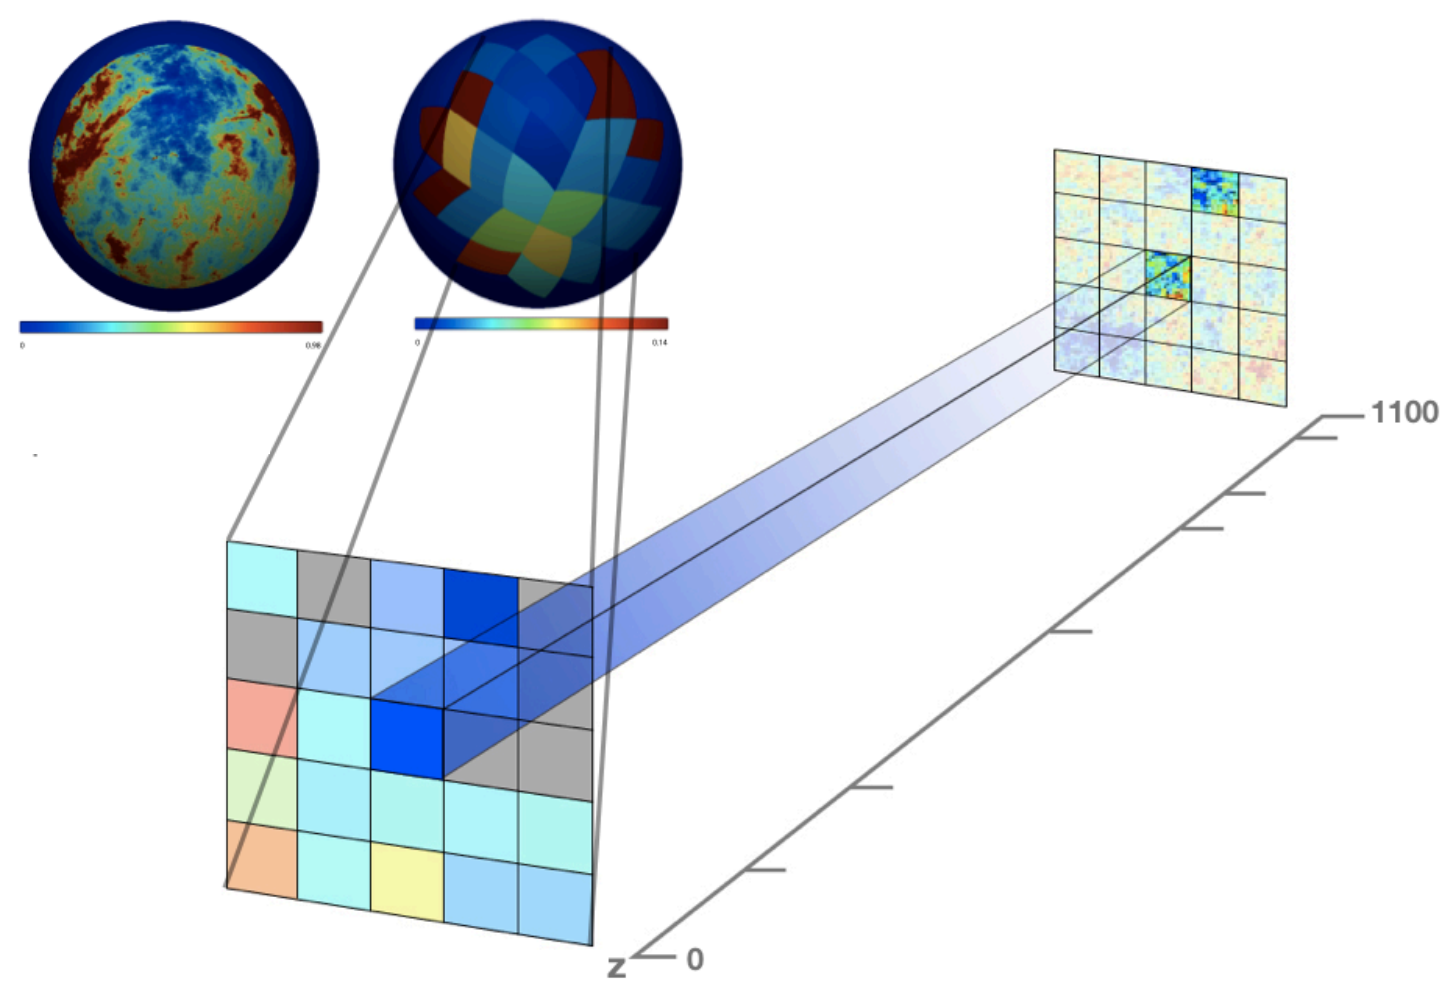

Based on: arXiv:1308.1404
and current work
with M. Kamionkowski (JHU)

Introduction: Exploration vs. Exploitation

Introduction: Exploration vs. Exploitation

- How to balance the tradeoff between *exploring* and *exploiting*?

Introduction: Exploration vs. Exploitation

- How to balance the tradeoff between *exploring* and *exploiting*?

  One way...

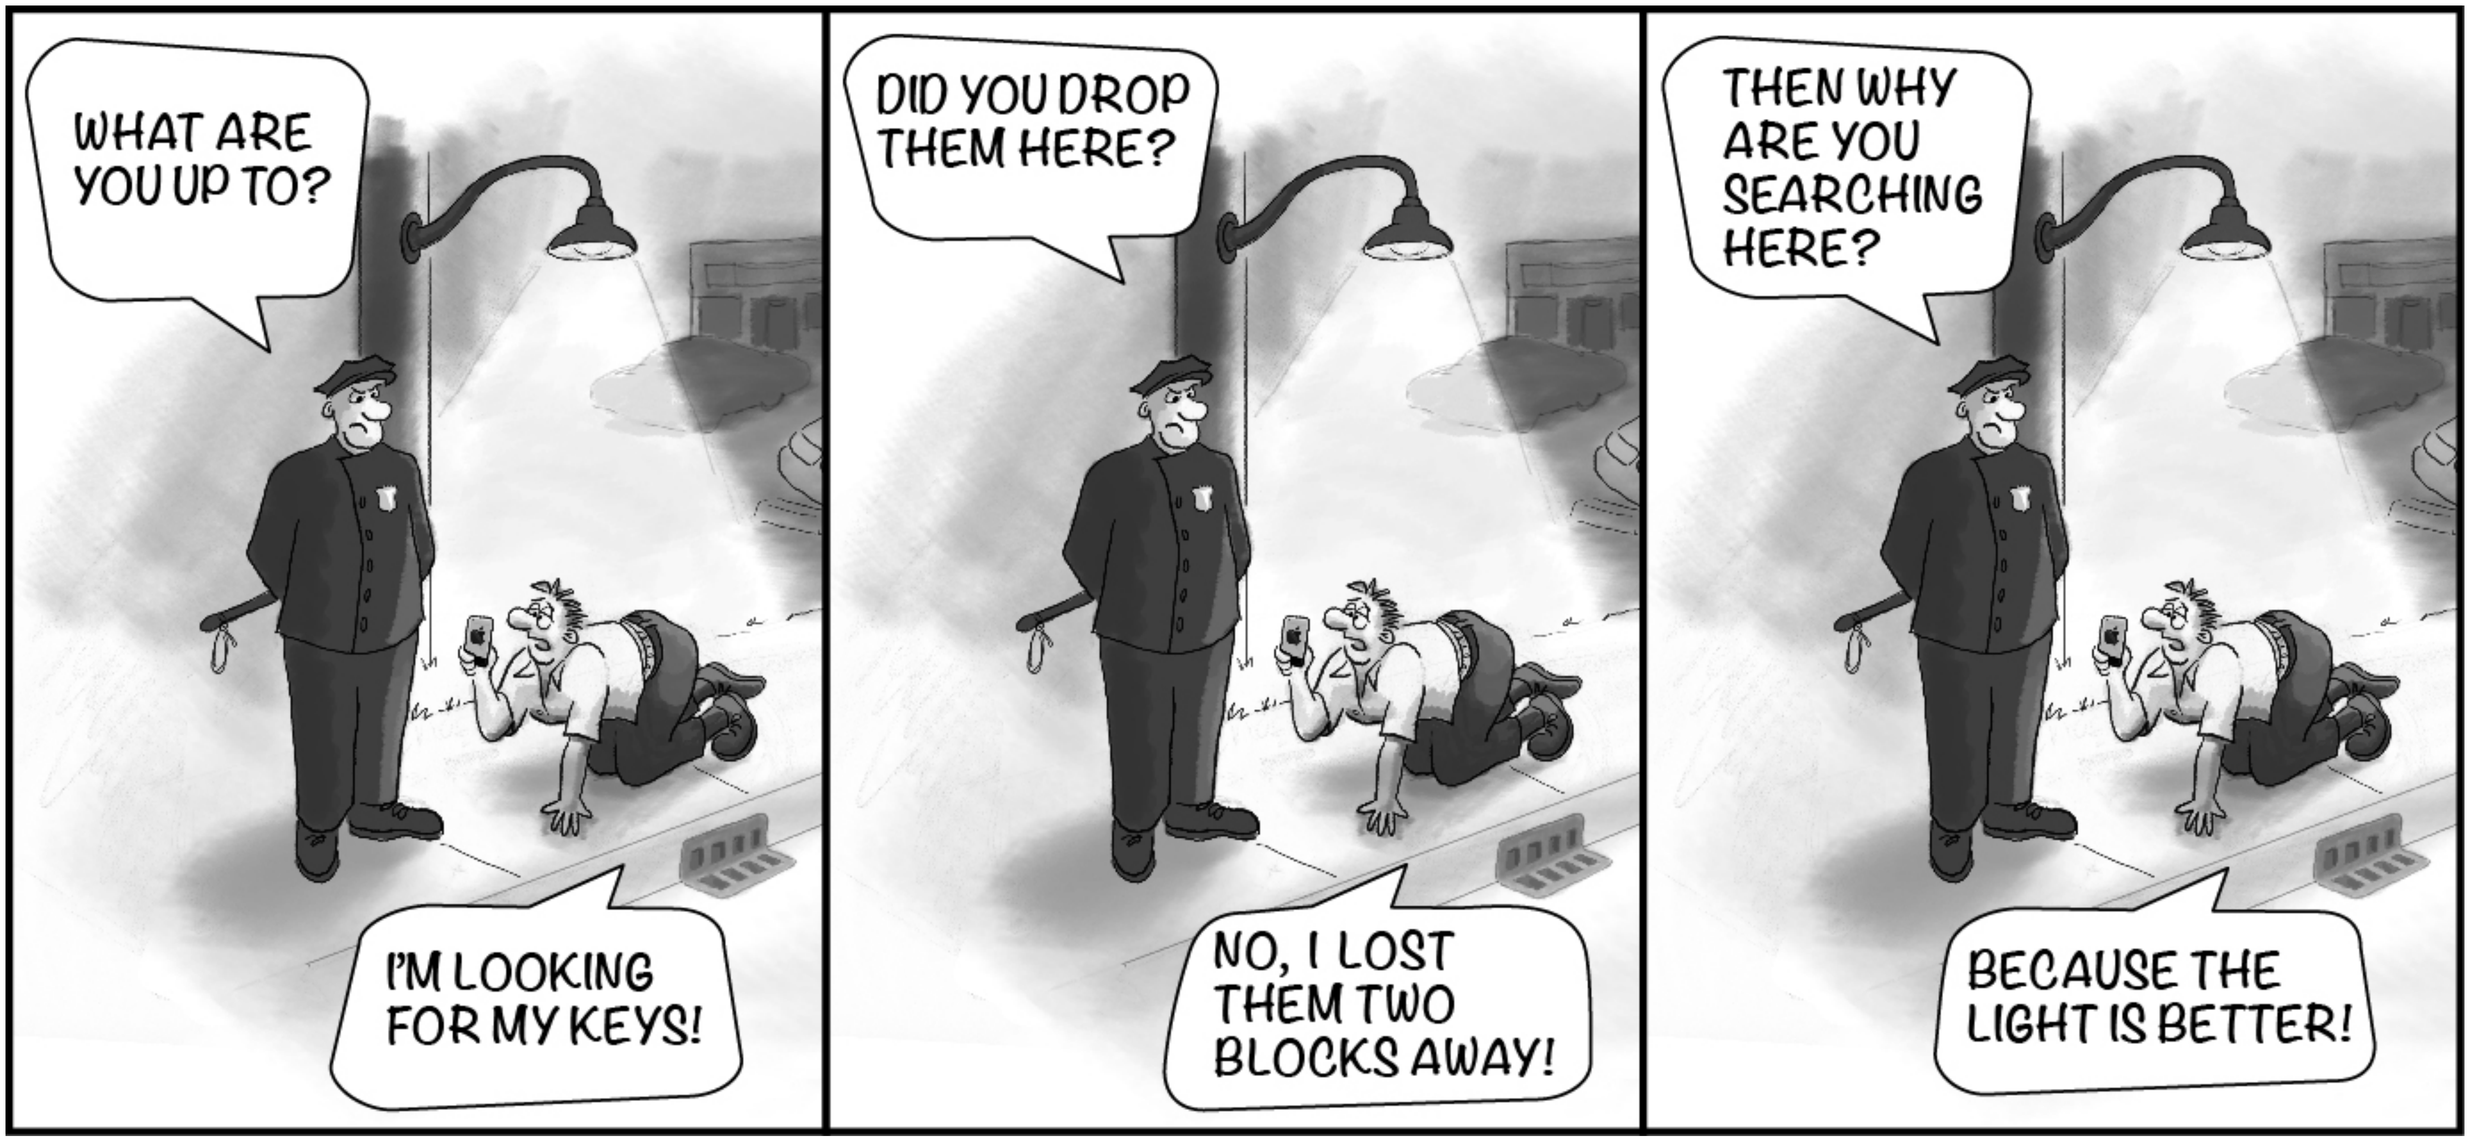

Introduction: Exploration vs. Exploitation

- How to balance the tradeoff between *exploring* and *exploiting*?

  One way...

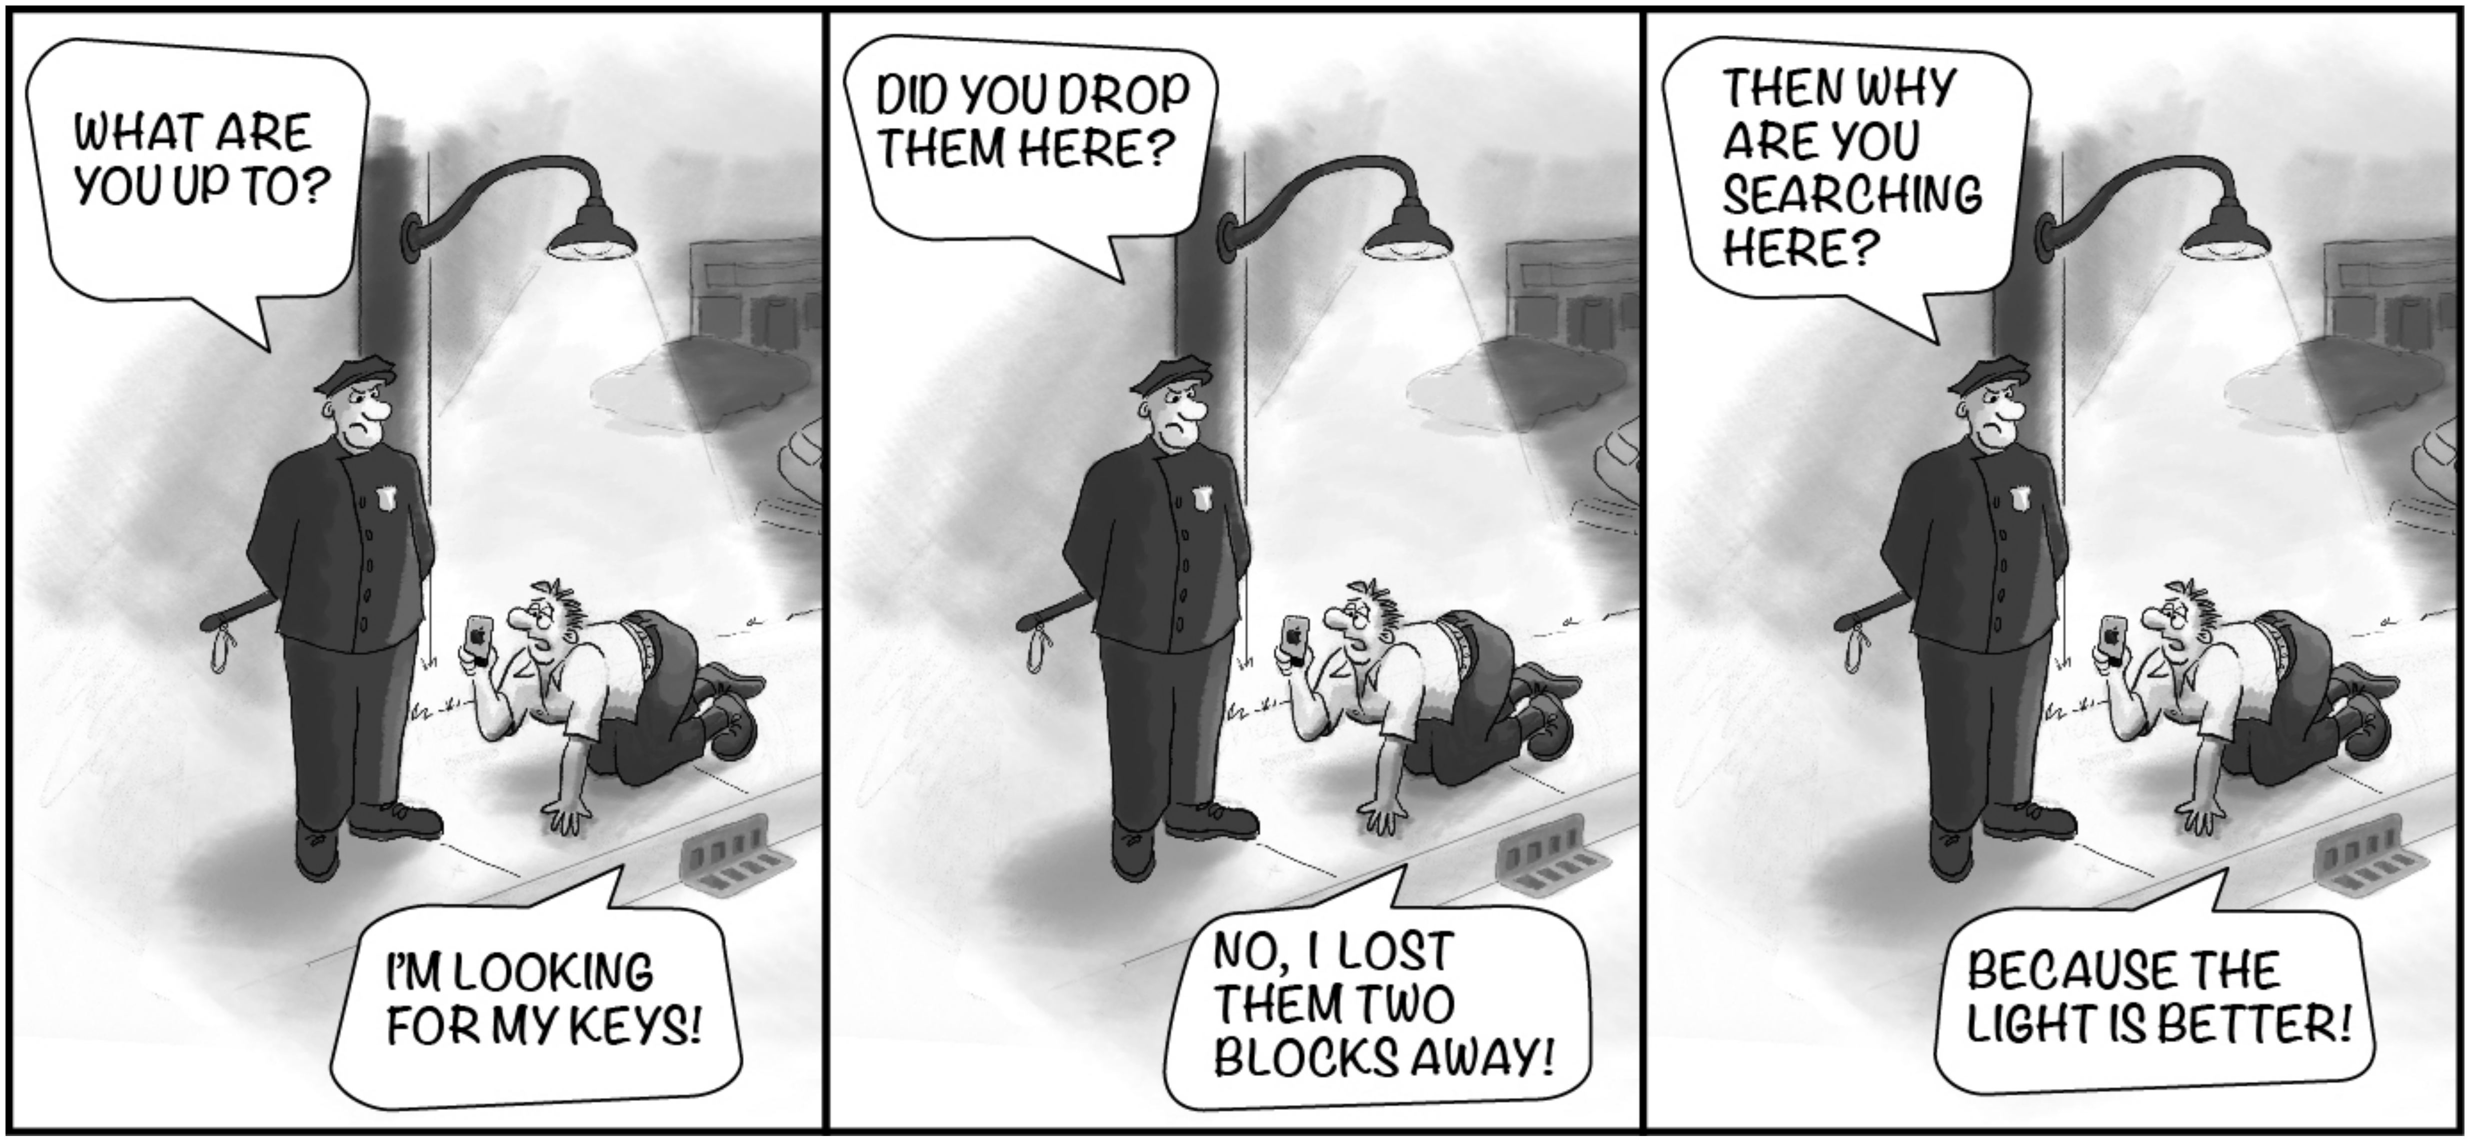

- Not necessarily stupid.

Introduction: Exploration vs. Exploitation

- How to balance the tradeoff between *exploring* and *exploiting*?

  One way...

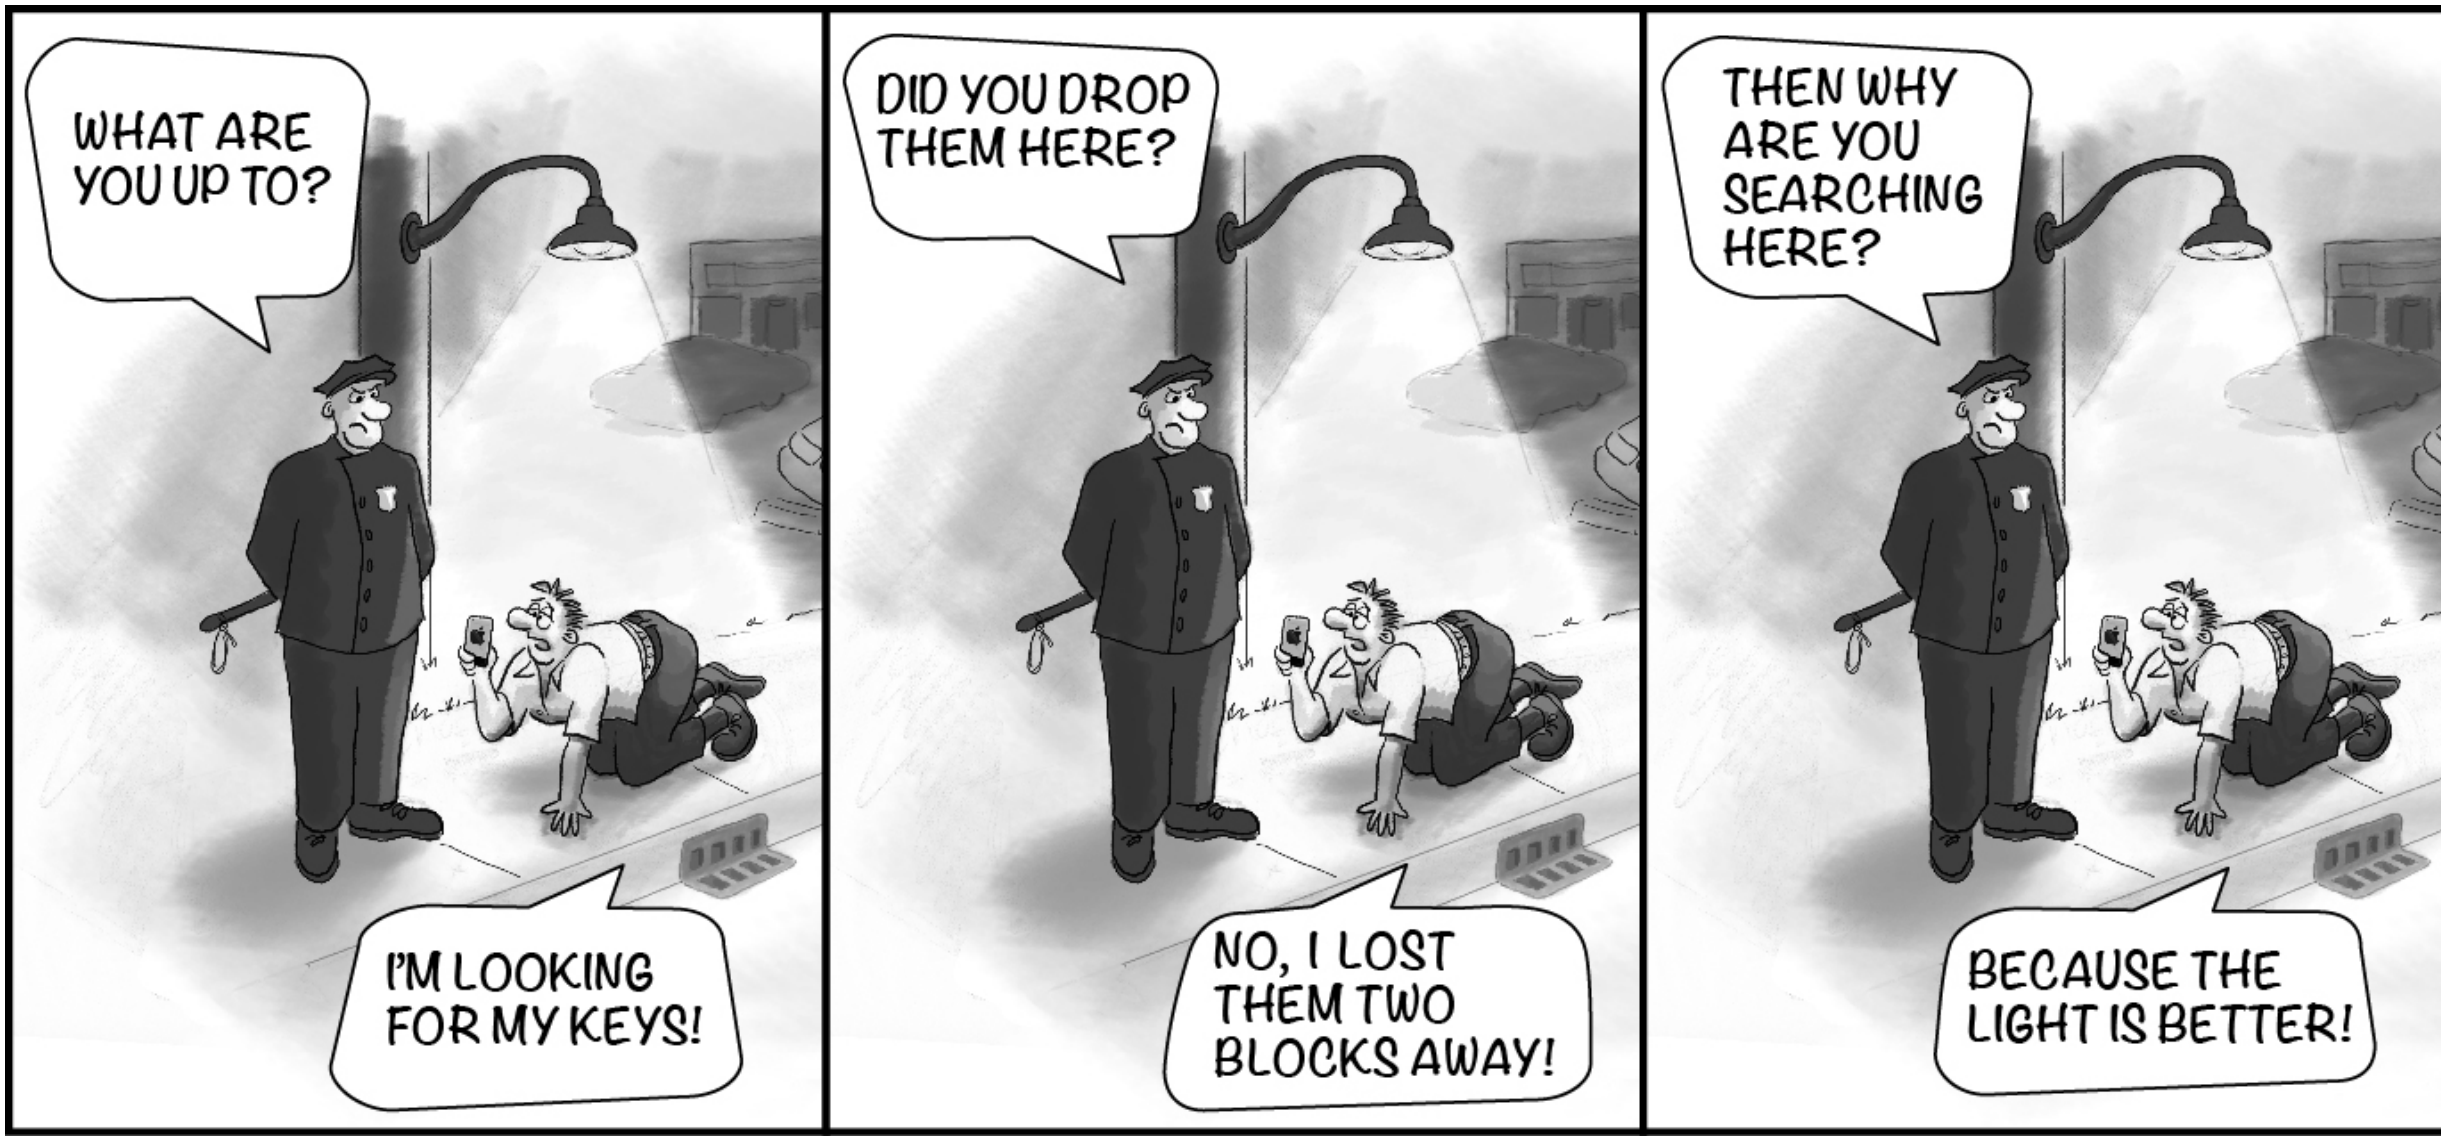

- Not necessarily stupid.
- Different targets call for different measurements.

Introduction: Exploration vs. Exploitation

Introduction: Exploration vs. Exploitation

Stochastic measurements and deep-field imaging call for different approaches.

Introduction: Exploration vs. Exploitation

Stochastic measurements and deep-field imaging call for different approaches.

- Deep-field imaging:

Exploration mostly wasted.

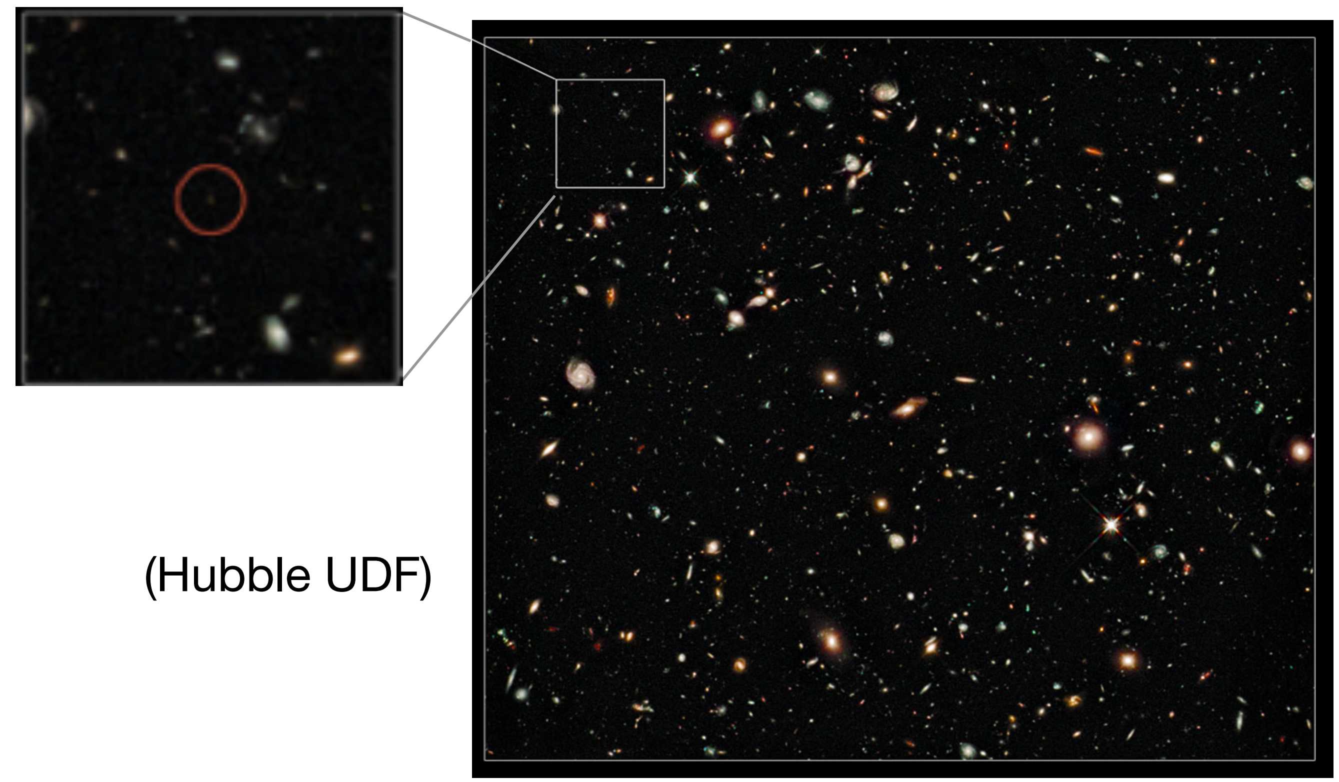

(Hubble UDF)

Introduction: Exploration vs. Exploitation

Stochastic measurements and deep-field imaging call for different approaches.

- Deep-field imaging:

Exploration mostly wasted.

Goal of adaptive strategy:

--> quickly converge and *exploit*.

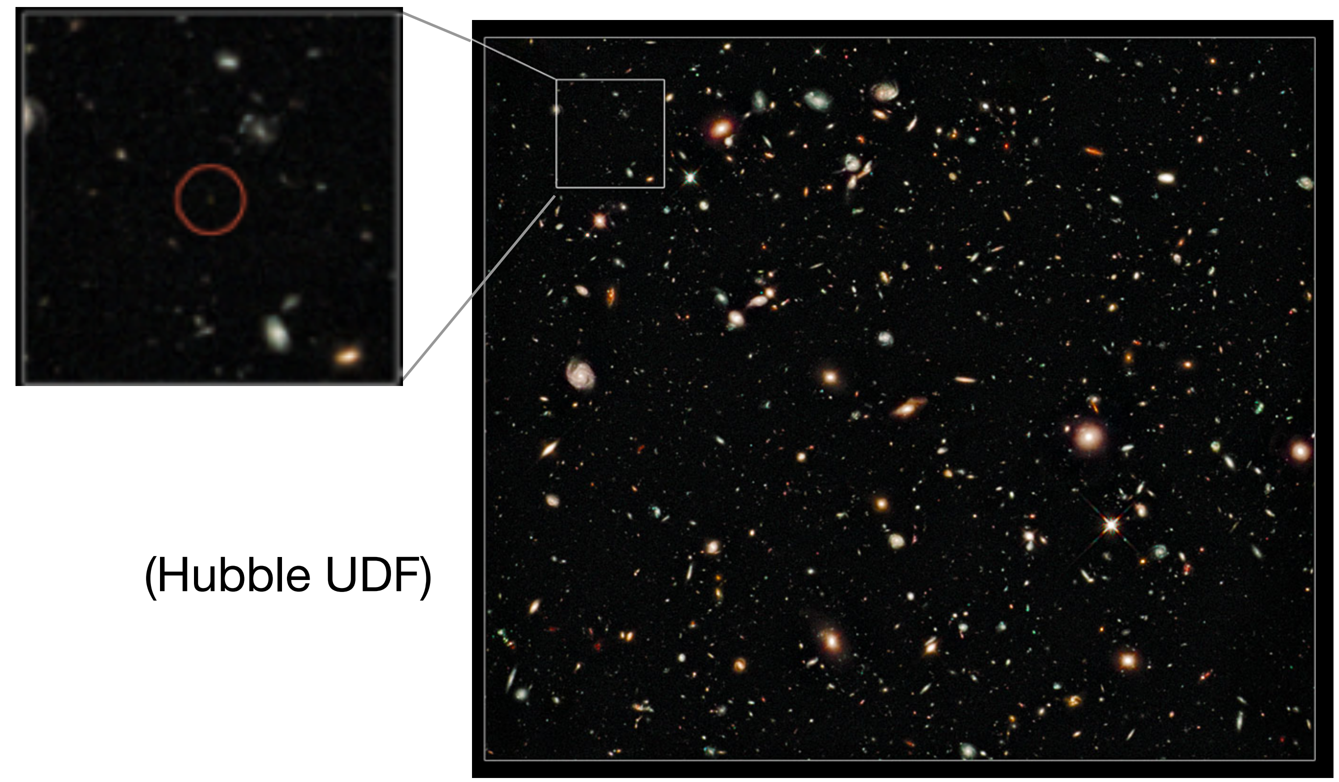

(Hubble UDF)

Introduction: Exploration vs. Exploitation

Stochastic measurements and deep-field imaging call for different approaches.

- Deep-field imaging:

Exploration mostly wasted.

Goal of adaptive strategy:

--> quickly converge and *exploit*.

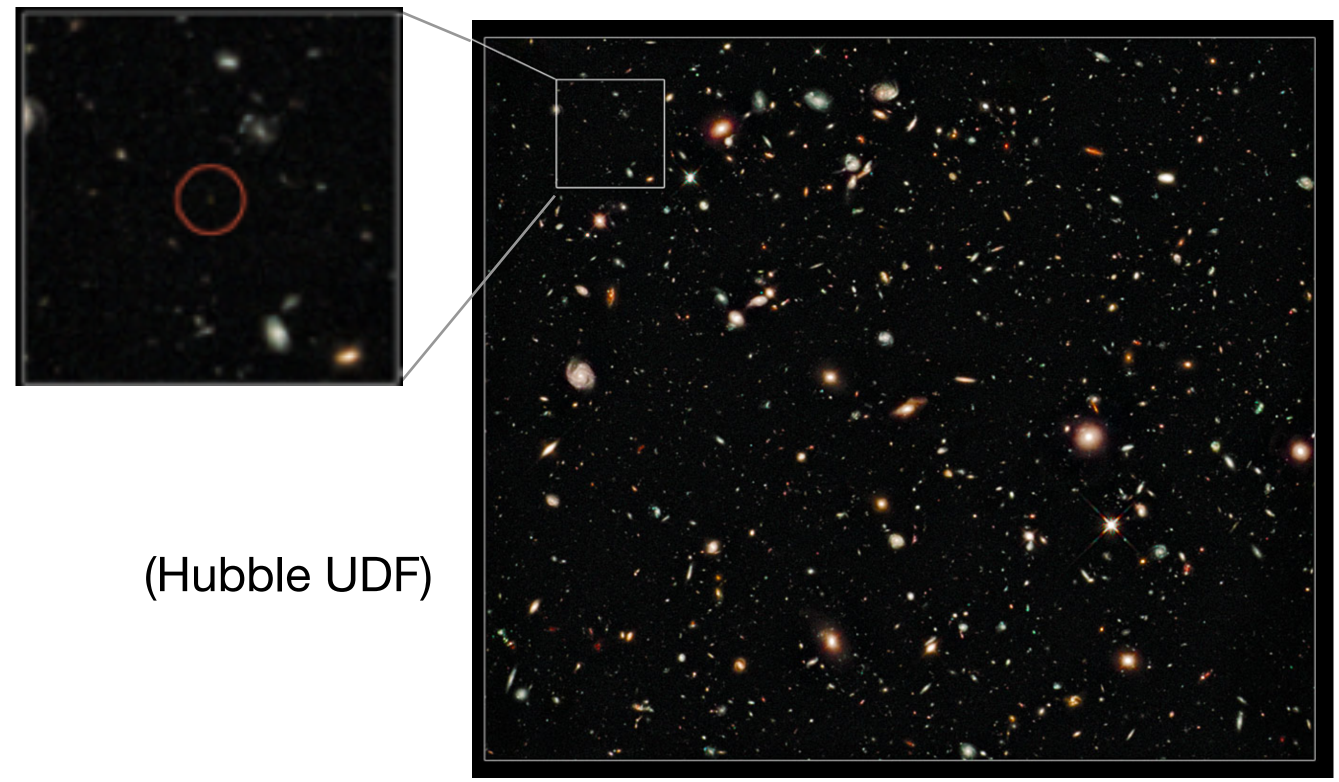

(Hubble UDF)

- Stochastic fluctuations:

Exploration mitigates cosmic variance.

(QUIET Fields)

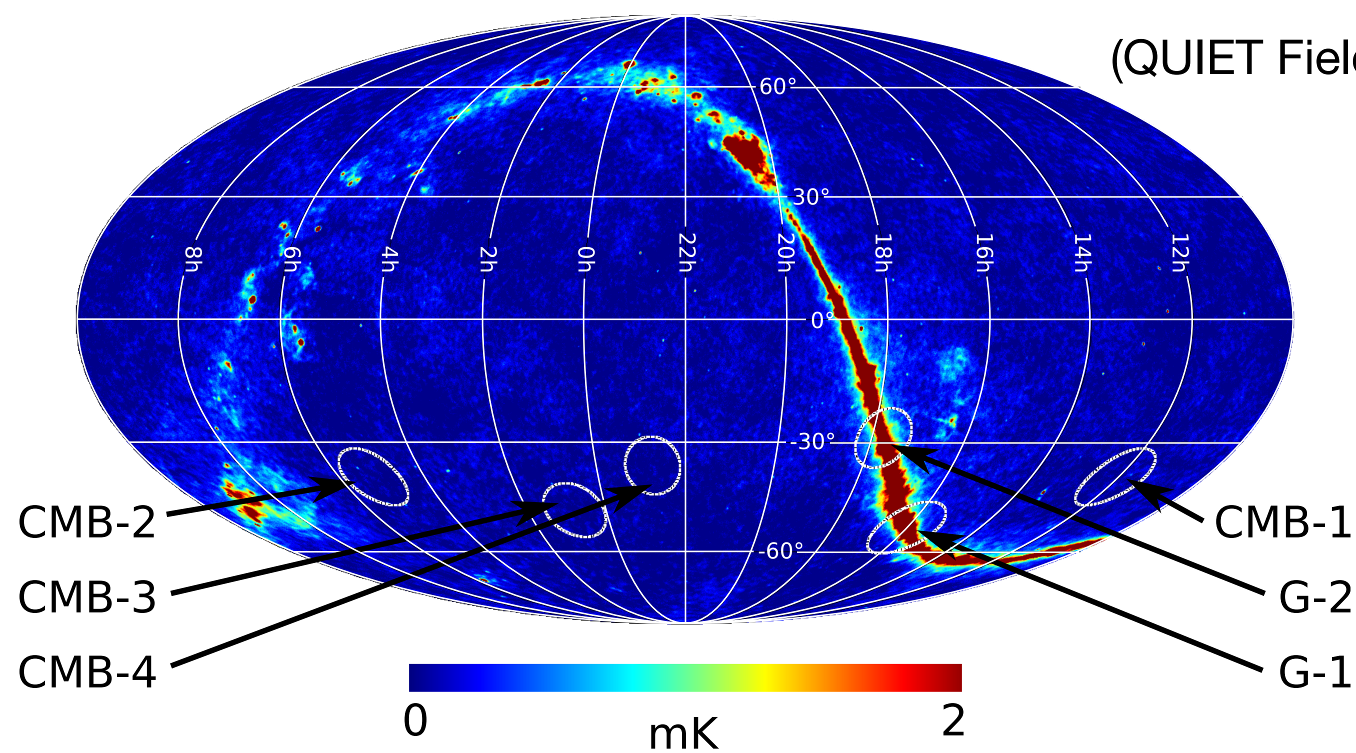

Introduction: Exploration vs. Exploitation

Stochastic measurements and deep-field imaging call for different approaches.

- Deep-field imaging:

Exploration mostly wasted.

Goal of adaptive strategy:

--> quickly converge and *exploit*.

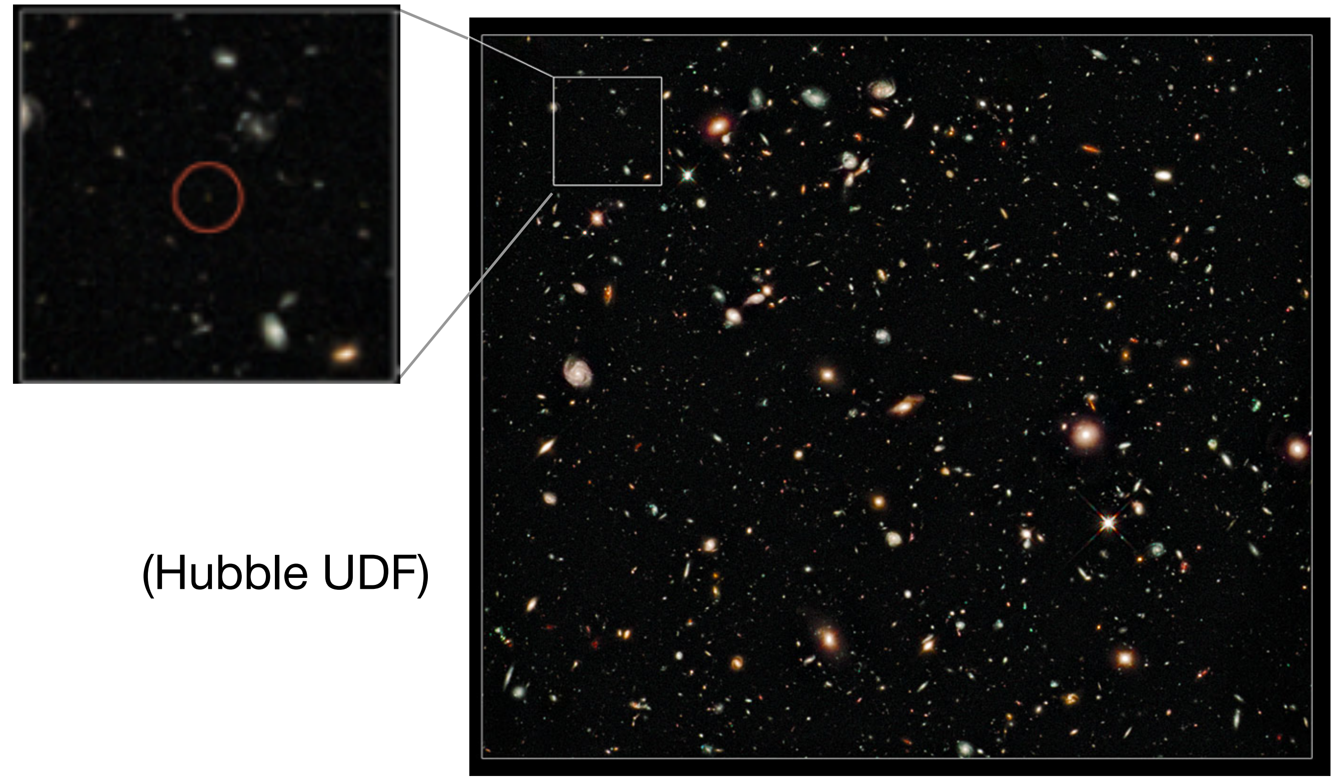

(Hubble UDF)

- Stochastic fluctuations:

Exploration mitigates cosmic variance.

Goal of adaptive strategy:

--> find ideal patches to *exploit*.

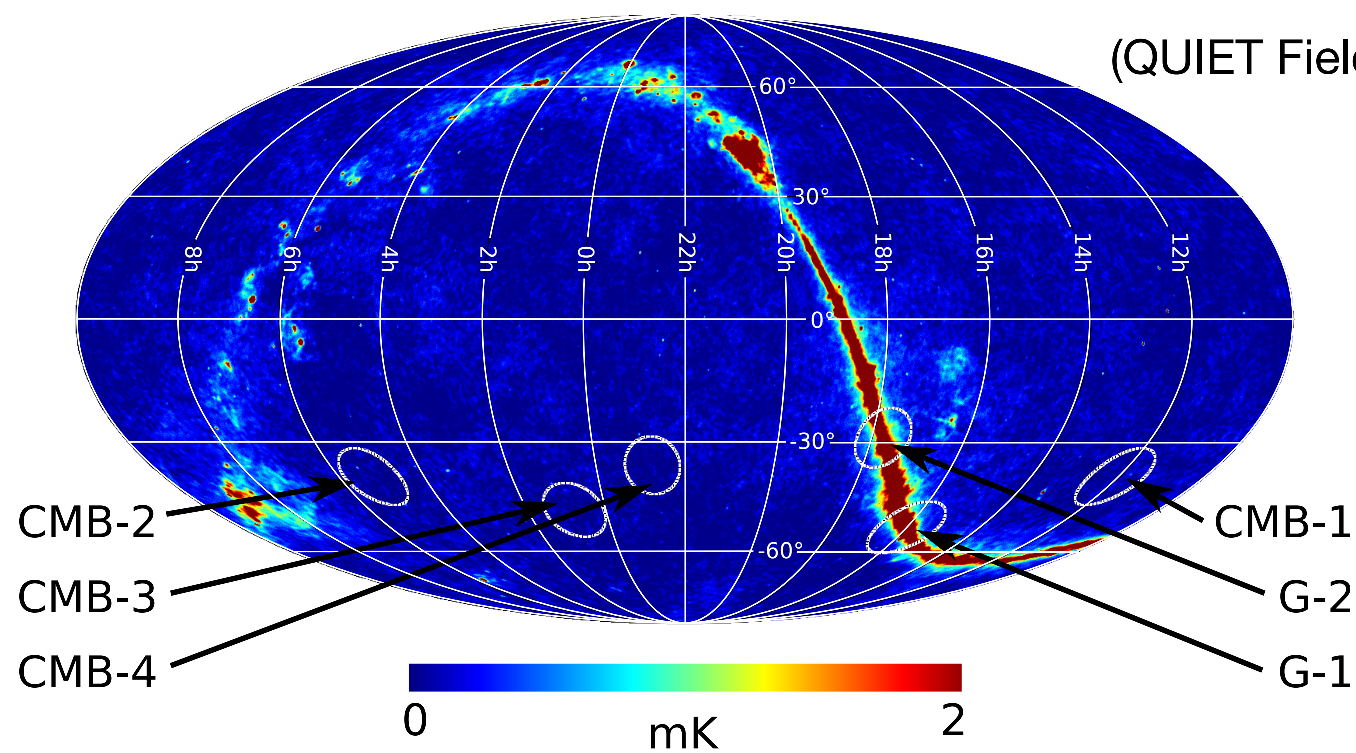

(QUIET Fields)

Exploration vs. Exploitation: B-mode Surveys

Exploration vs. Exploitation: B-mode Surveys

- A preferred method of detection is prolonged integration over a small sky patch.

Exploration vs. Exploitation: B-mode Surveys

- A preferred method of detection is prolonged integration over a small sky patch.
- Tradeoff is between *finding lower-foreground patches* and *integrating over them*.

Exploration vs. Exploitation: B-mode Surveys

- A preferred method of detection is prolonged integration over a small sky patch.
- Tradeoff is between *finding lower-foreground patches* and *integrating over them*.

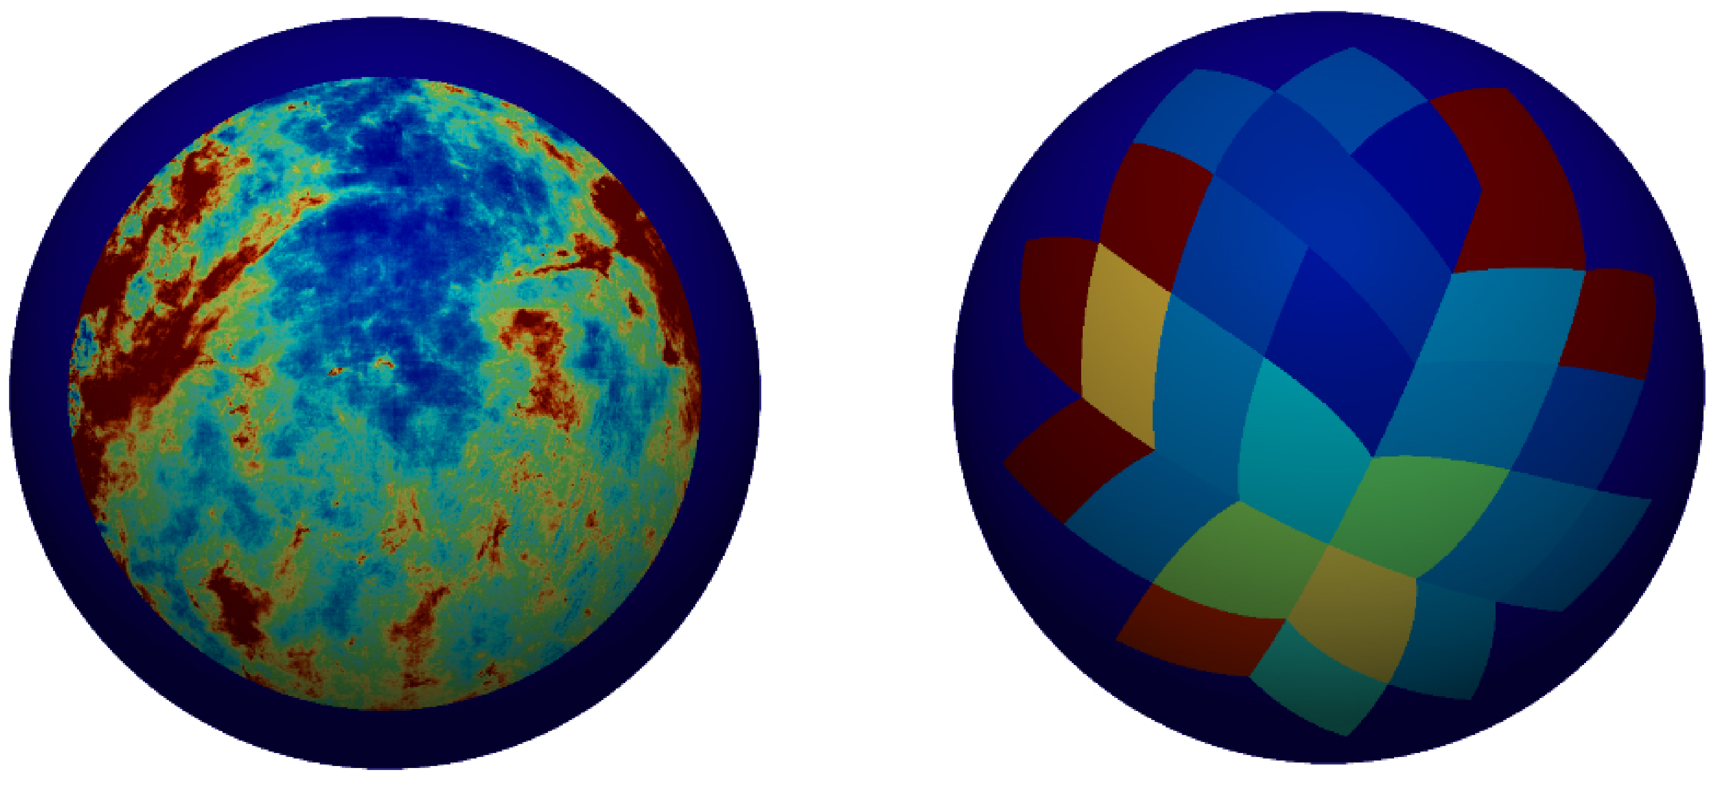

Templates for polarized emission from dust (PED) in the Galaxy at 150GHz

(Clark et al. arXiv:1211.6404)

Exploration vs. Exploitation: B-mode Surveys

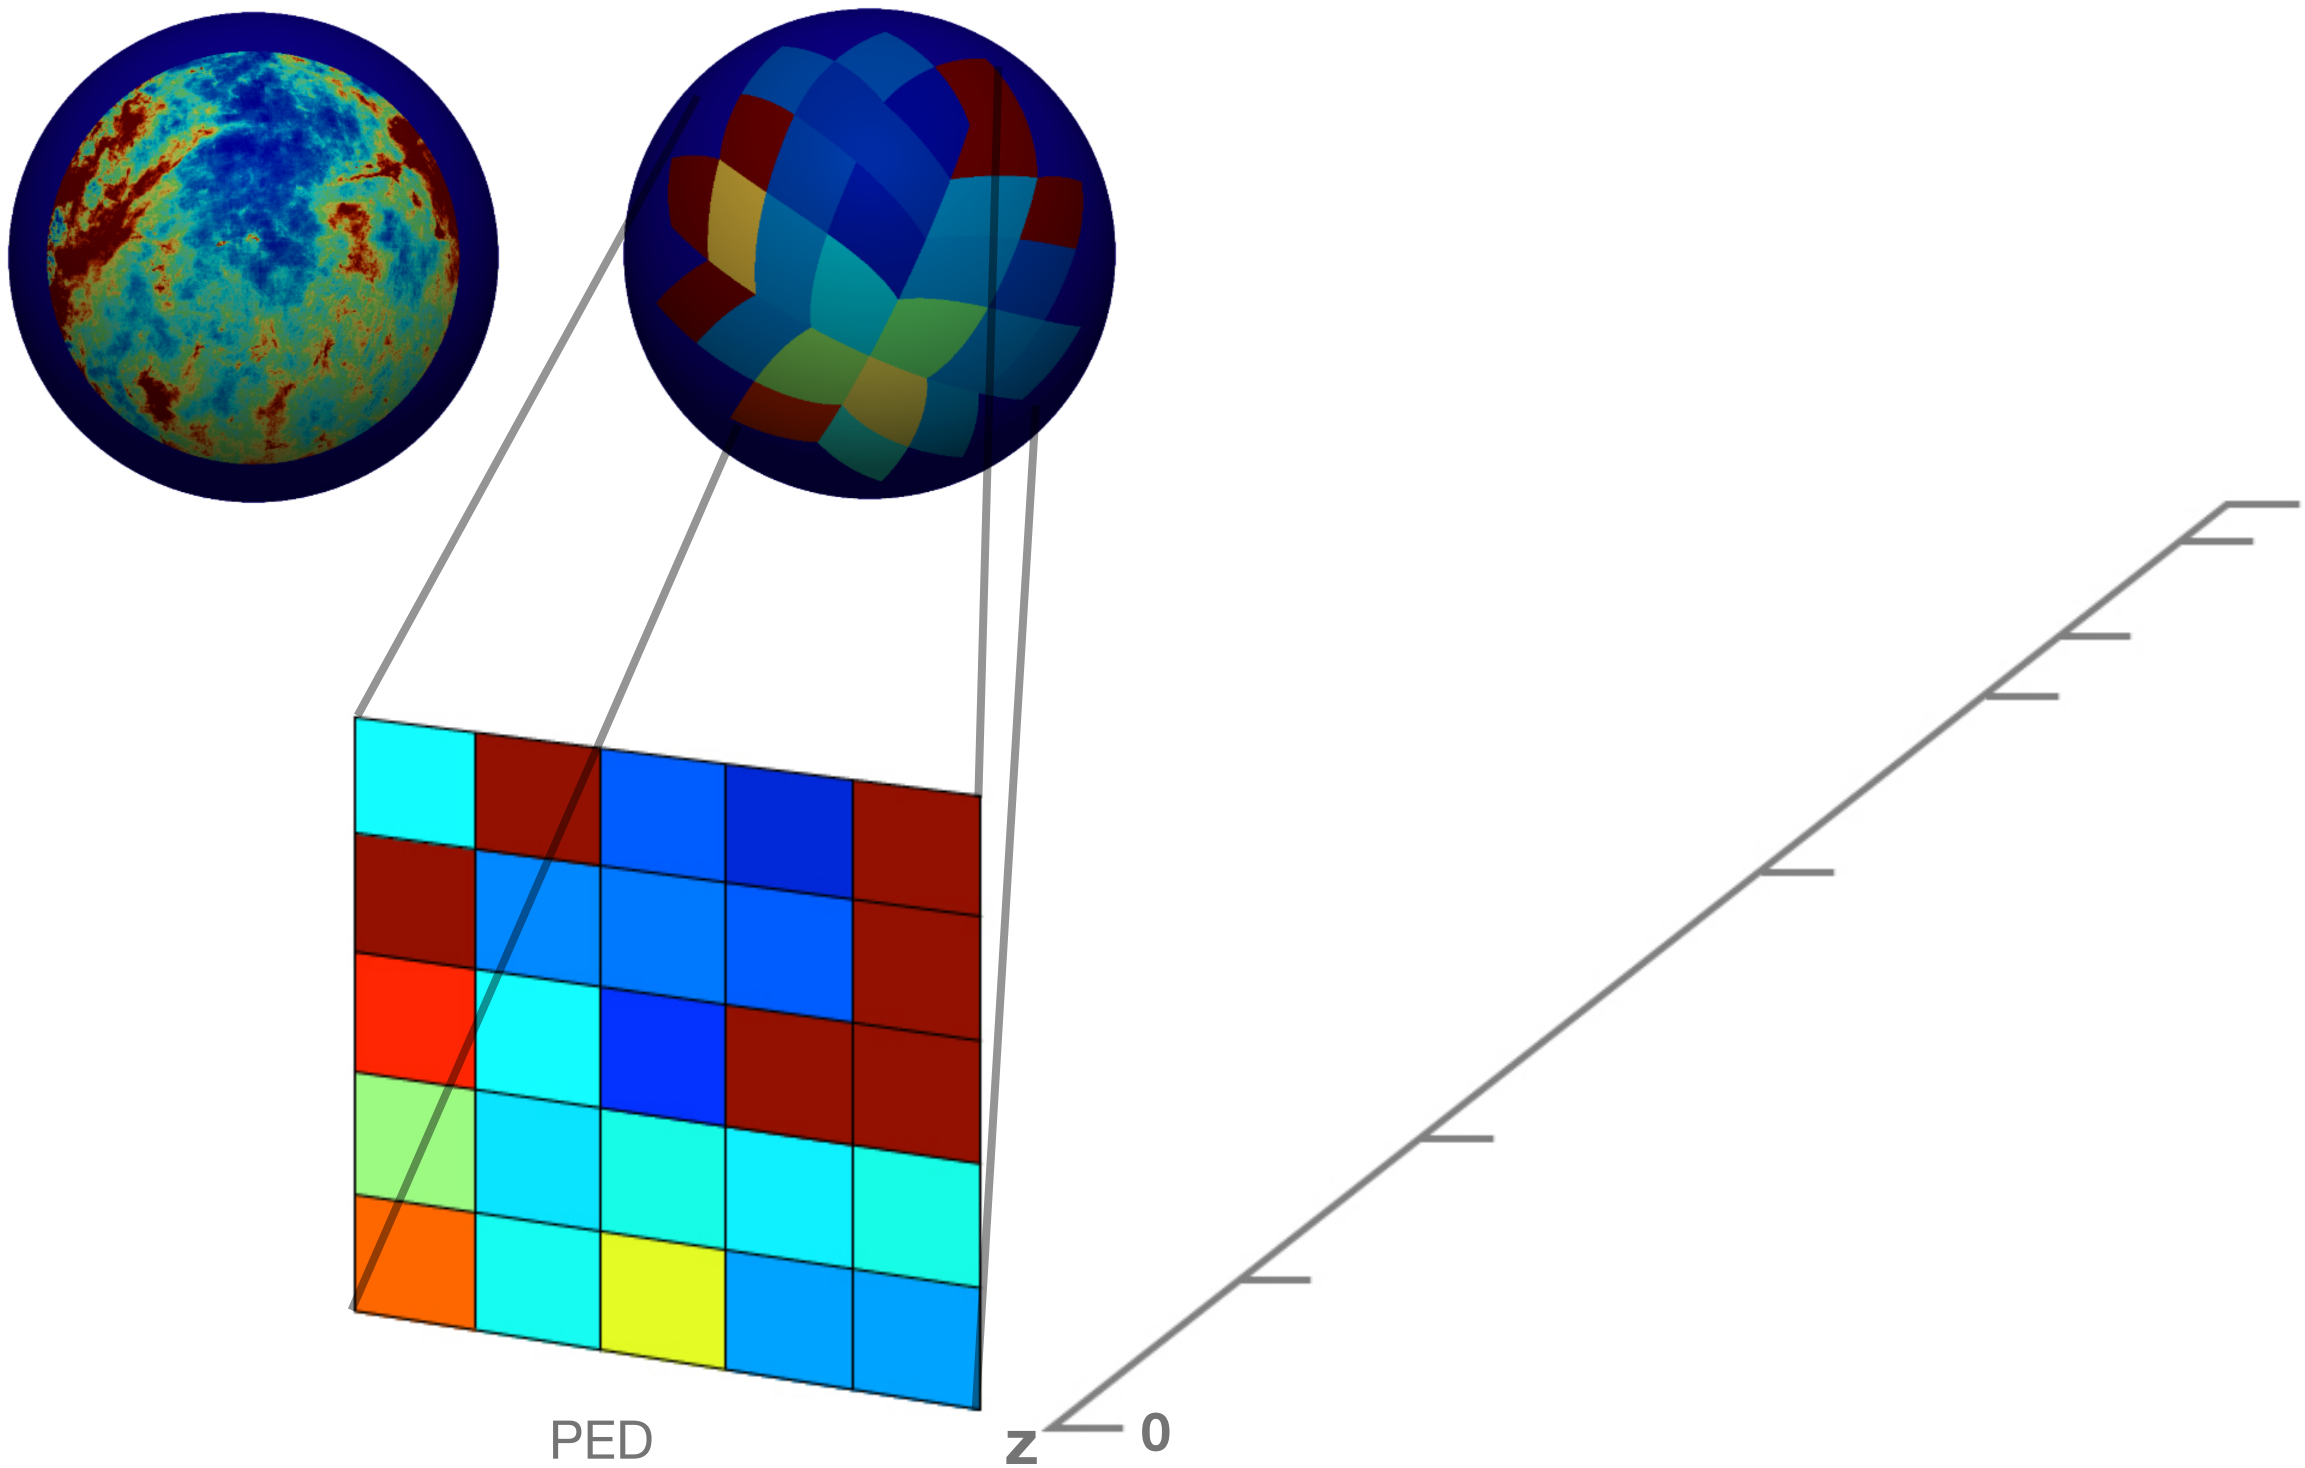

Exploration vs. Exploitation: B-mode Surveys

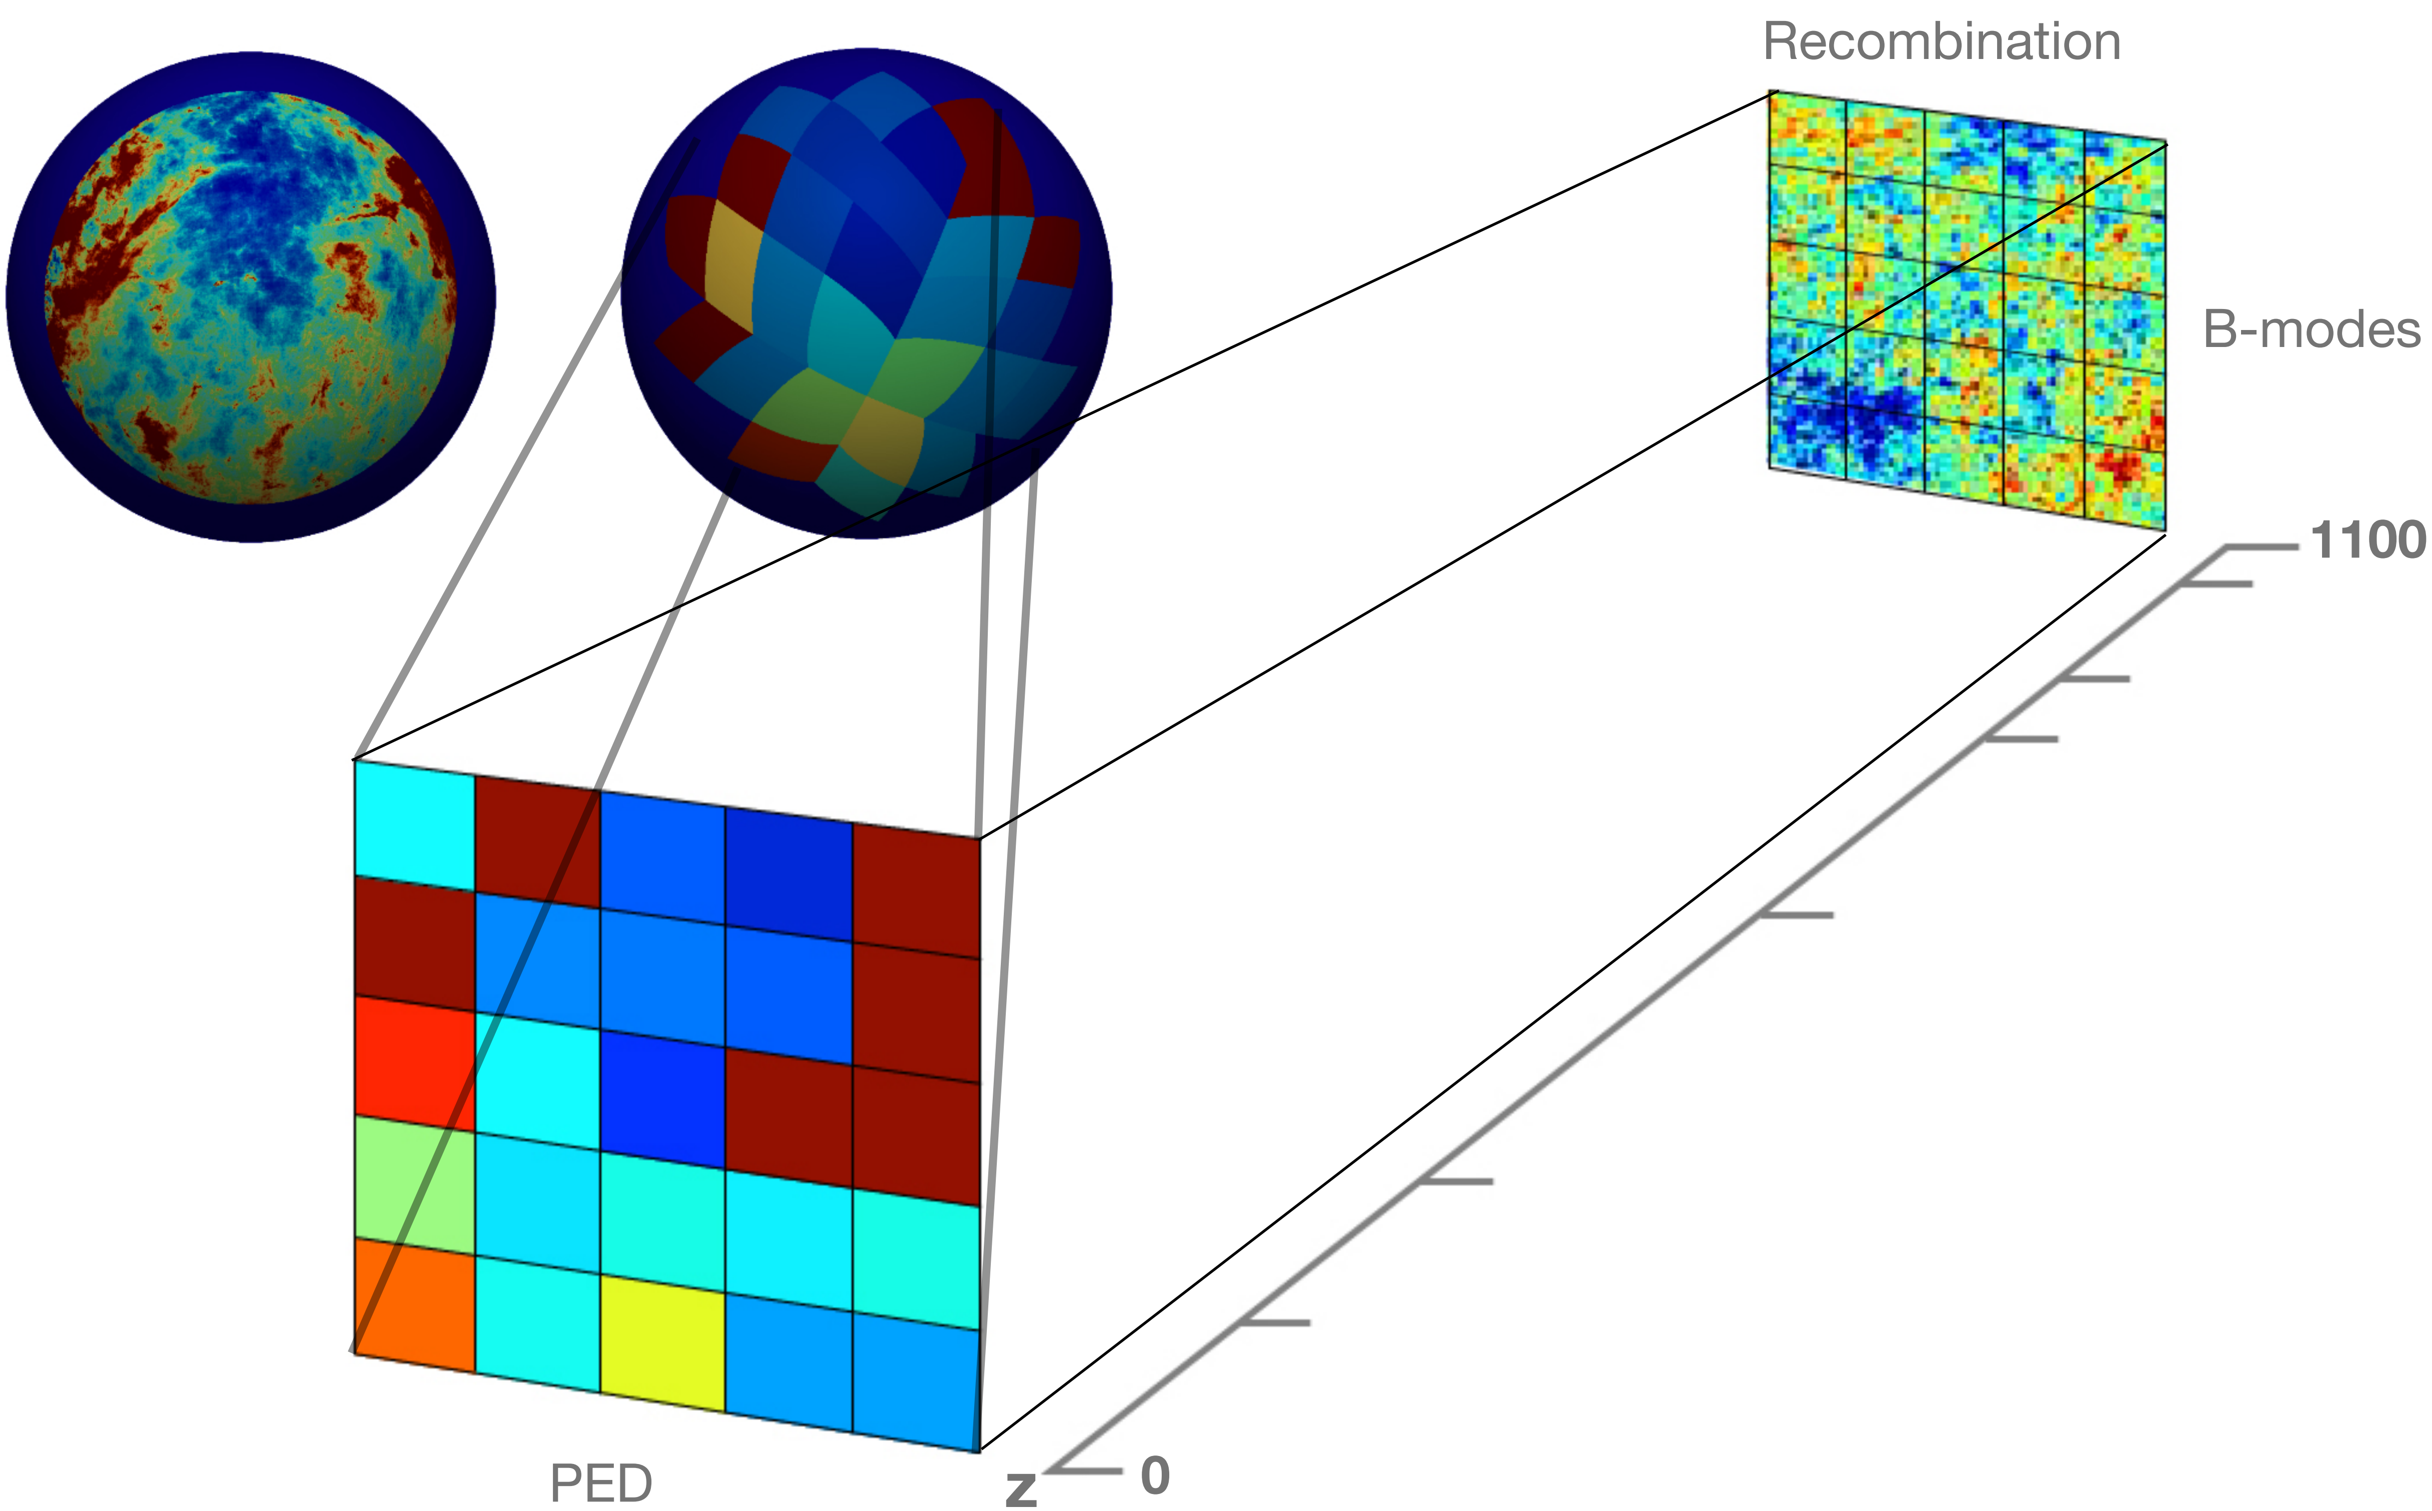

Exploration vs. Exploitation: B-mode Surveys

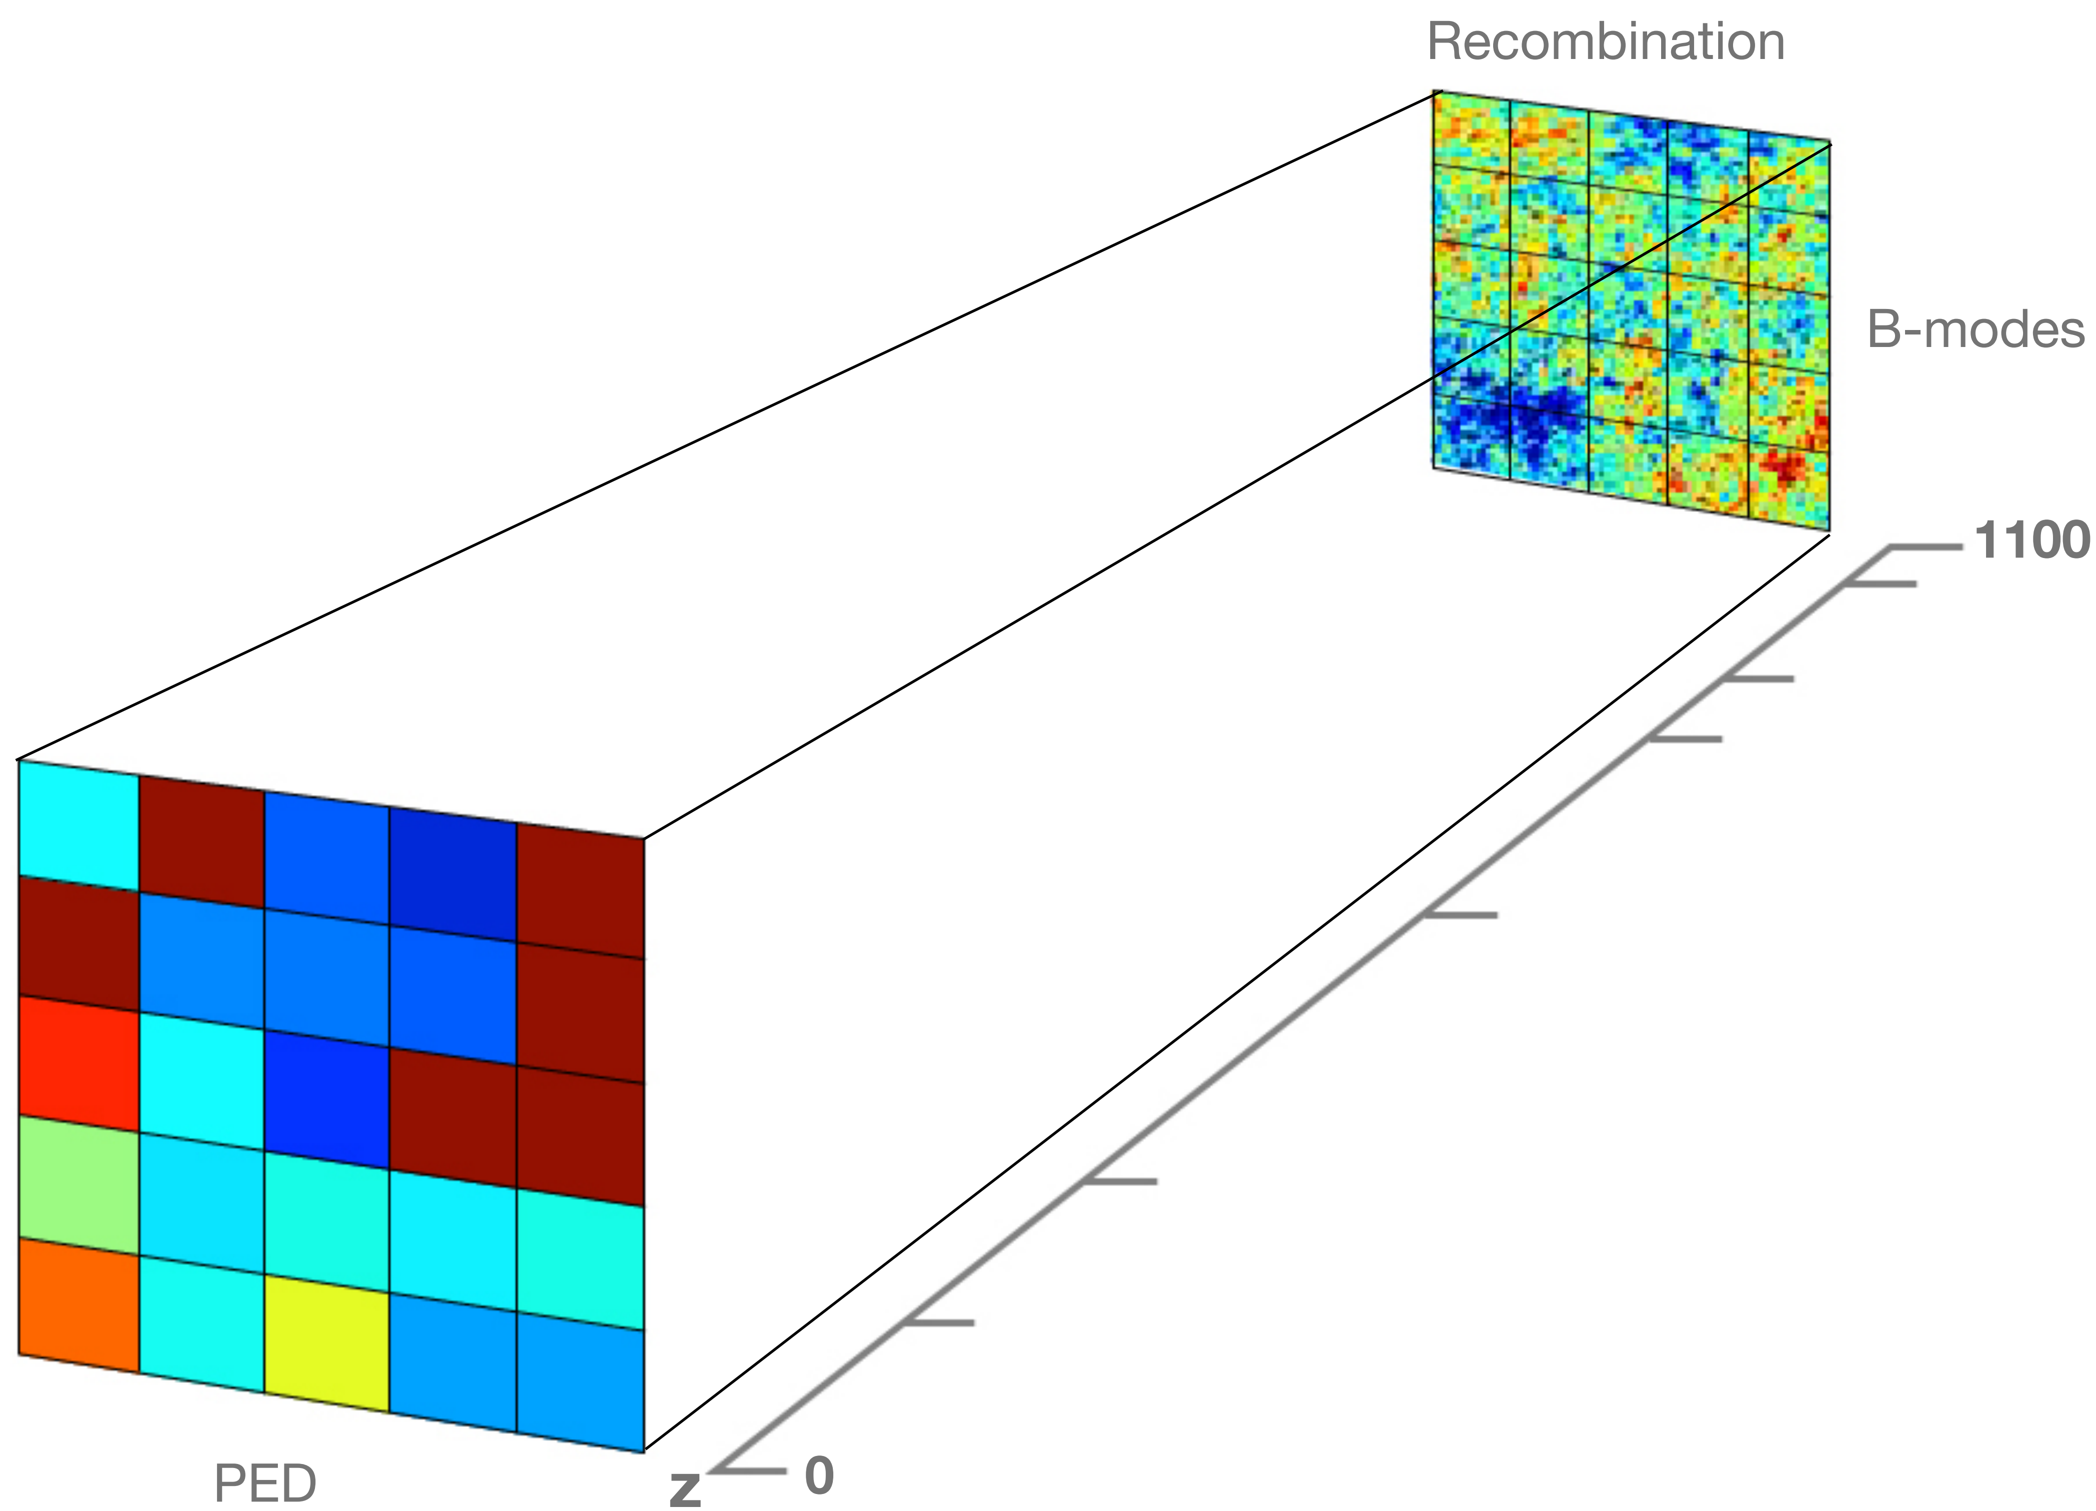

Exploration vs. Exploitation: B-mode Surveys

- Worst patches can probably be avoided.

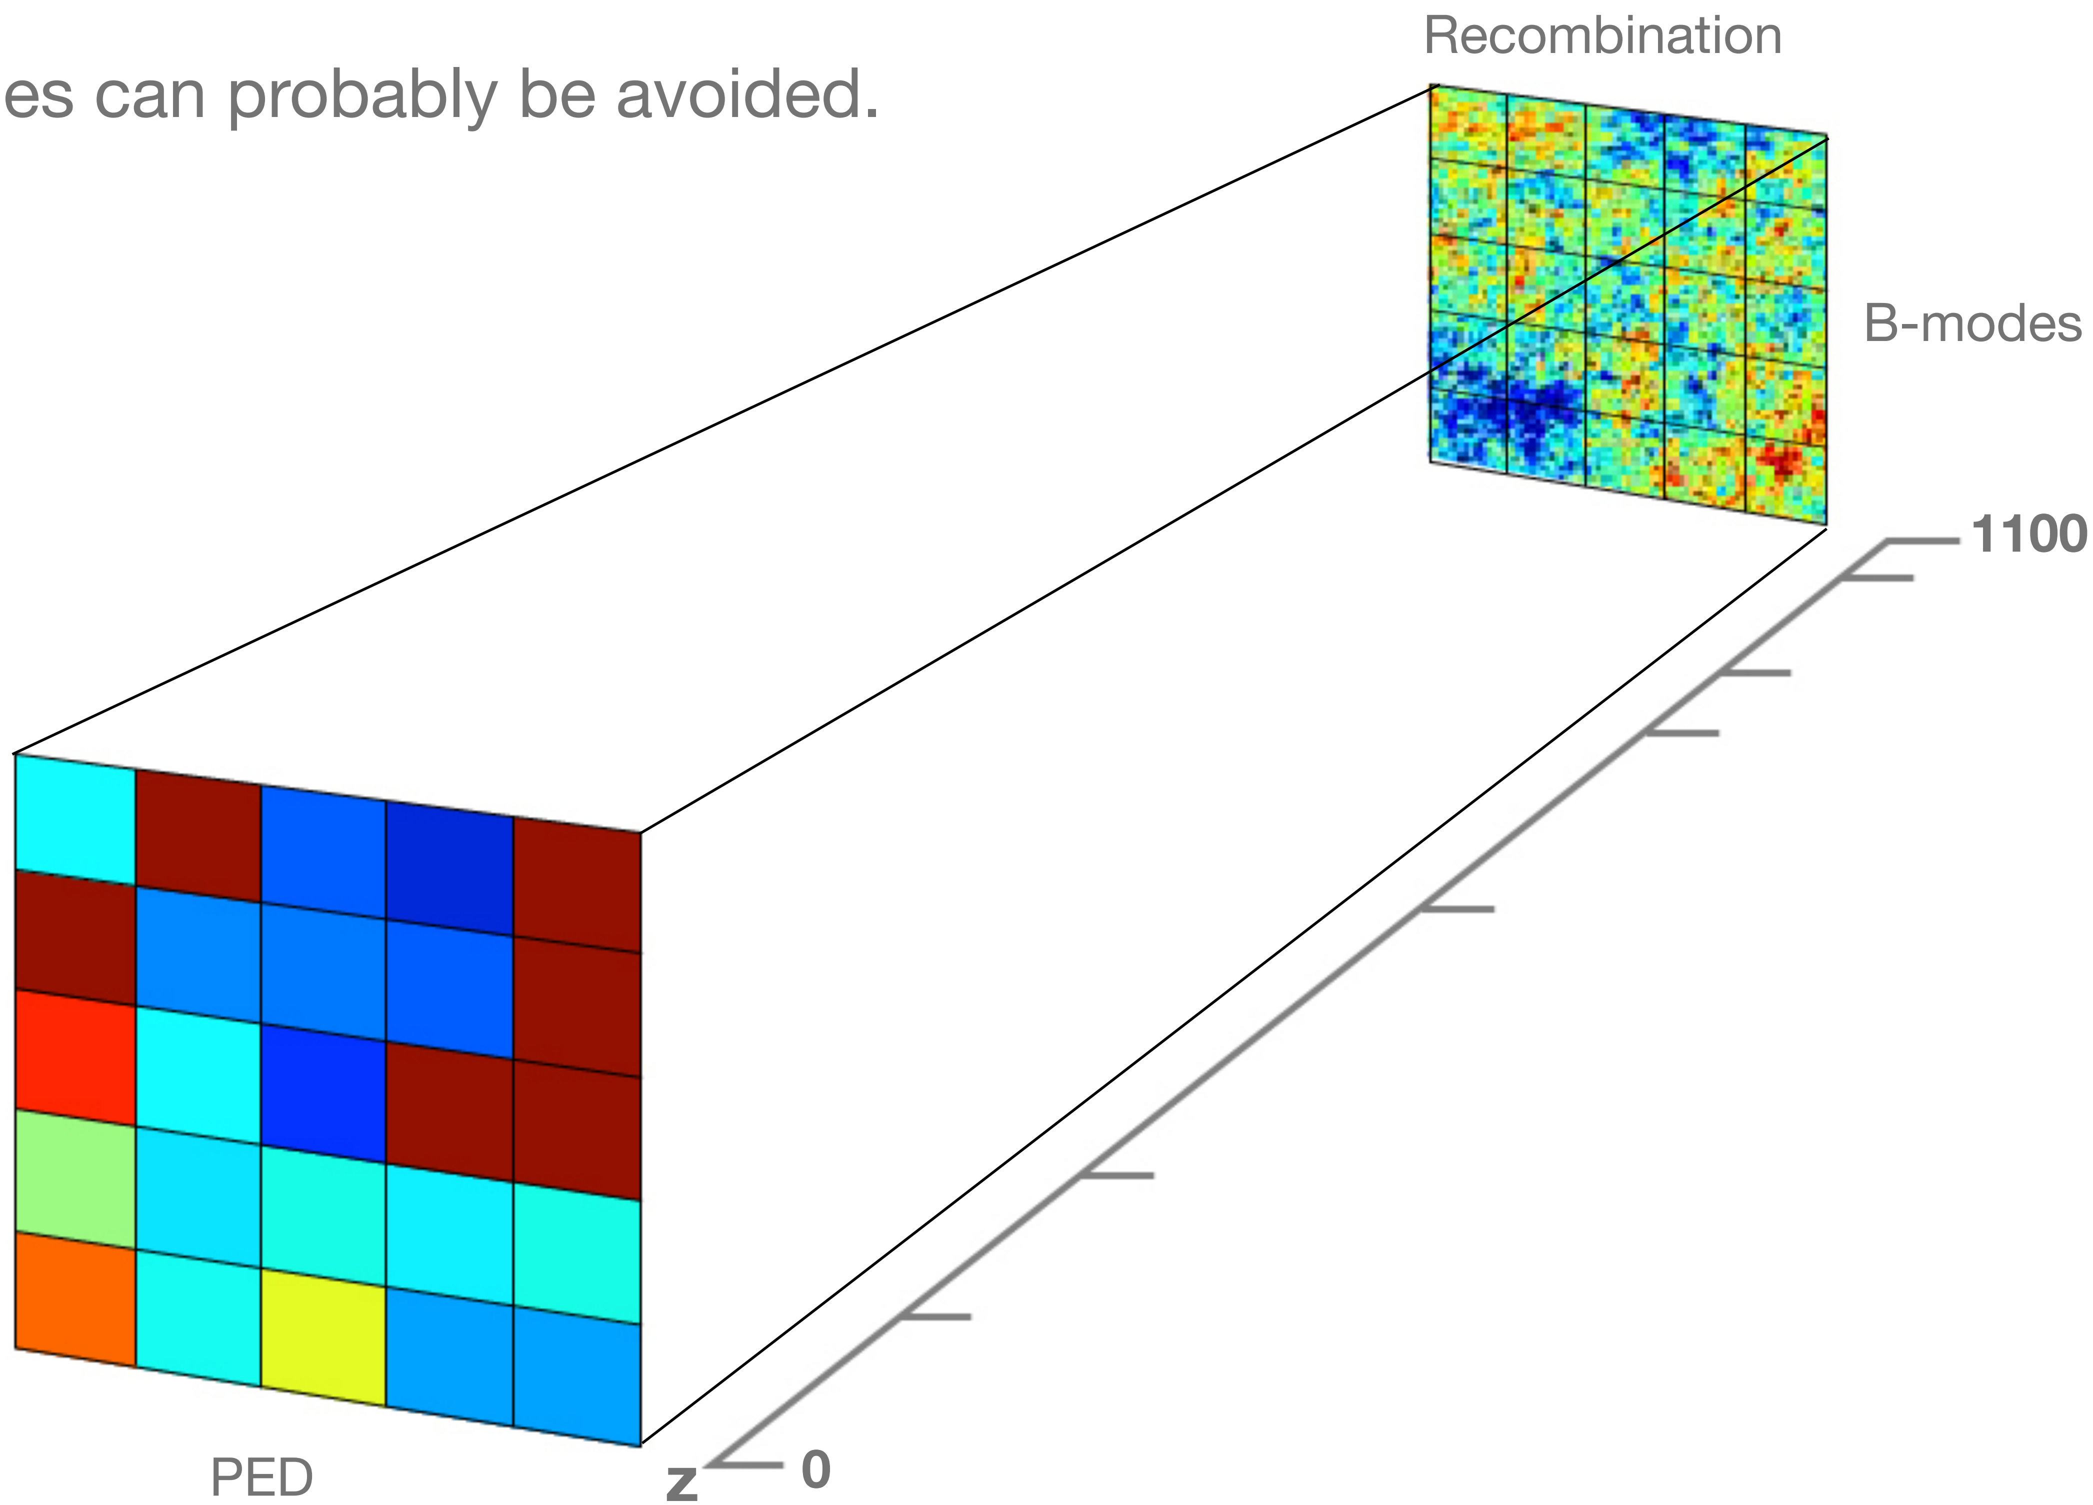

Exploration vs. Exploitation: B-mode Surveys

- Worst patches can probably be avoided.

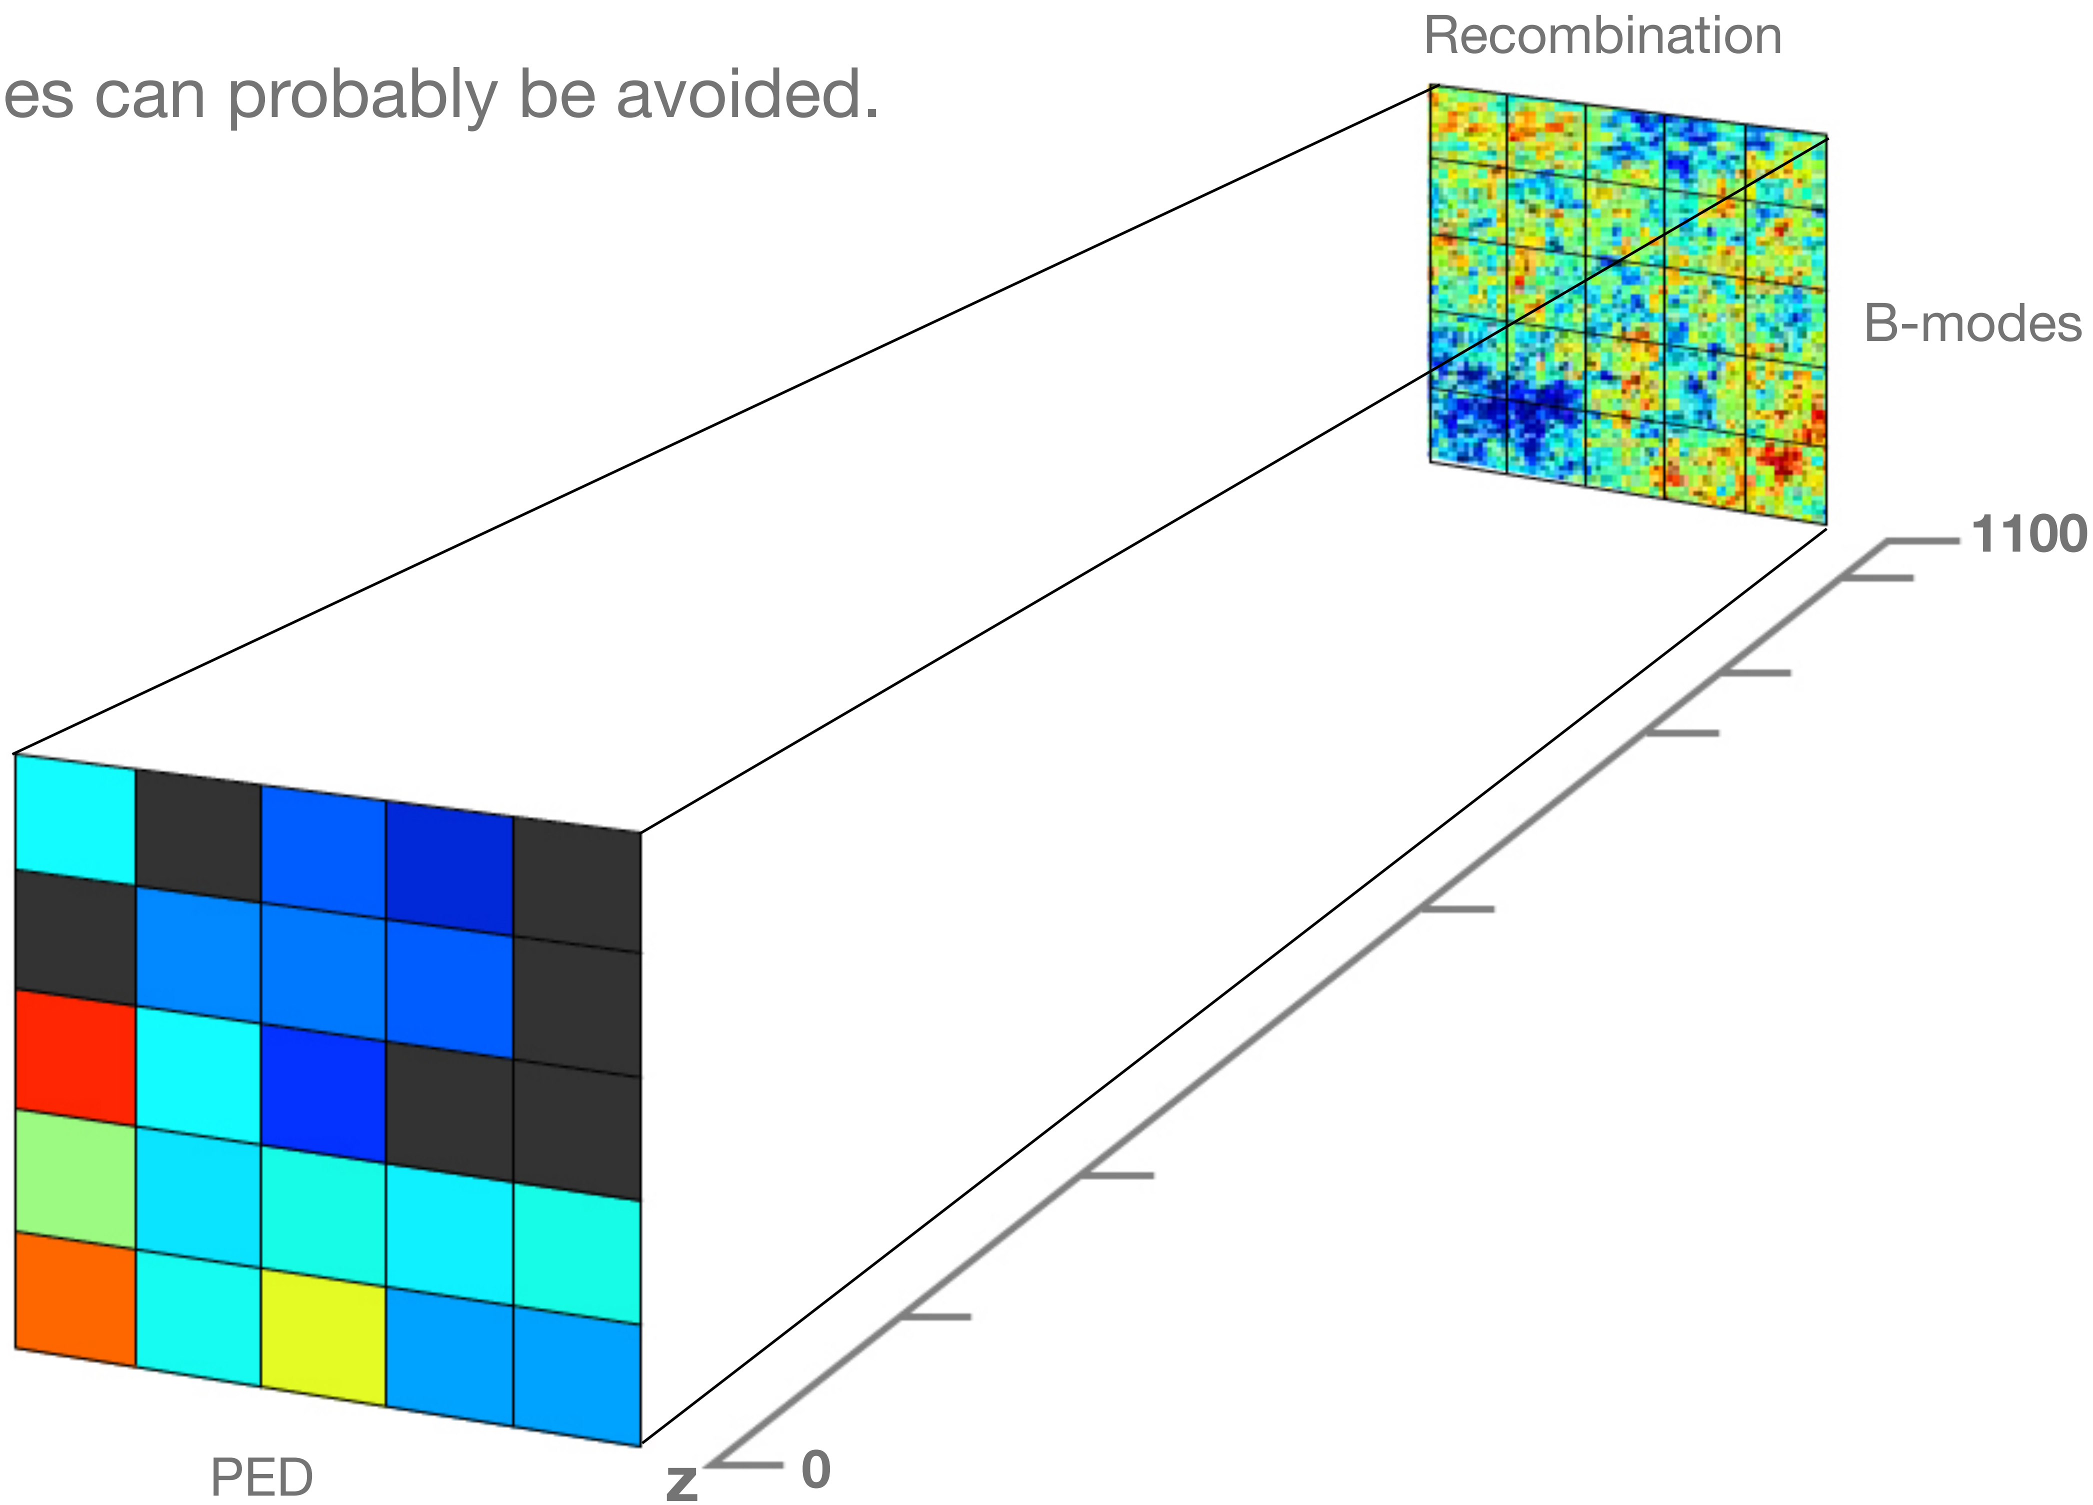

Exploration vs. Exploitation: B-mode Surveys

- Worst patches can probably be avoided.
- But cleanest patches not yet detected.

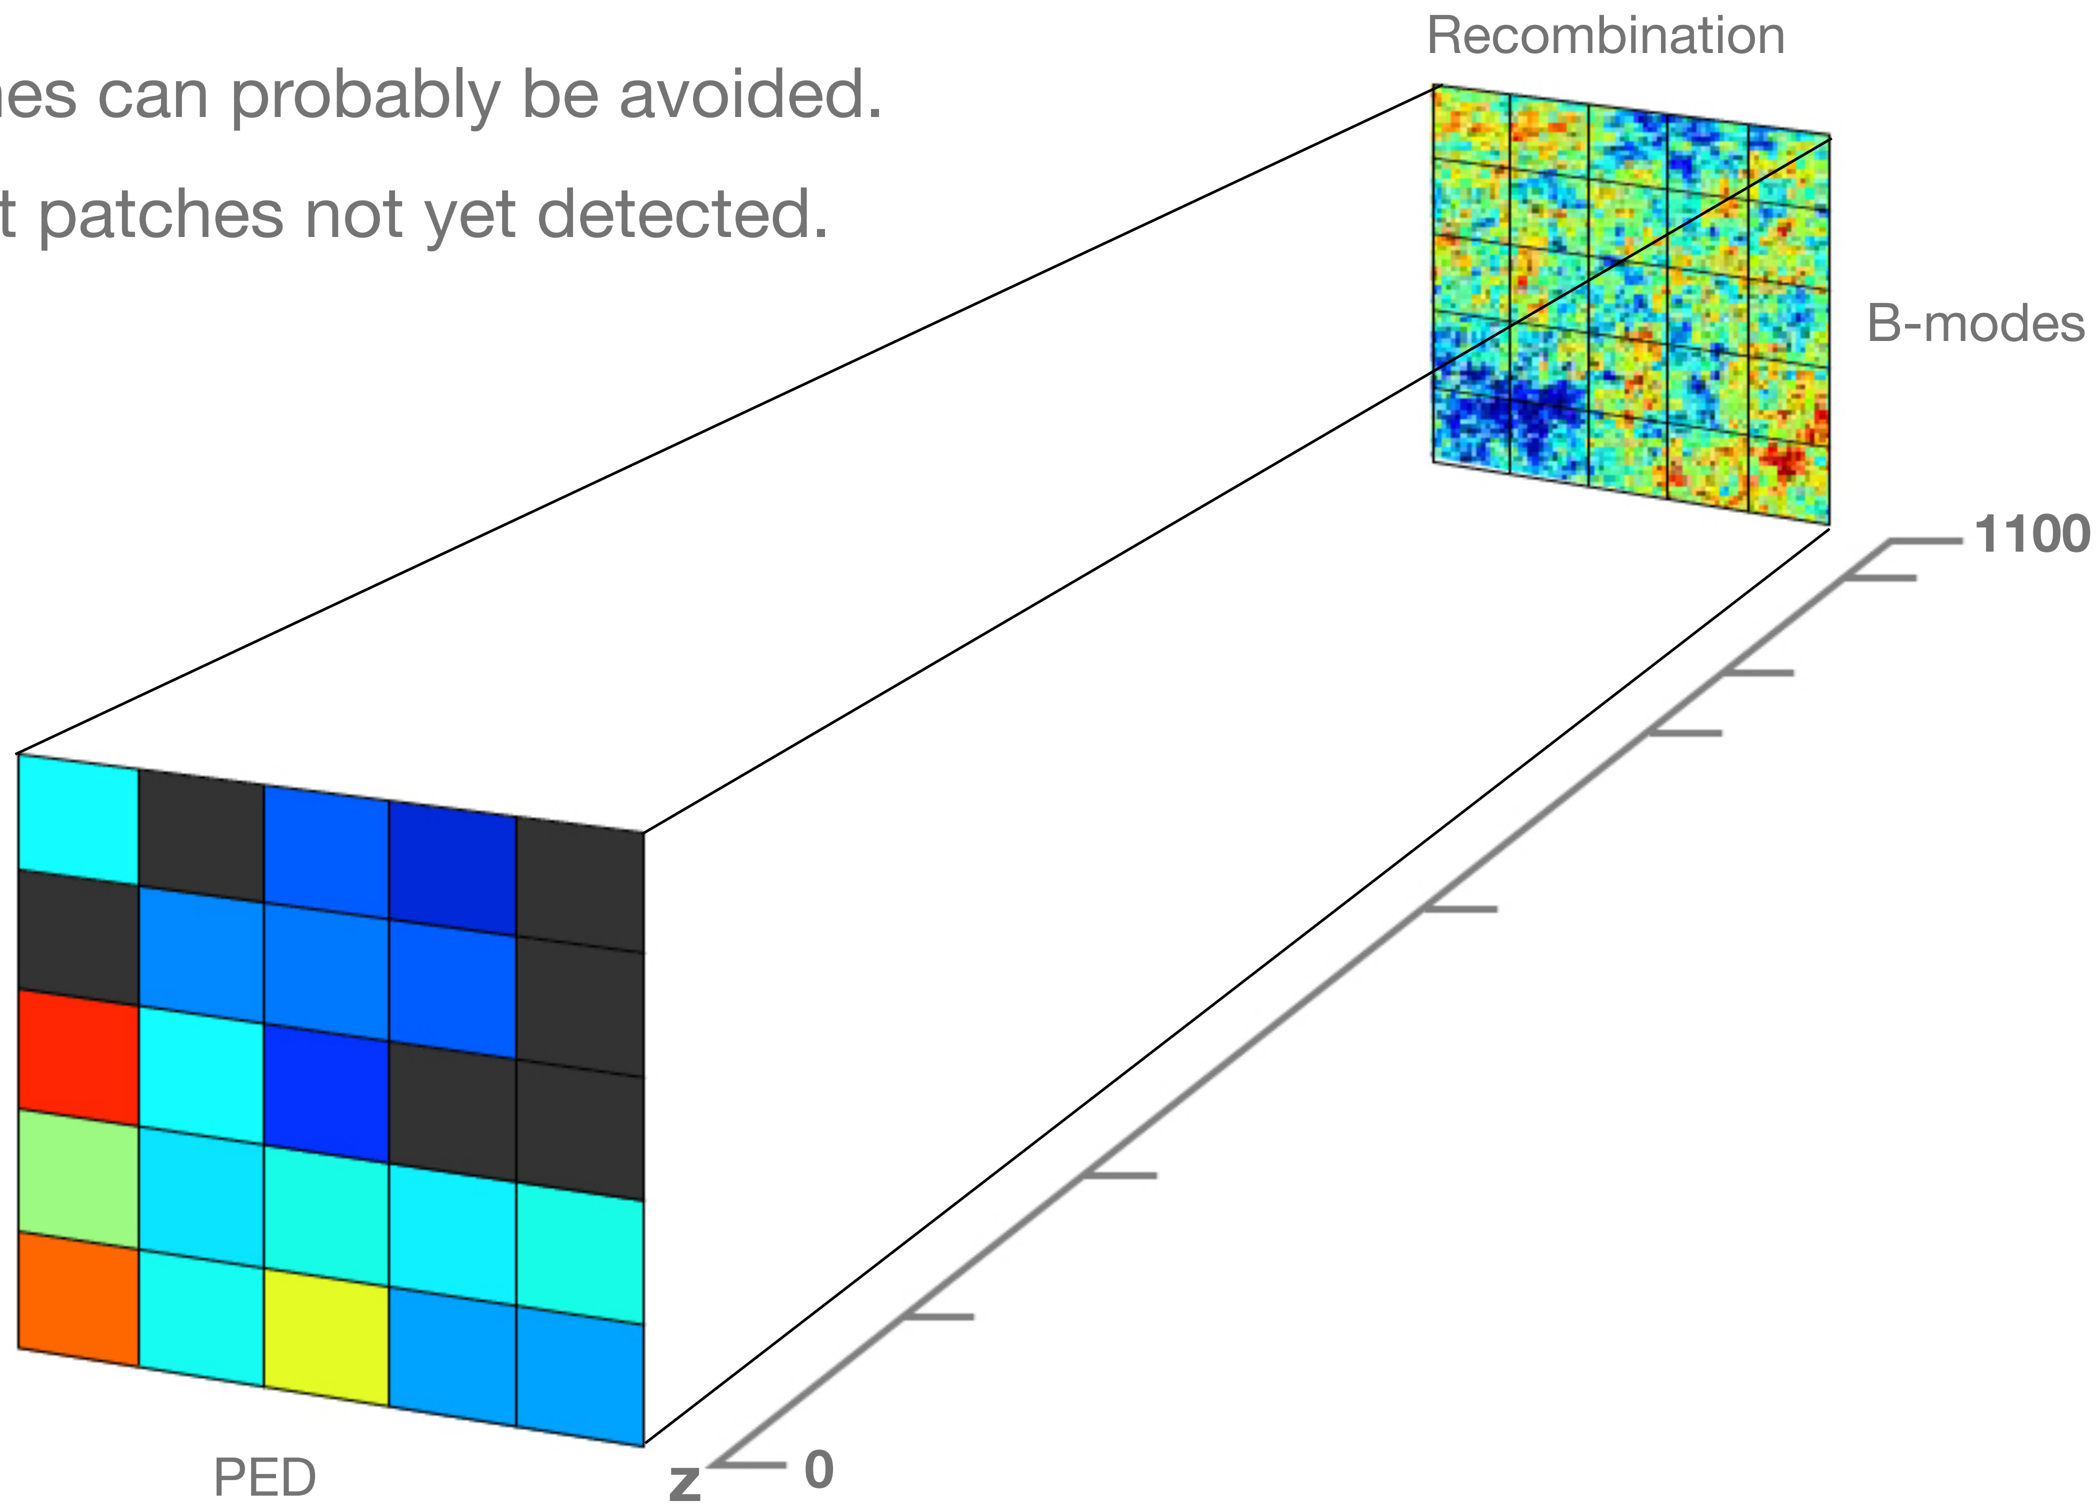

Exploration vs. Exploitation: B-mode Surveys

- Worst patches can probably be avoided.
- But cleanest patches not yet detected.

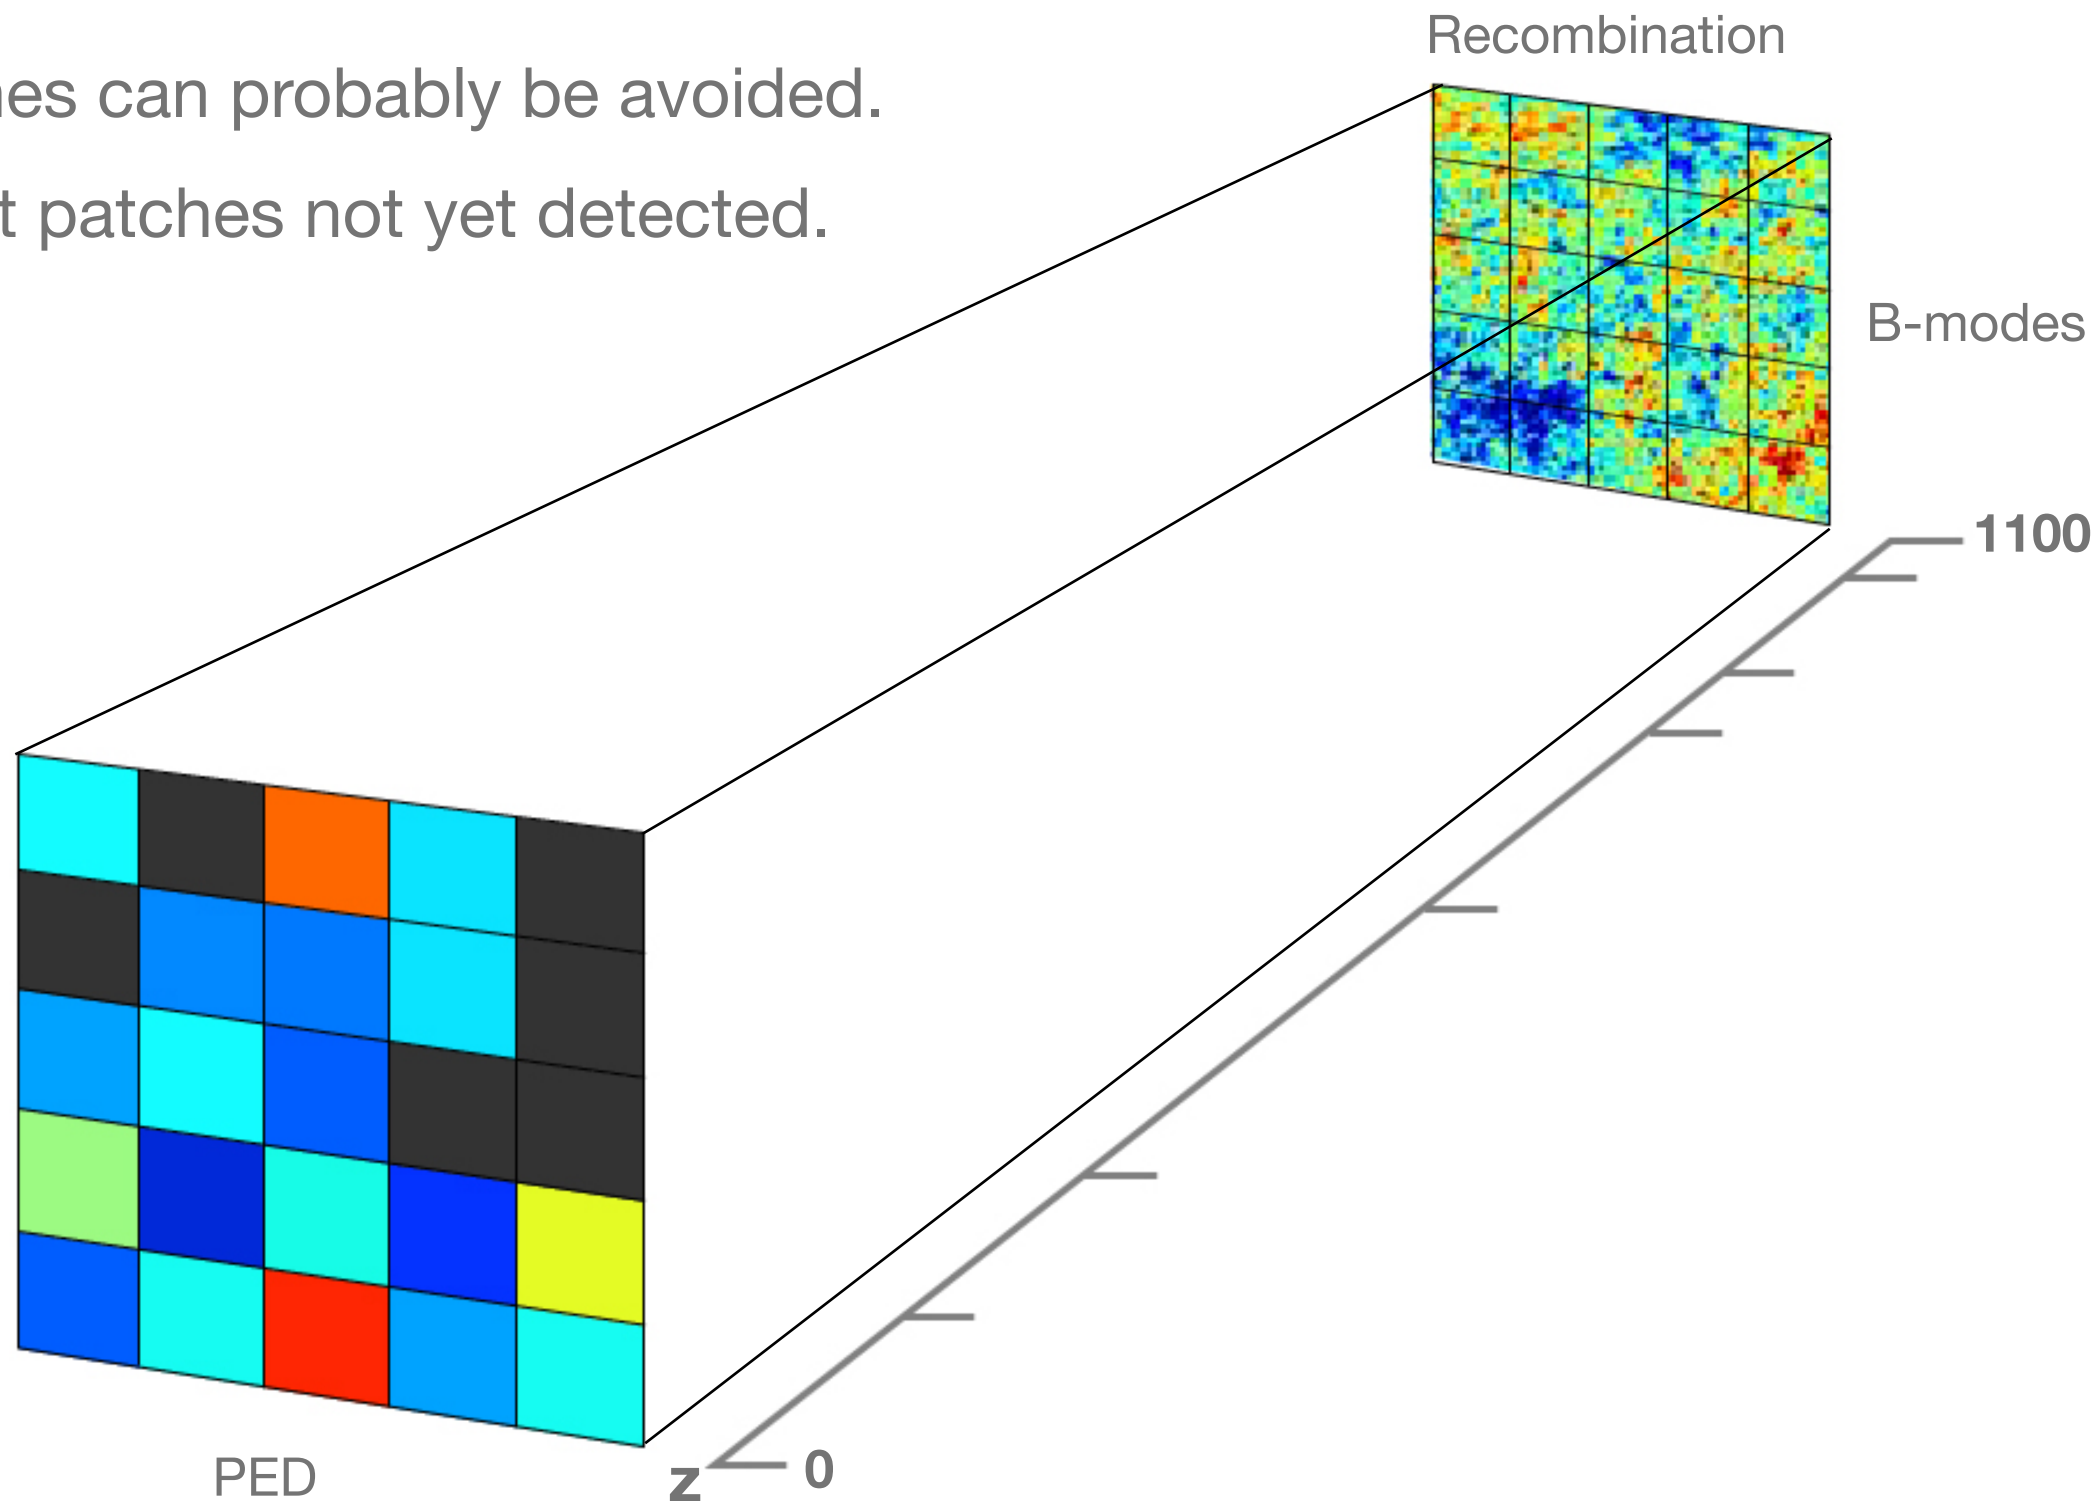

Exploration vs. Exploitation: B-mode Surveys

- Worst patches can probably be avoided.
- But cleanest patches not yet detected.

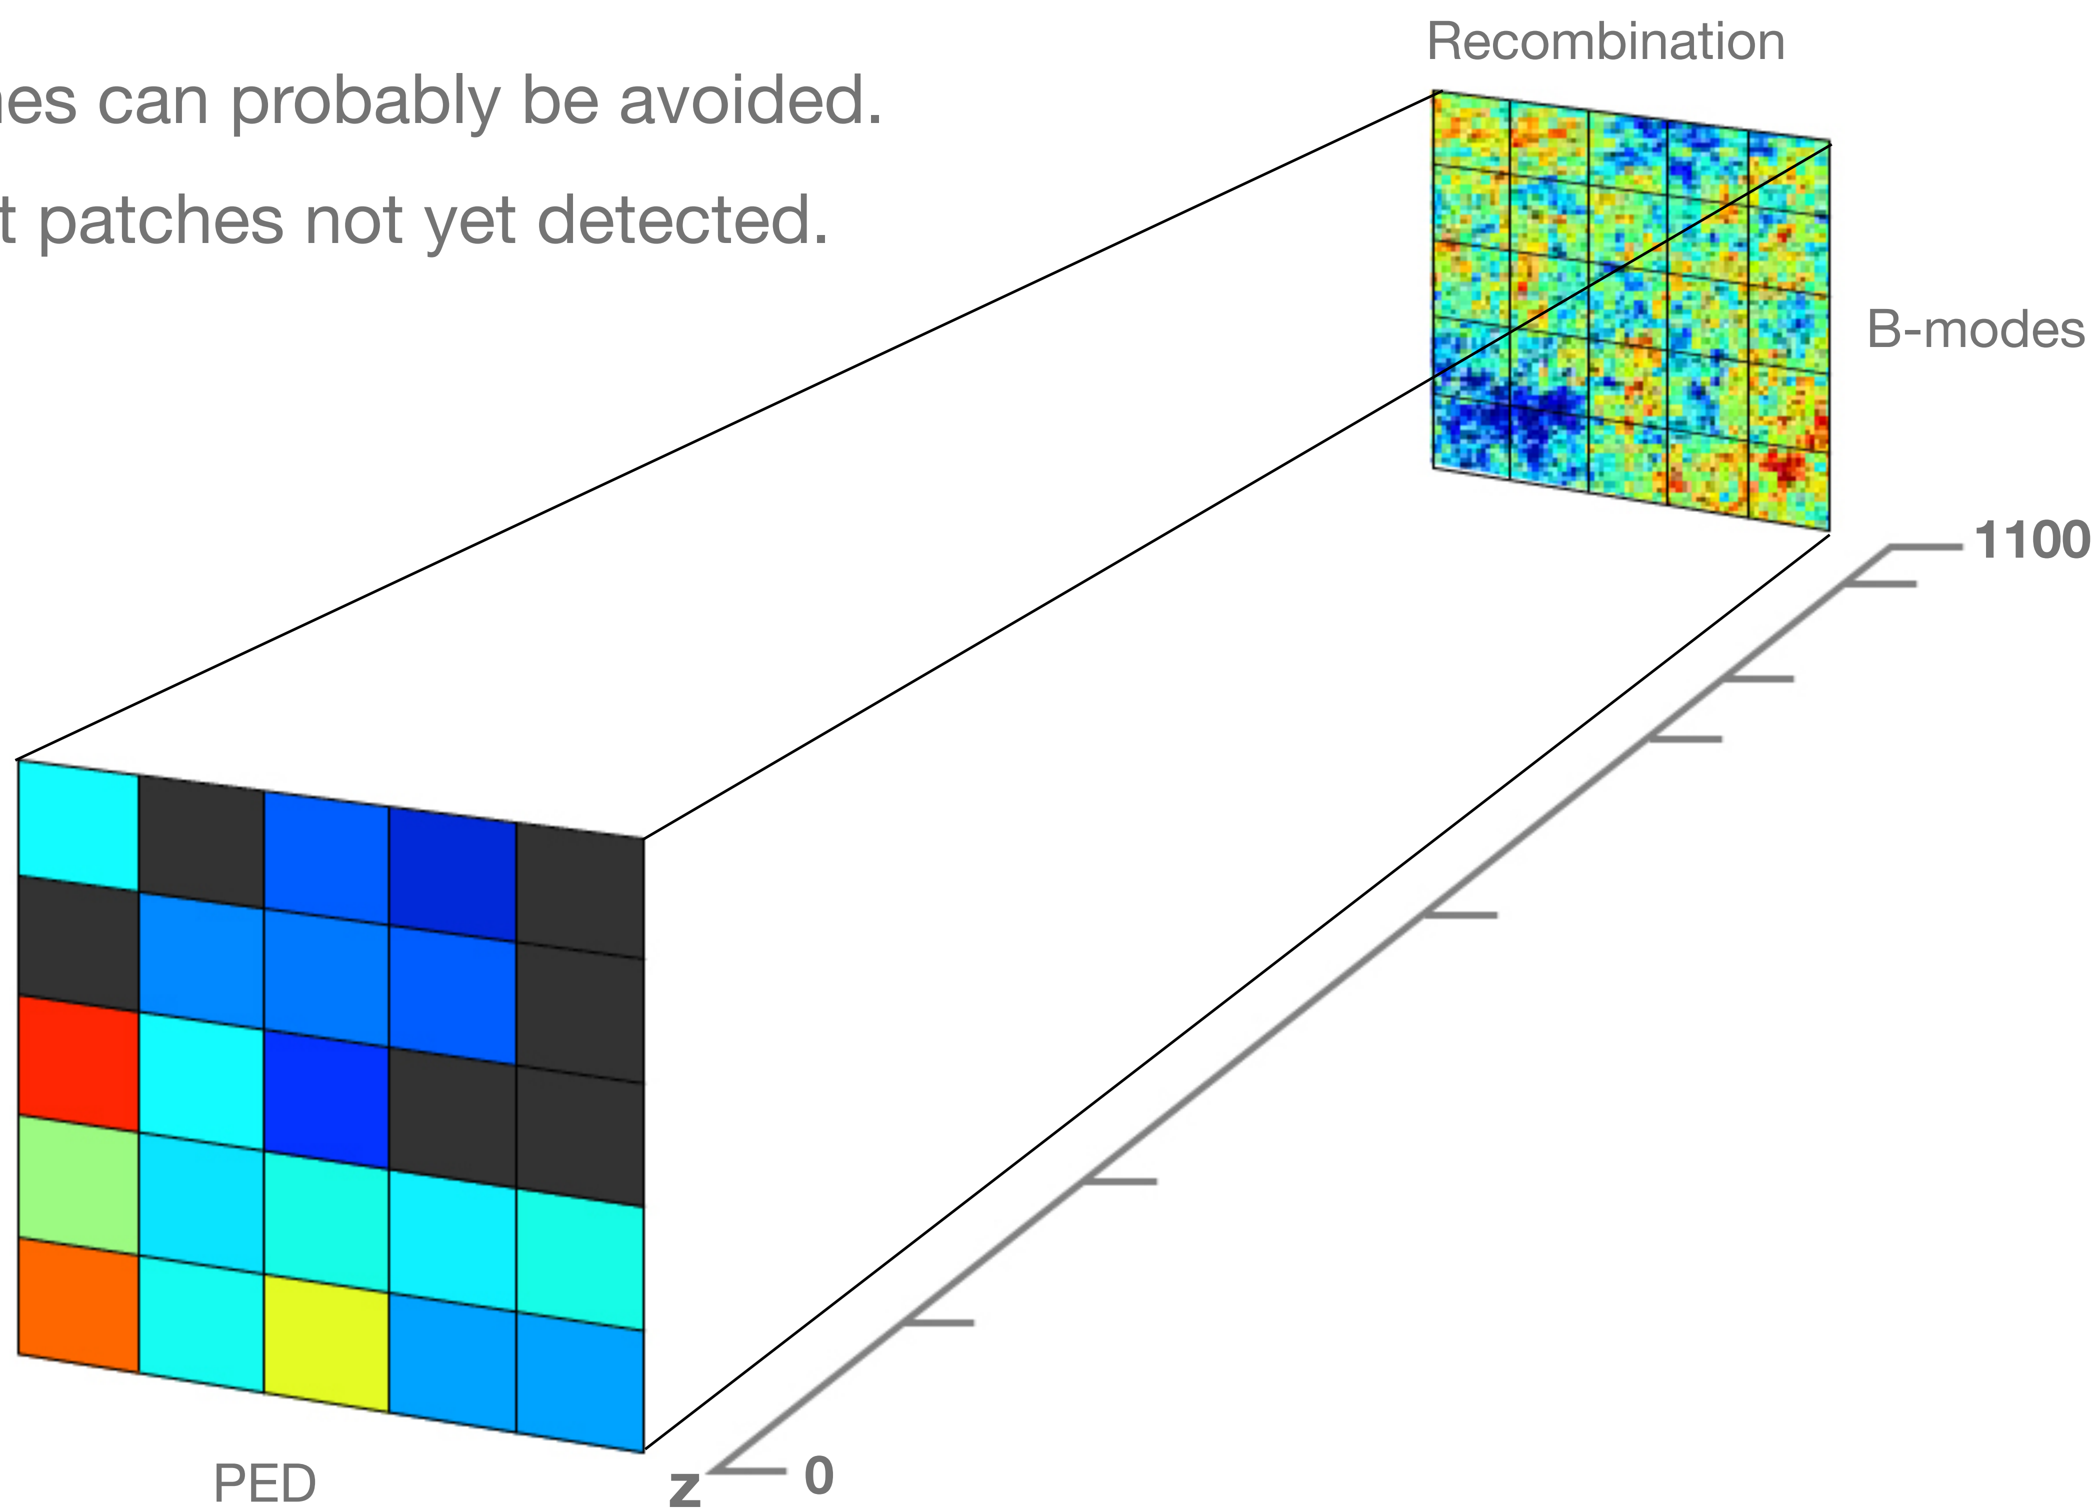

Exploration vs. Exploitation: B-mode Surveys

- Worst patches can probably be avoided.
- But cleanest patches not yet detected.

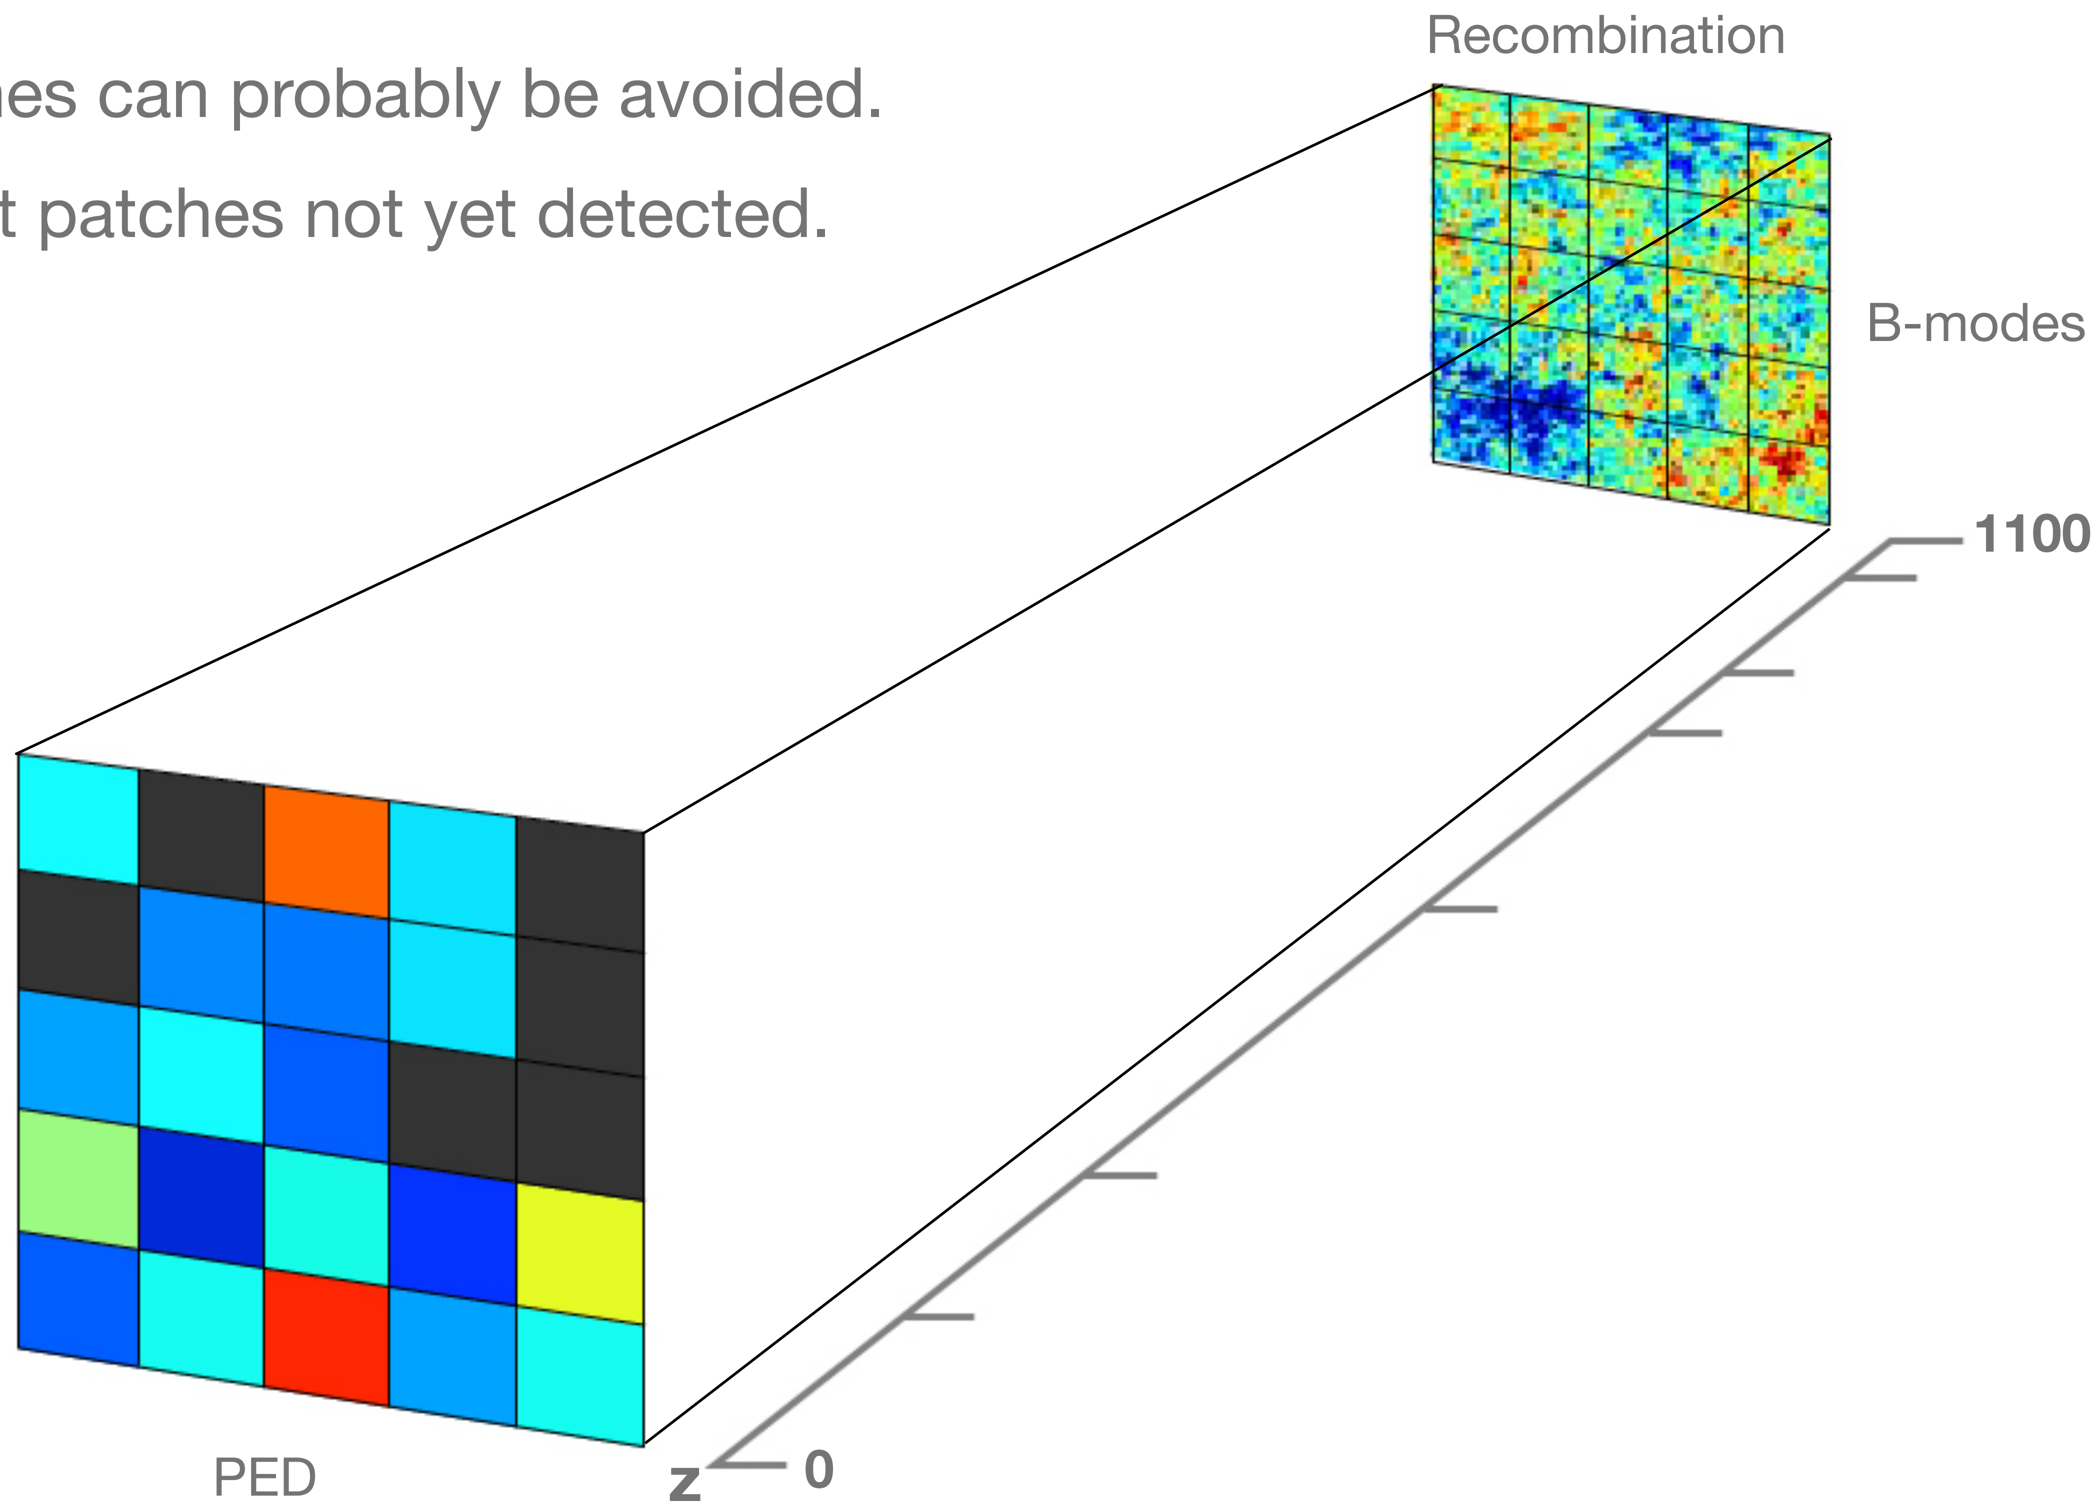

Exploration vs. Exploitation: B-mode Surveys

- Worst patches can probably be avoided.
- But cleanest patches not yet detected.

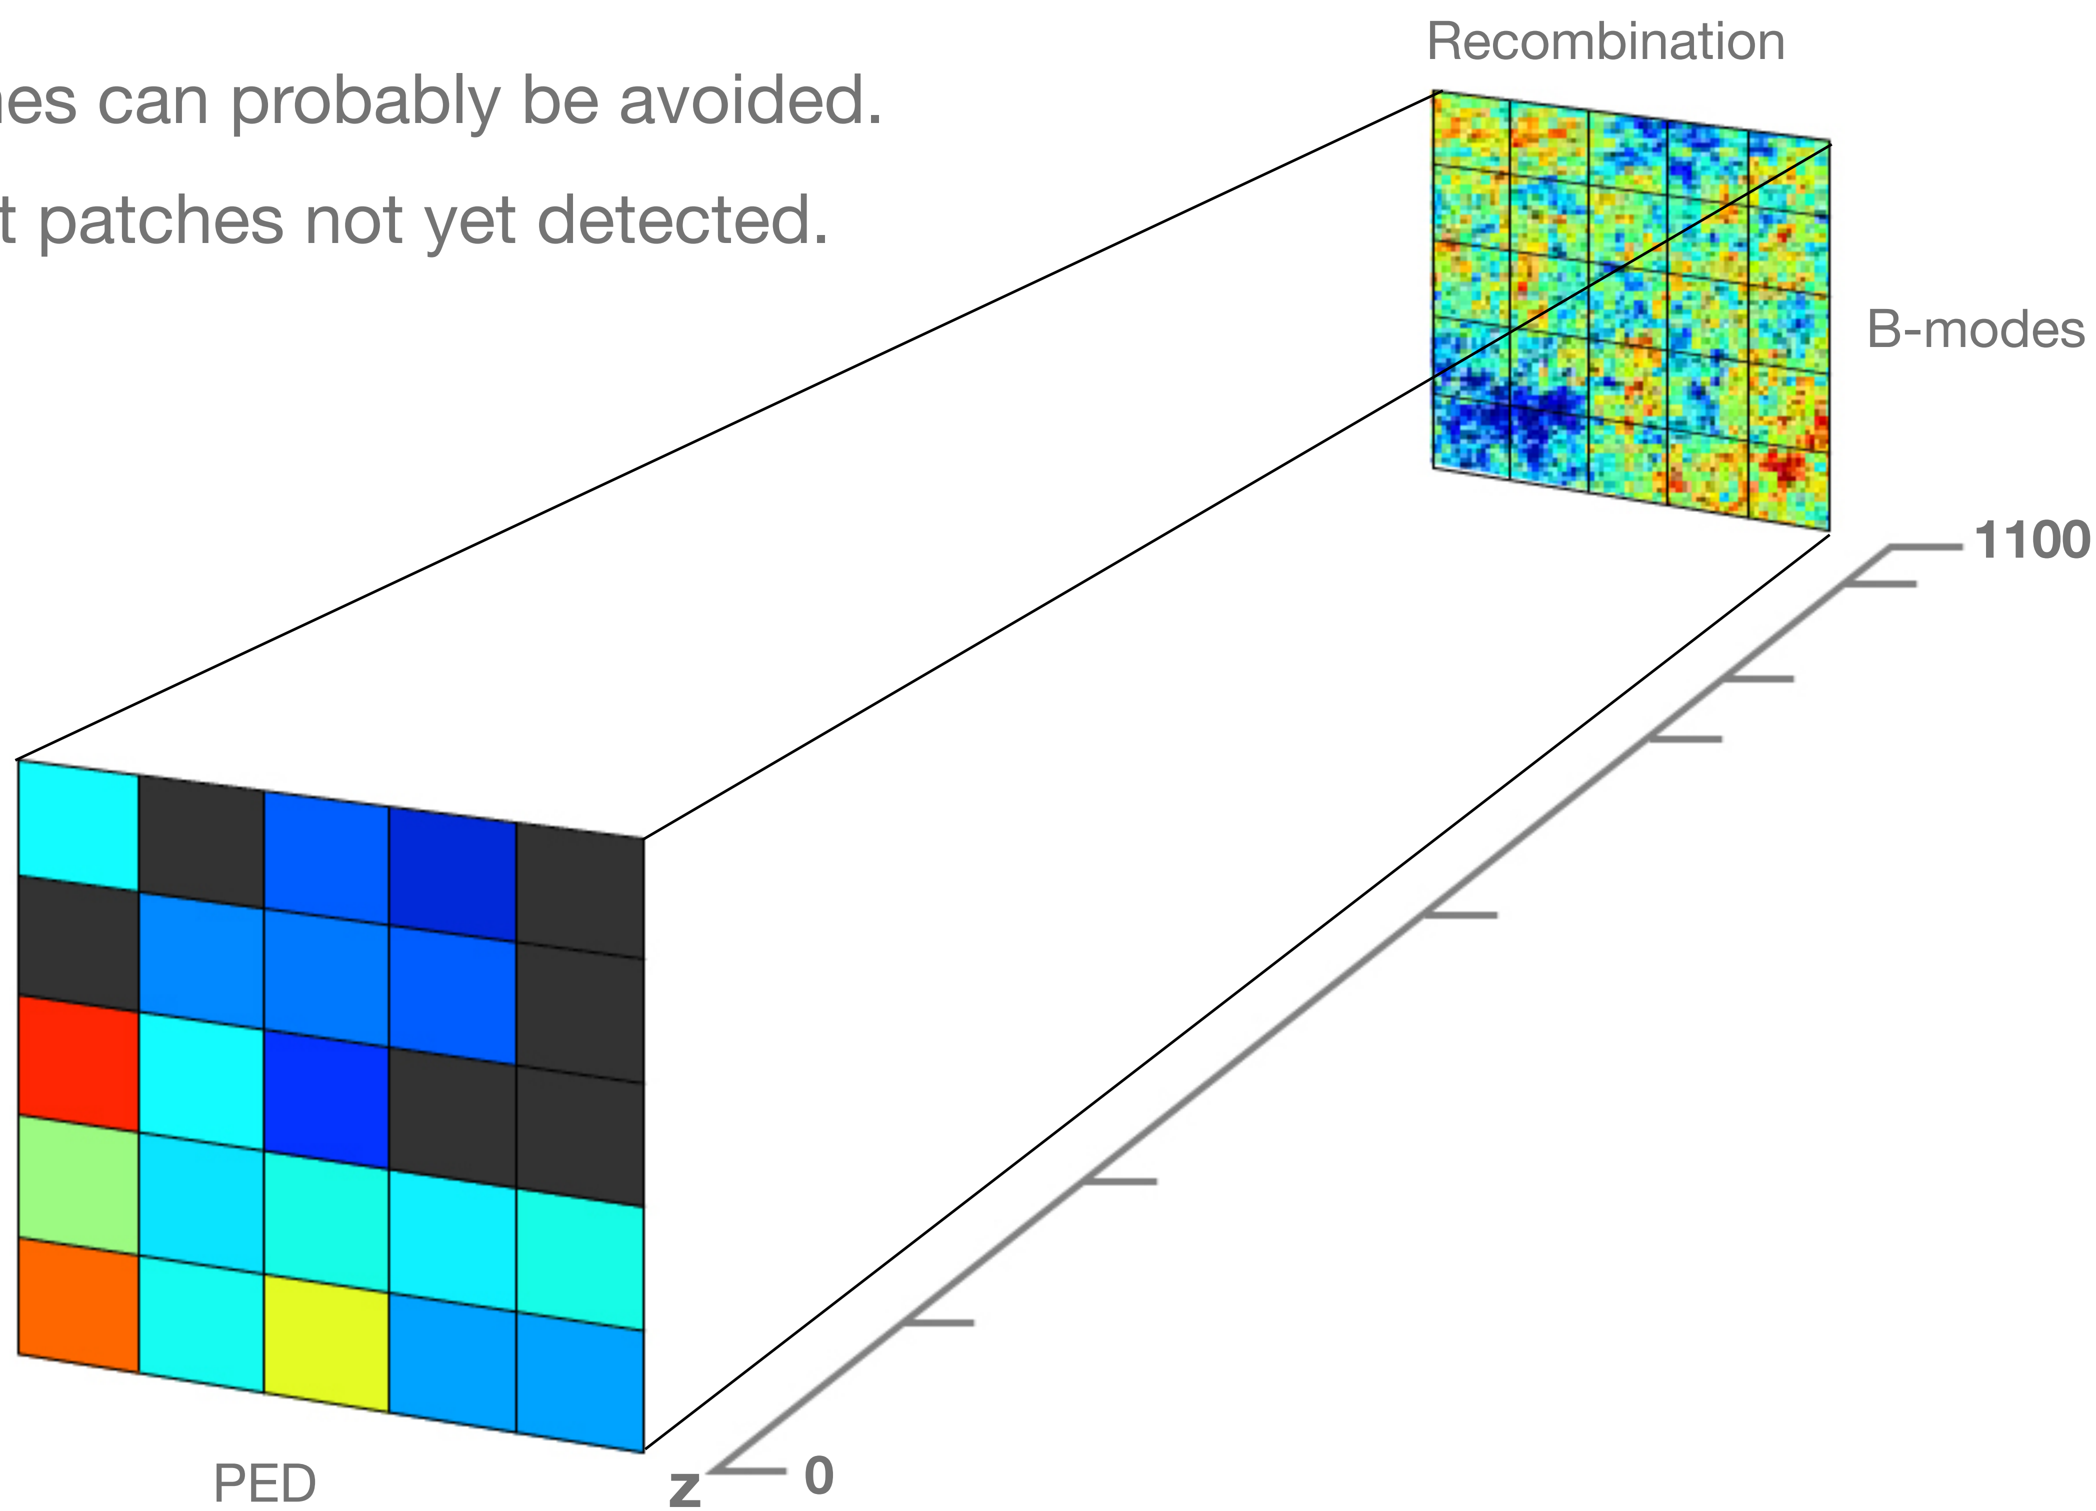

Exploration vs. Exploitation: B-mode Surveys

- Worst patches can probably be avoided.
- But cleanest patches not yet detected.
- If we find the cleanest patches...

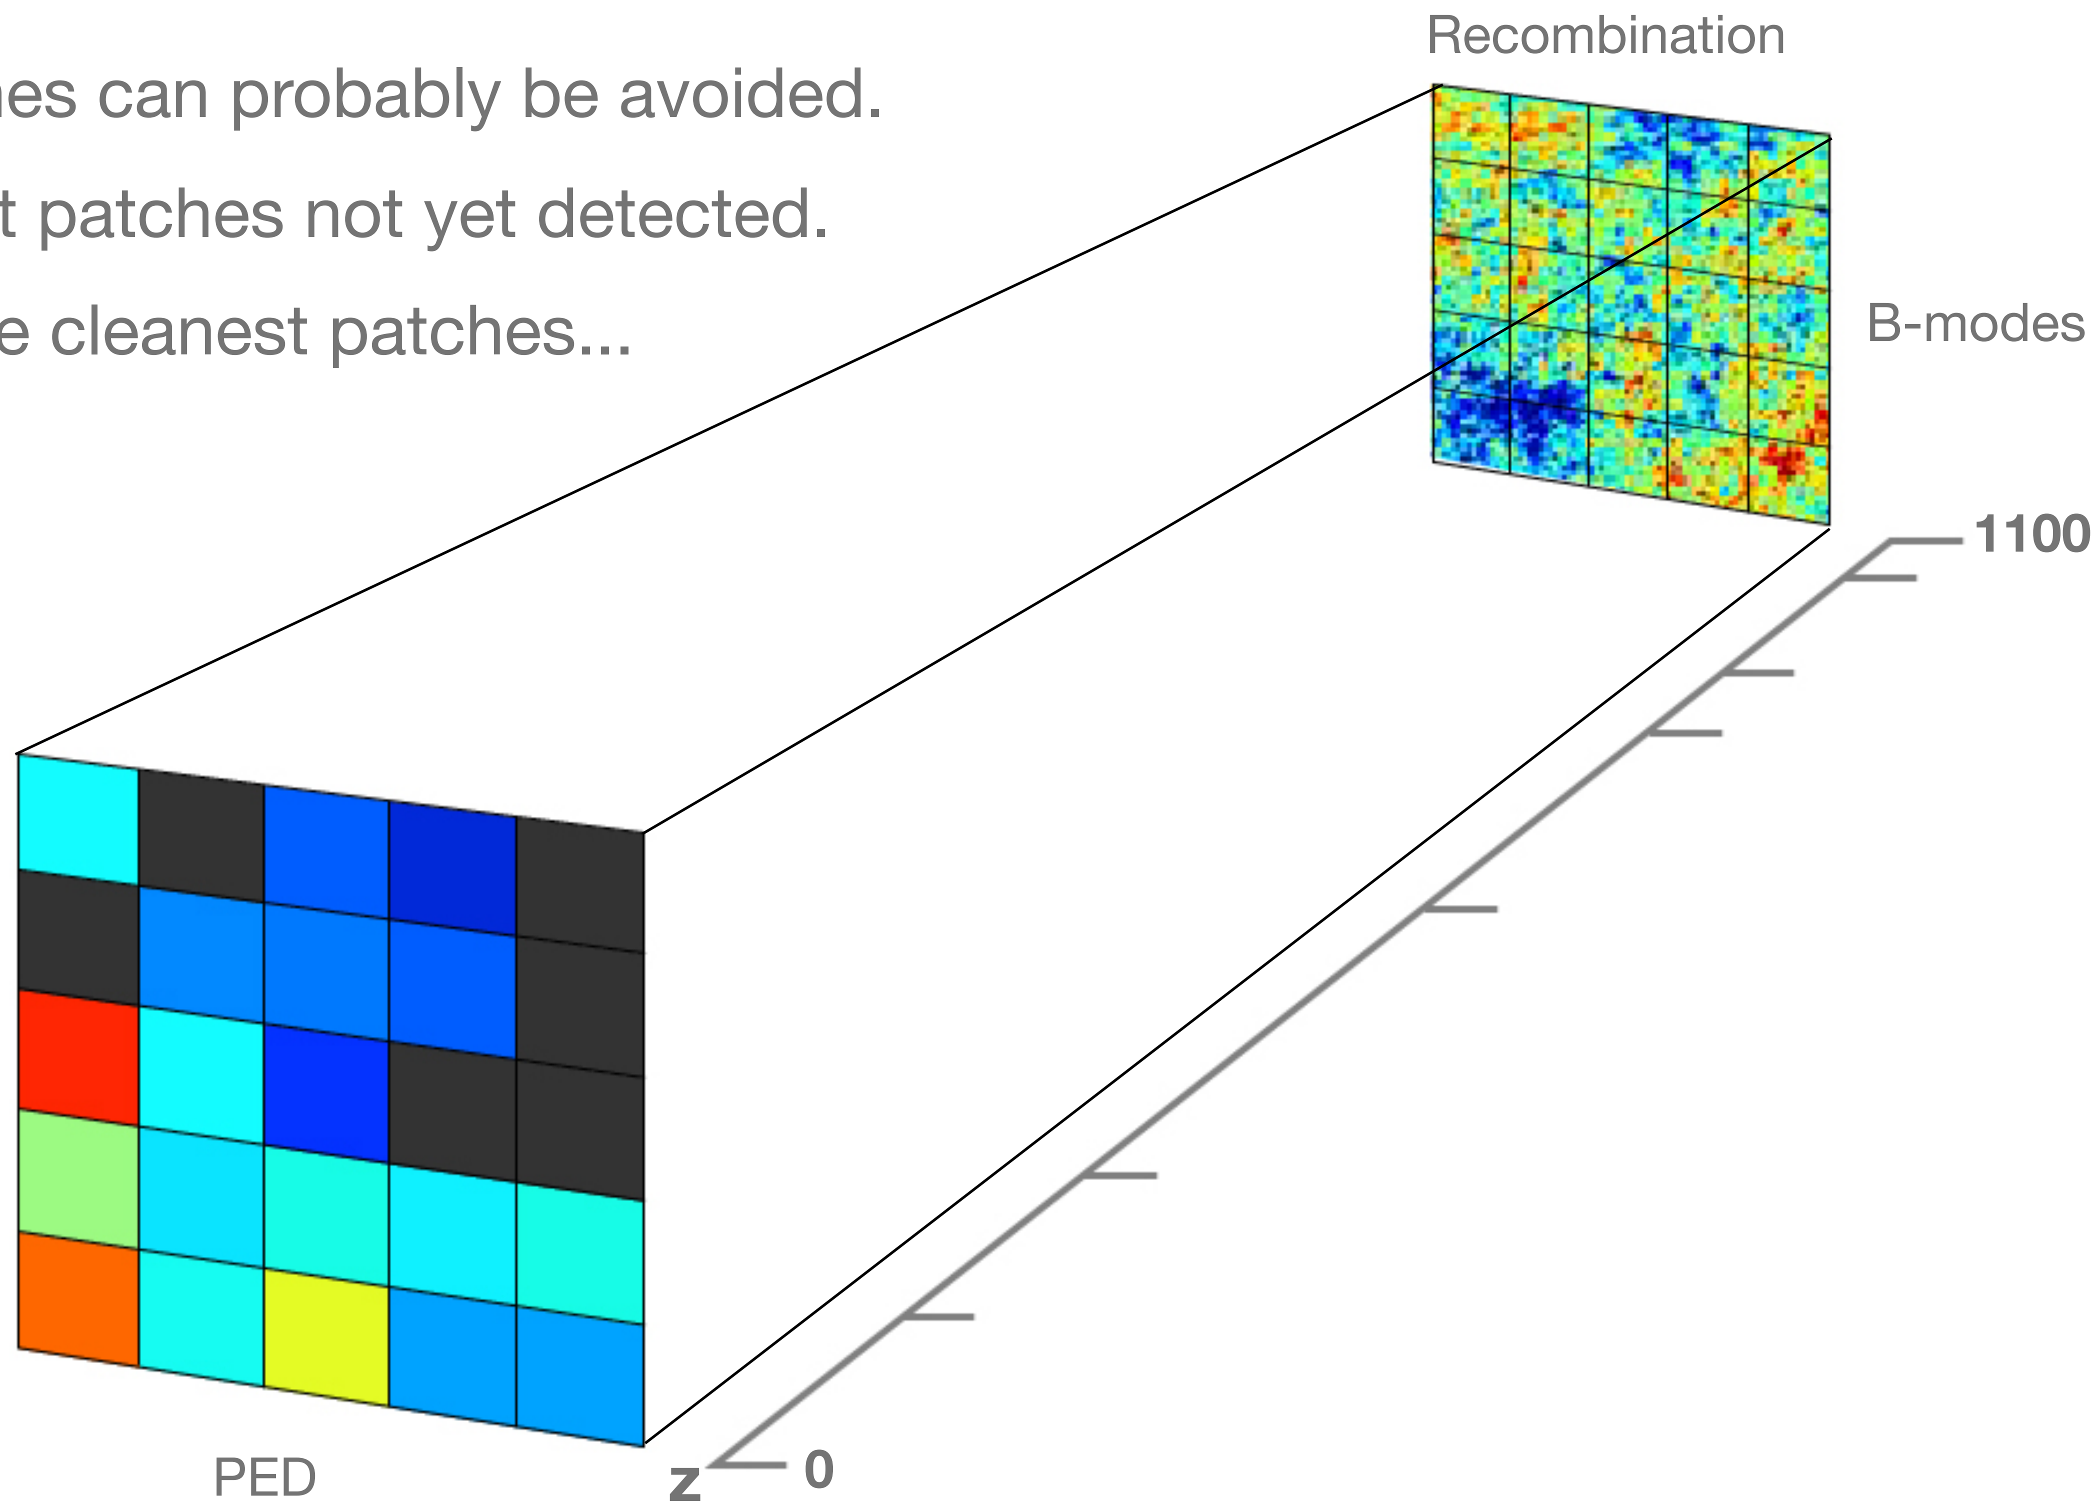

Exploration vs. Exploitation: B-mode Surveys

- Worst patches can probably be avoided.
- But cleanest patches not yet detected.
- If we find the cleanest patches...

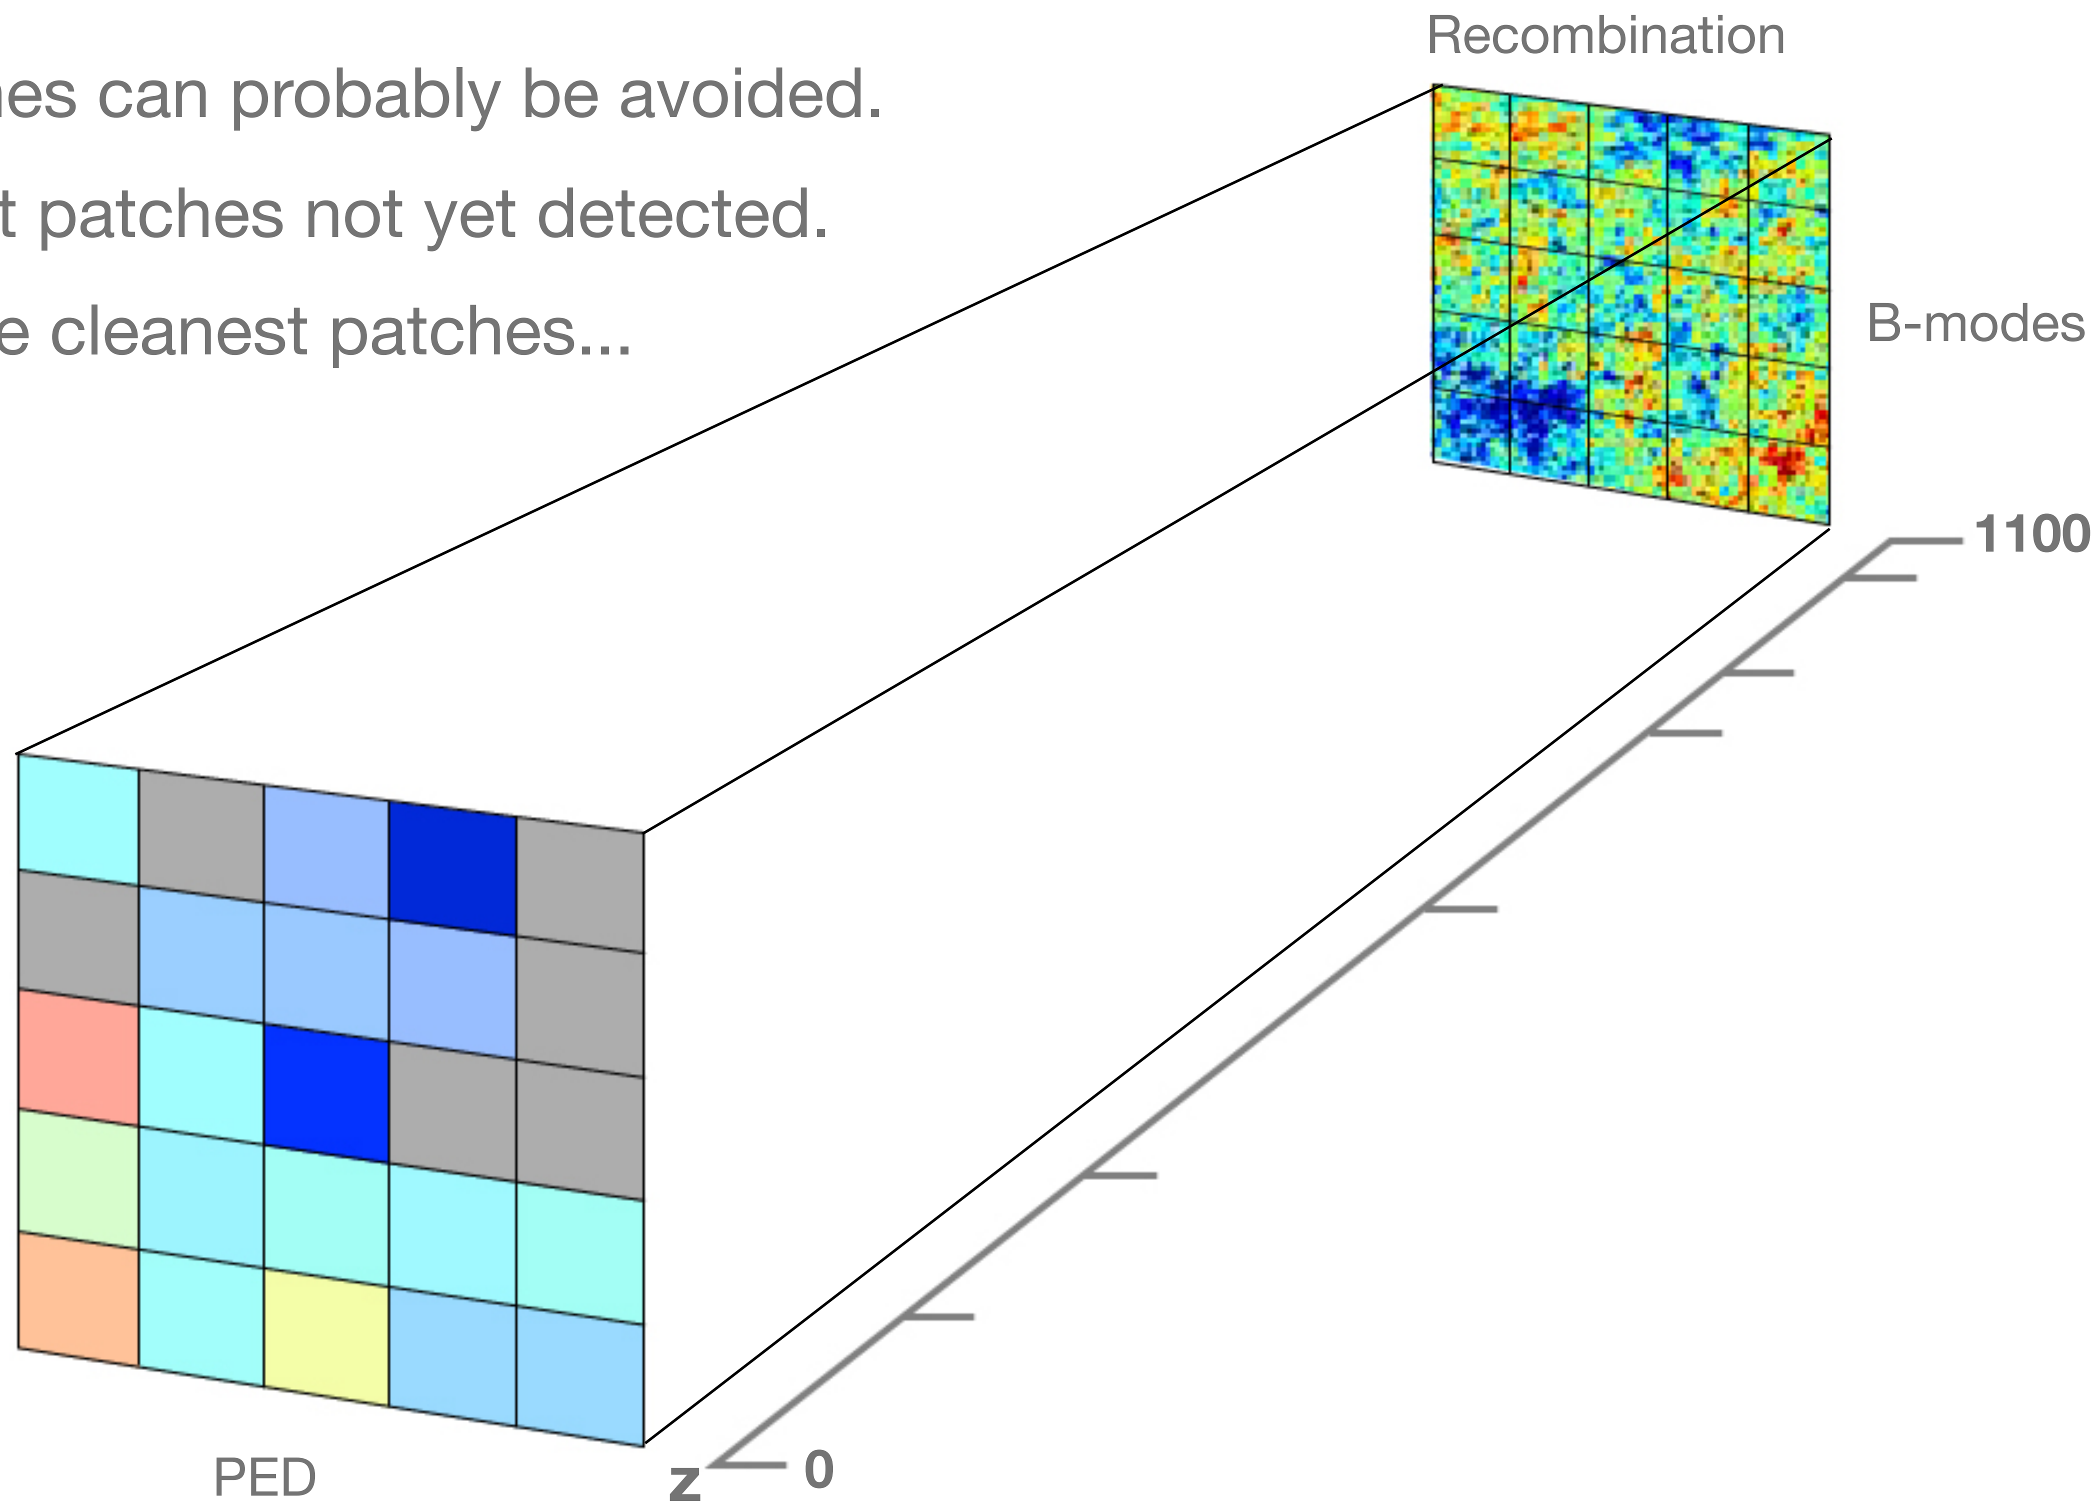

Exploration vs. Exploitation: B-mode Surveys

- Worst patches can probably be avoided.
- But cleanest patches not yet detected.
- If we find the cleanest patches...
- Sensitivity might improve!

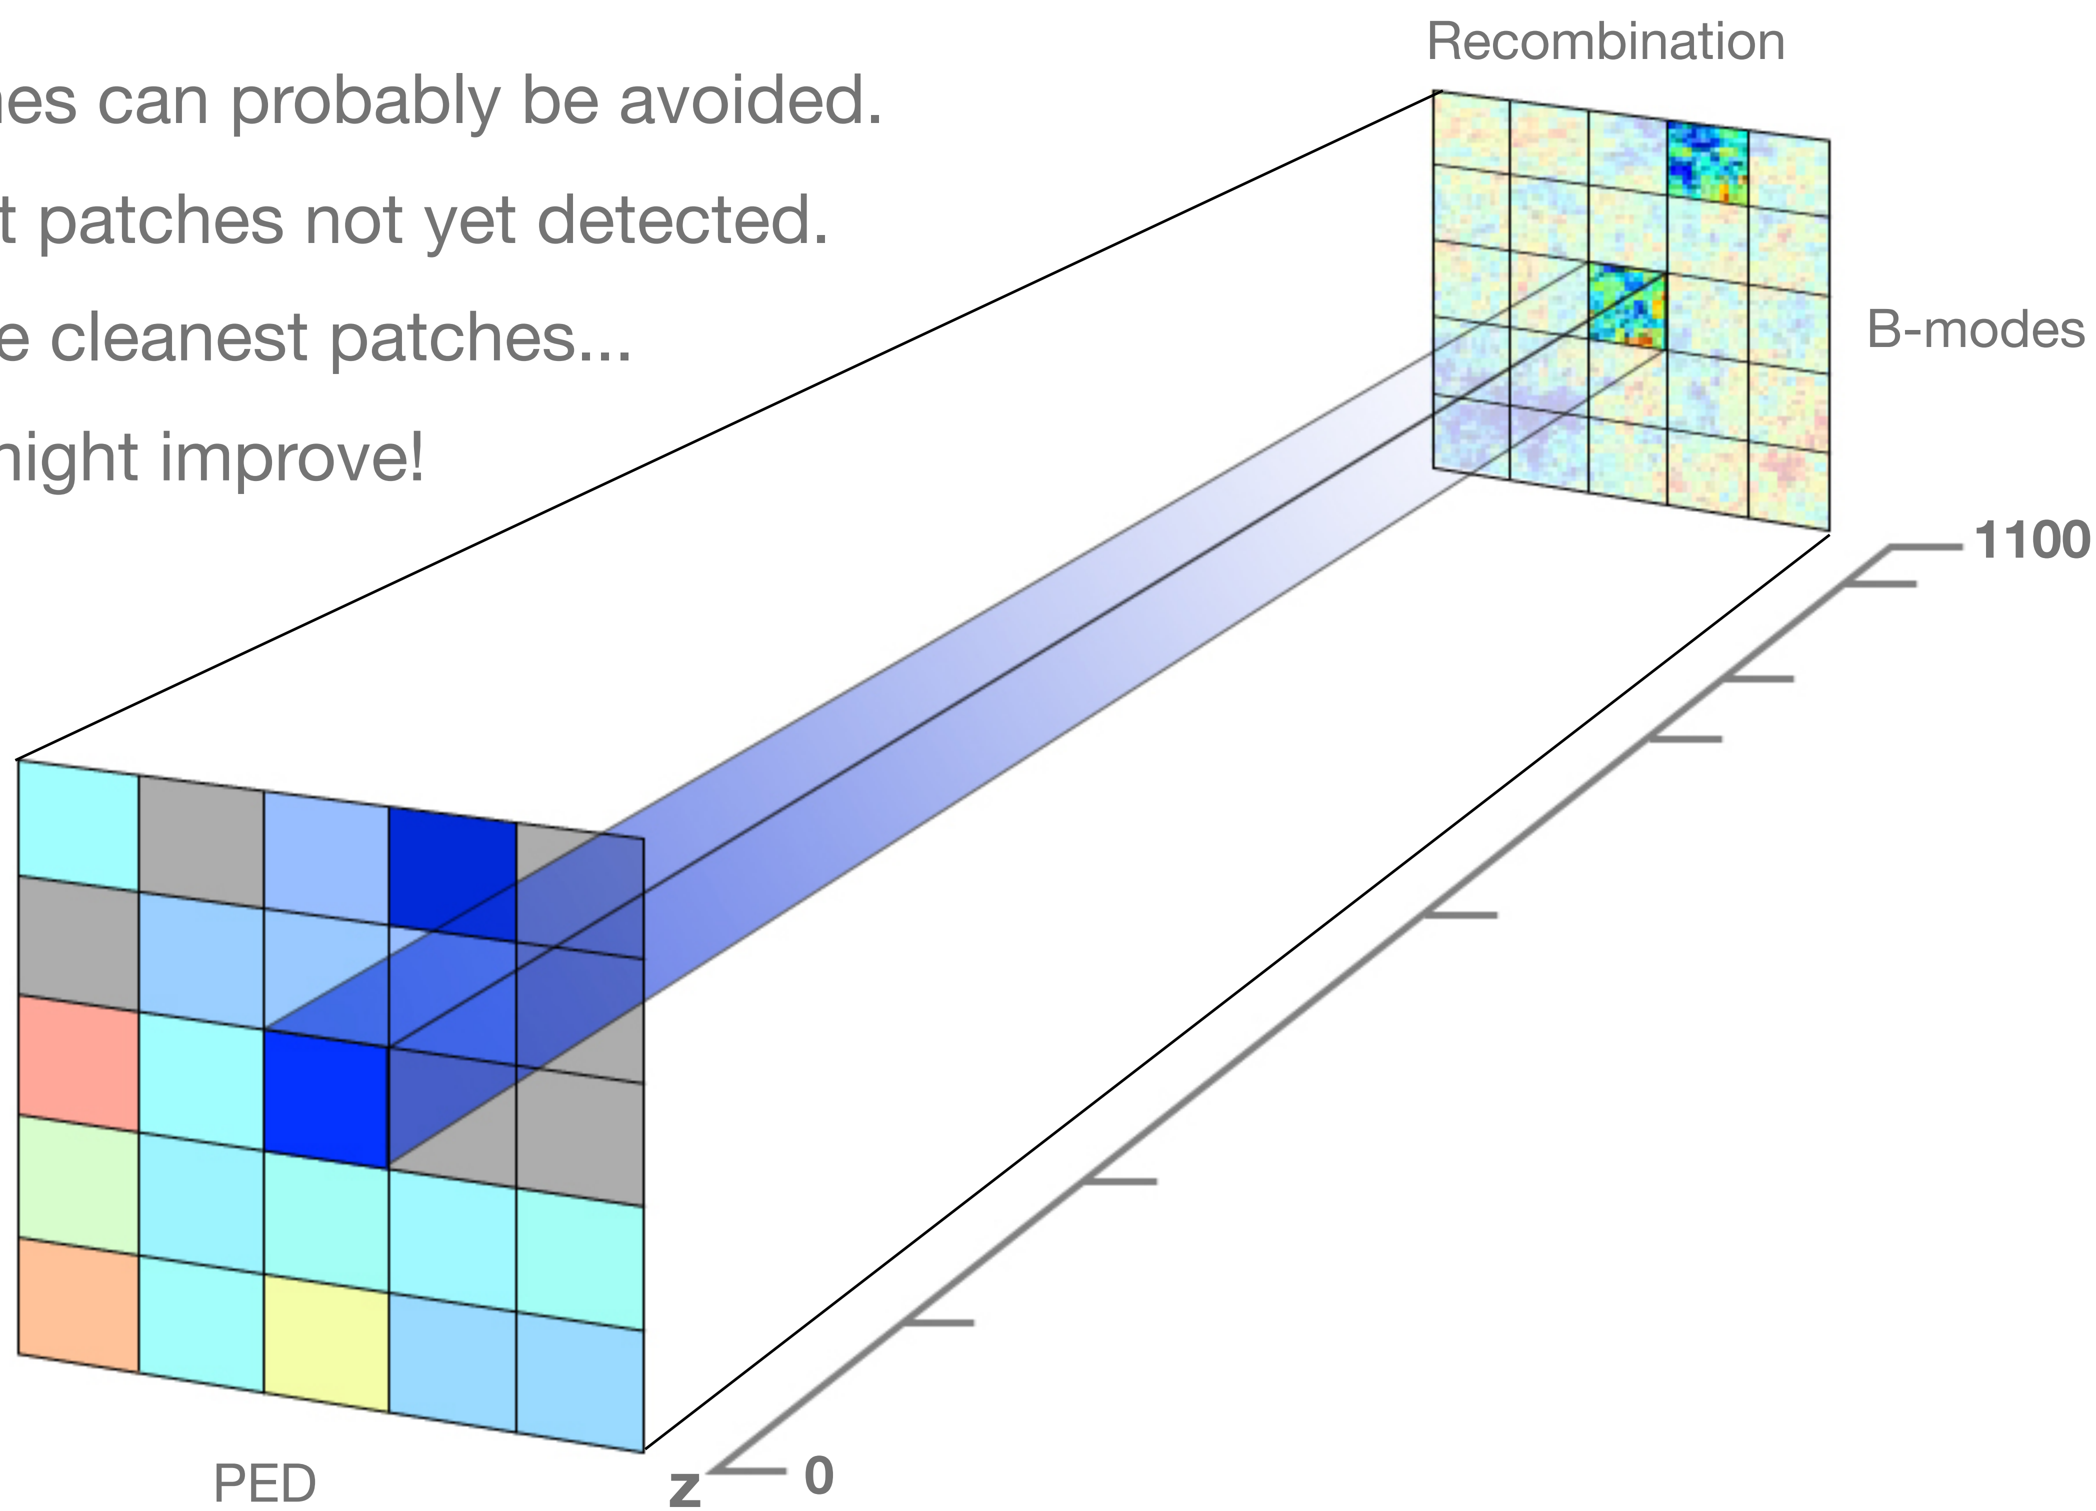

Exploration vs. Exploitation: B-mode Surveys

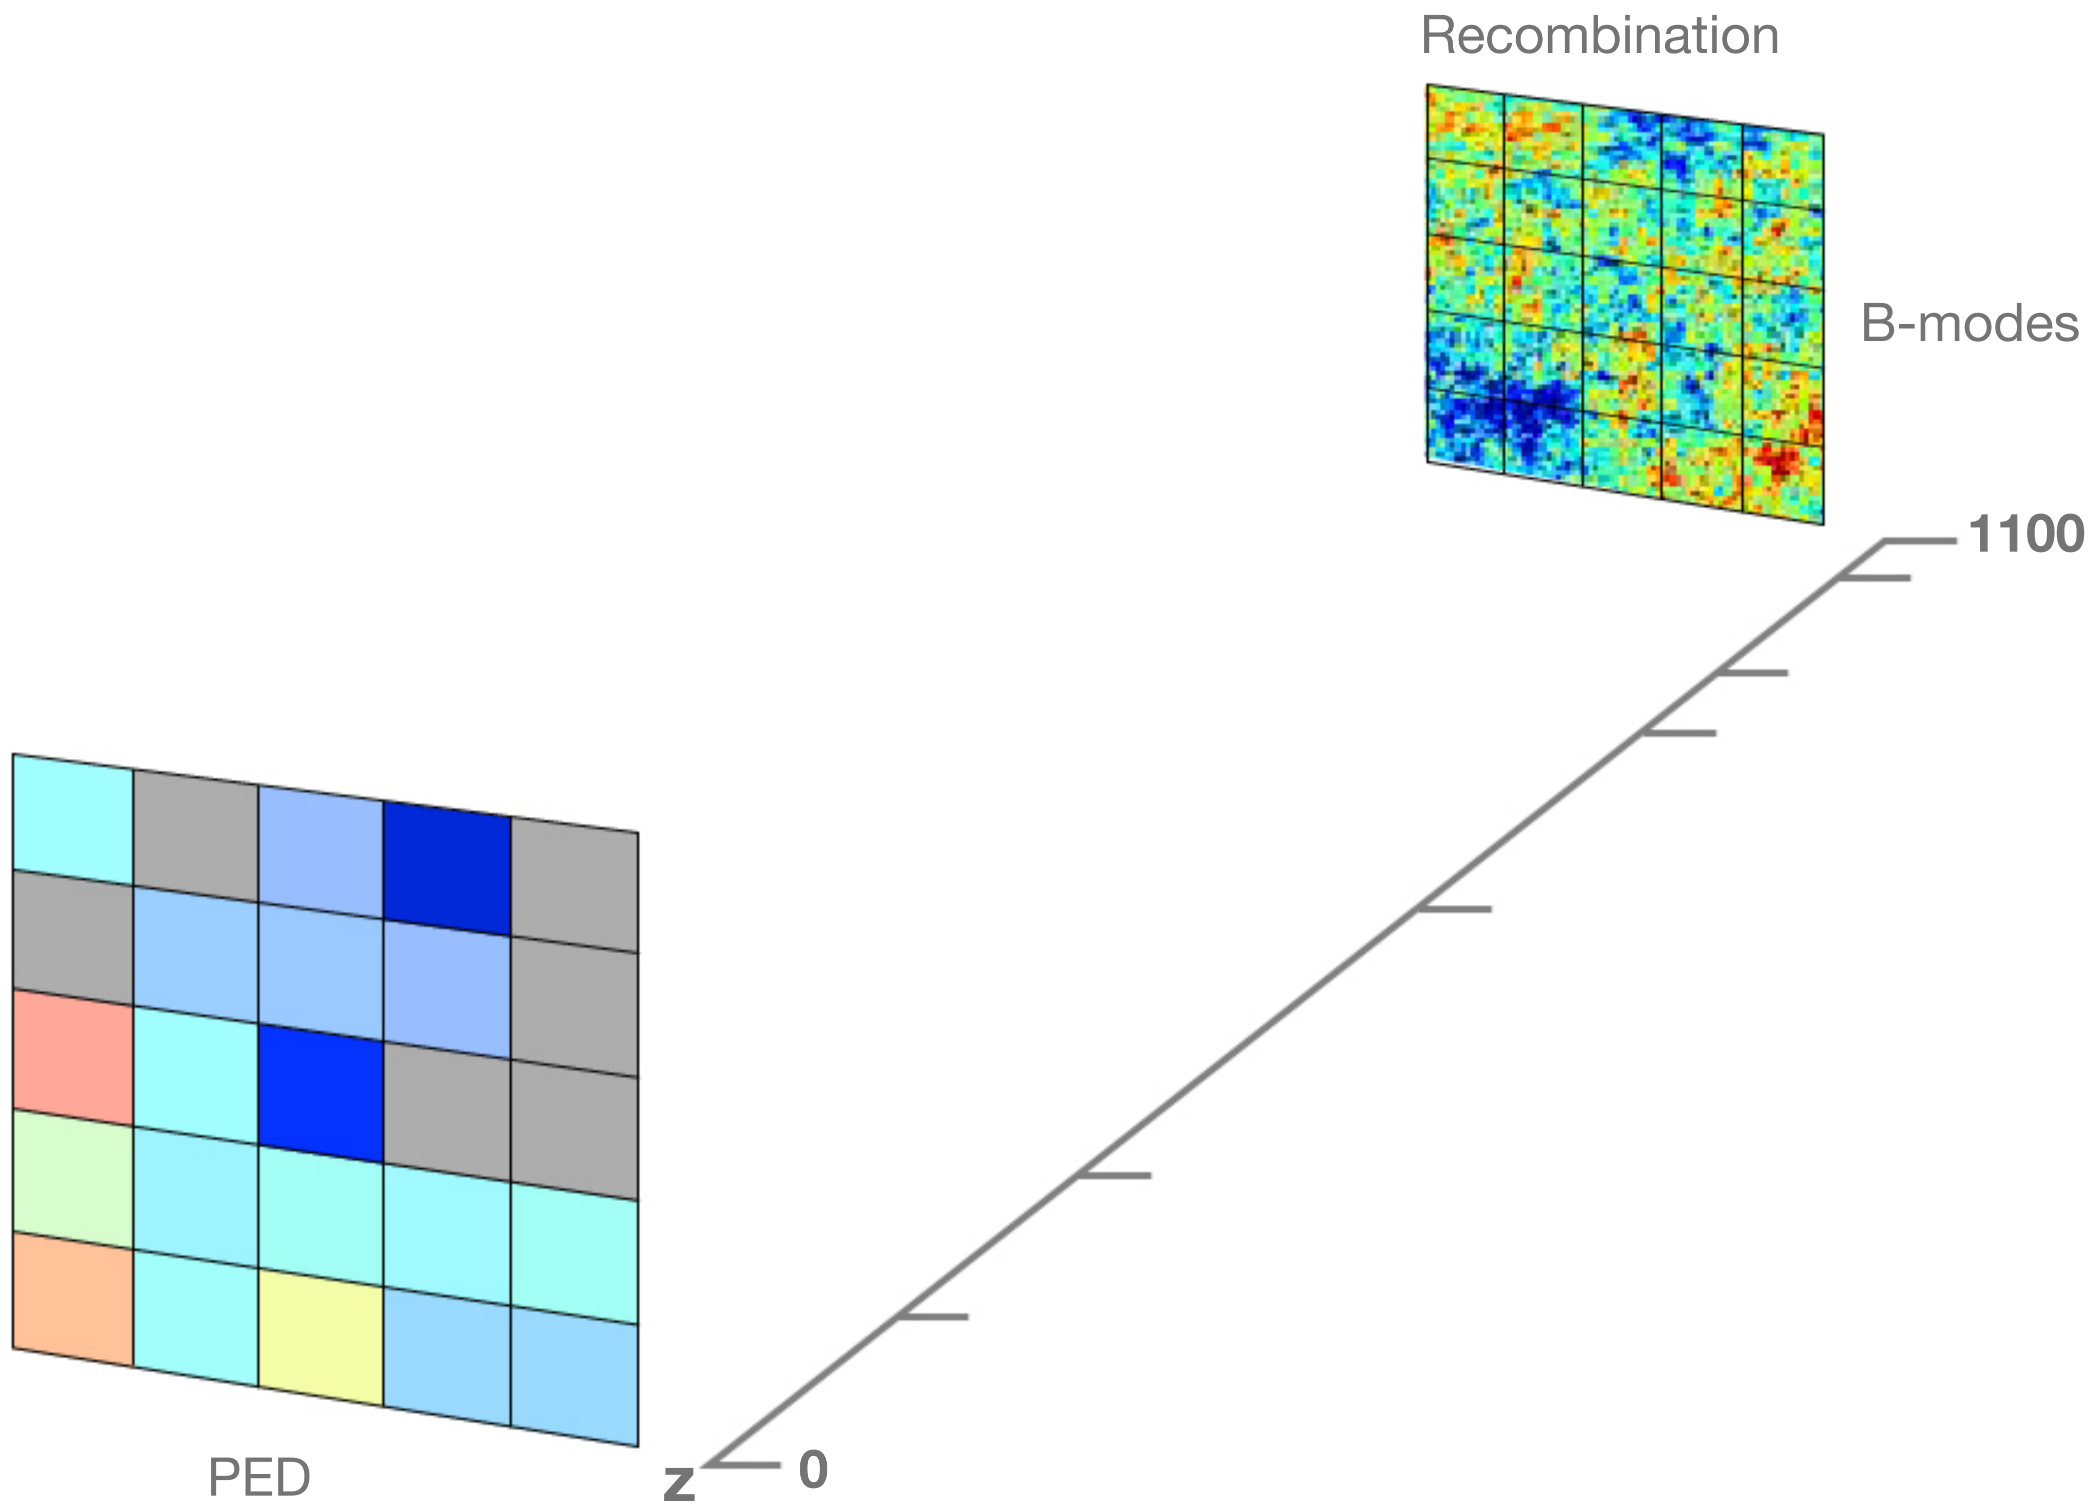

Exploration vs. Exploitation: B-mode Surveys

- Competing signals?

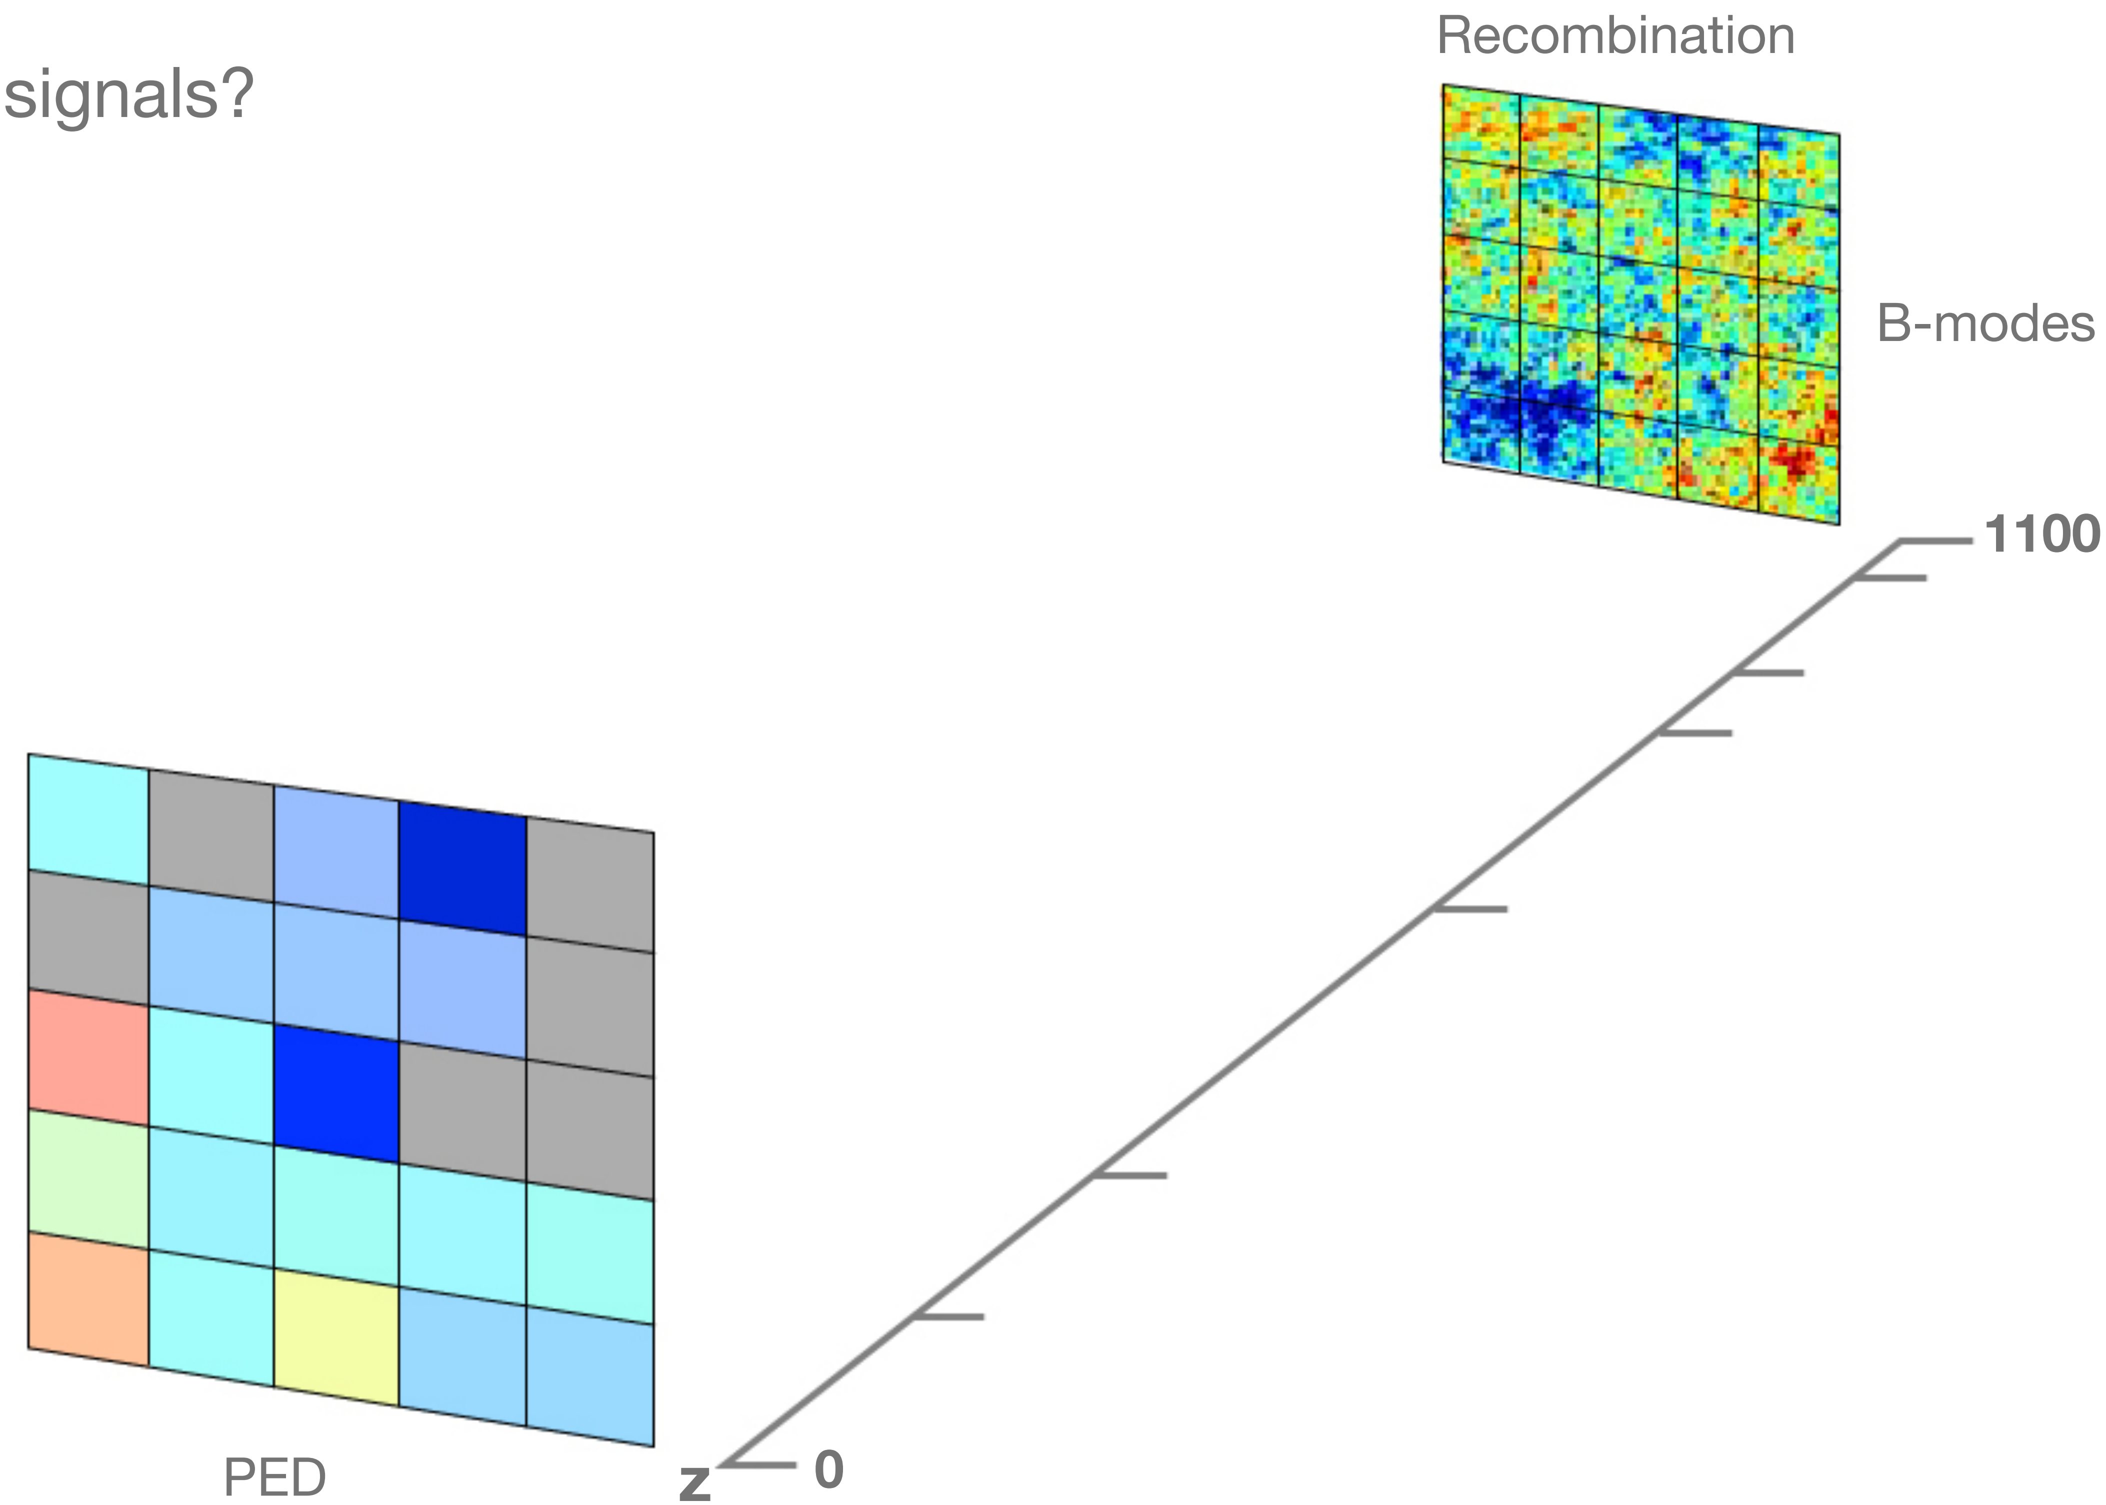

Exploration vs. Exploitation: B-mode Surveys

- Competing signals?
- Lensing converts E-modes into B-modes.

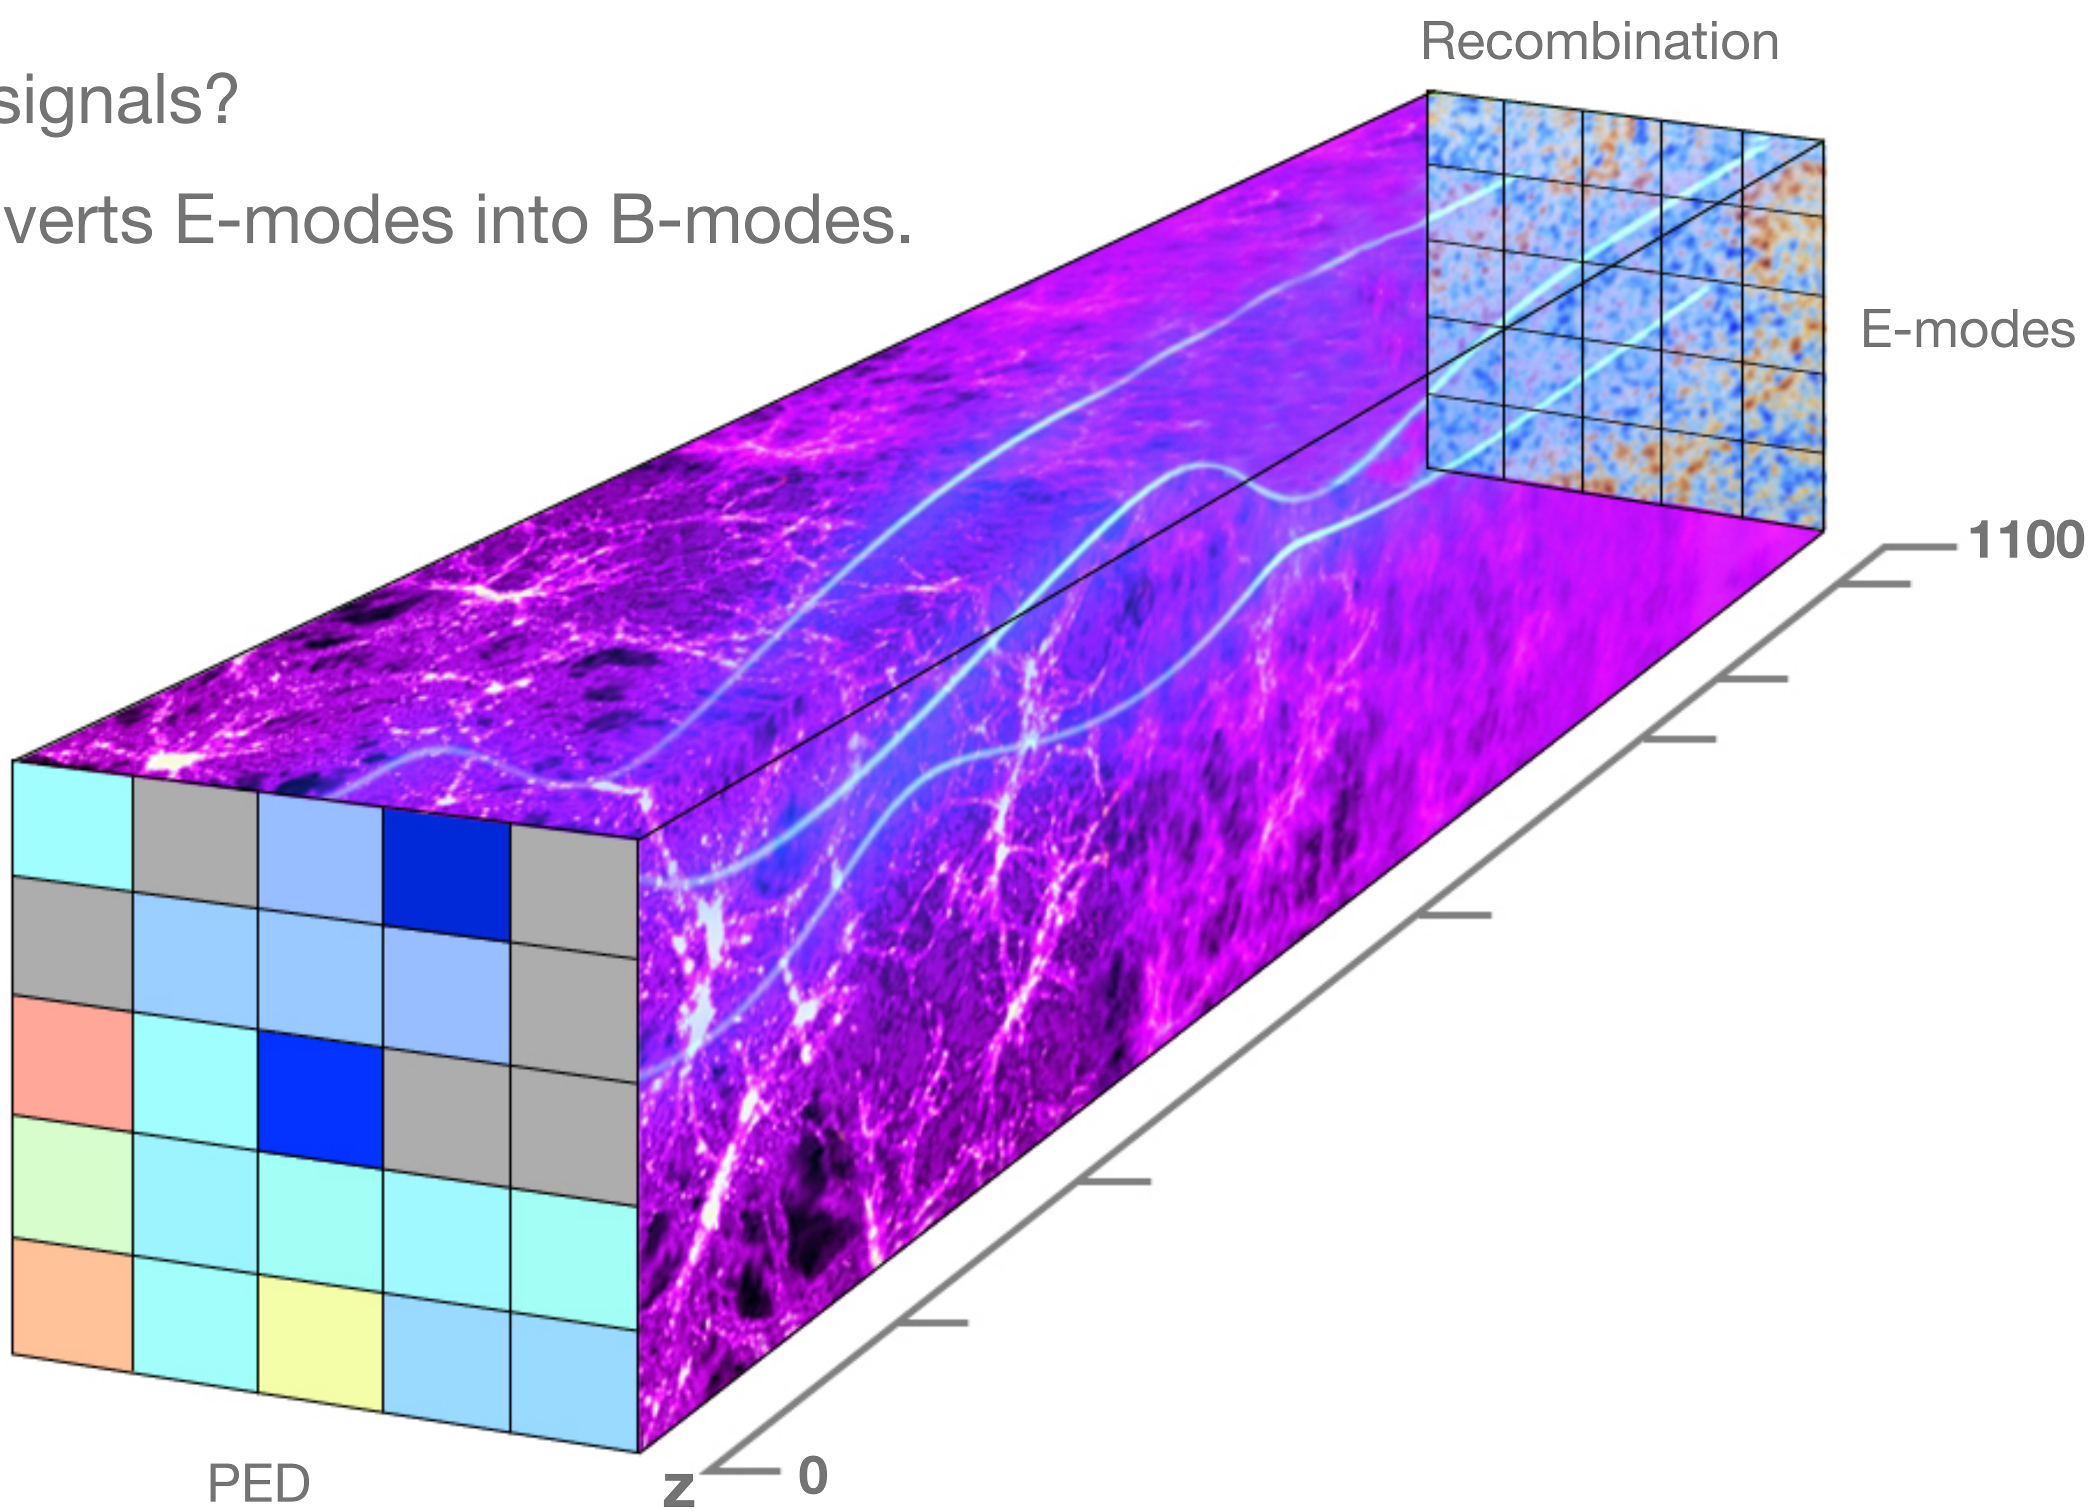

Exploration vs. Exploitation: B-mode Surveys

- Competing signals?
- Lensing converts E-modes into B-modes.
- Detected! (Hanson et al. arXiv:1307.5830)

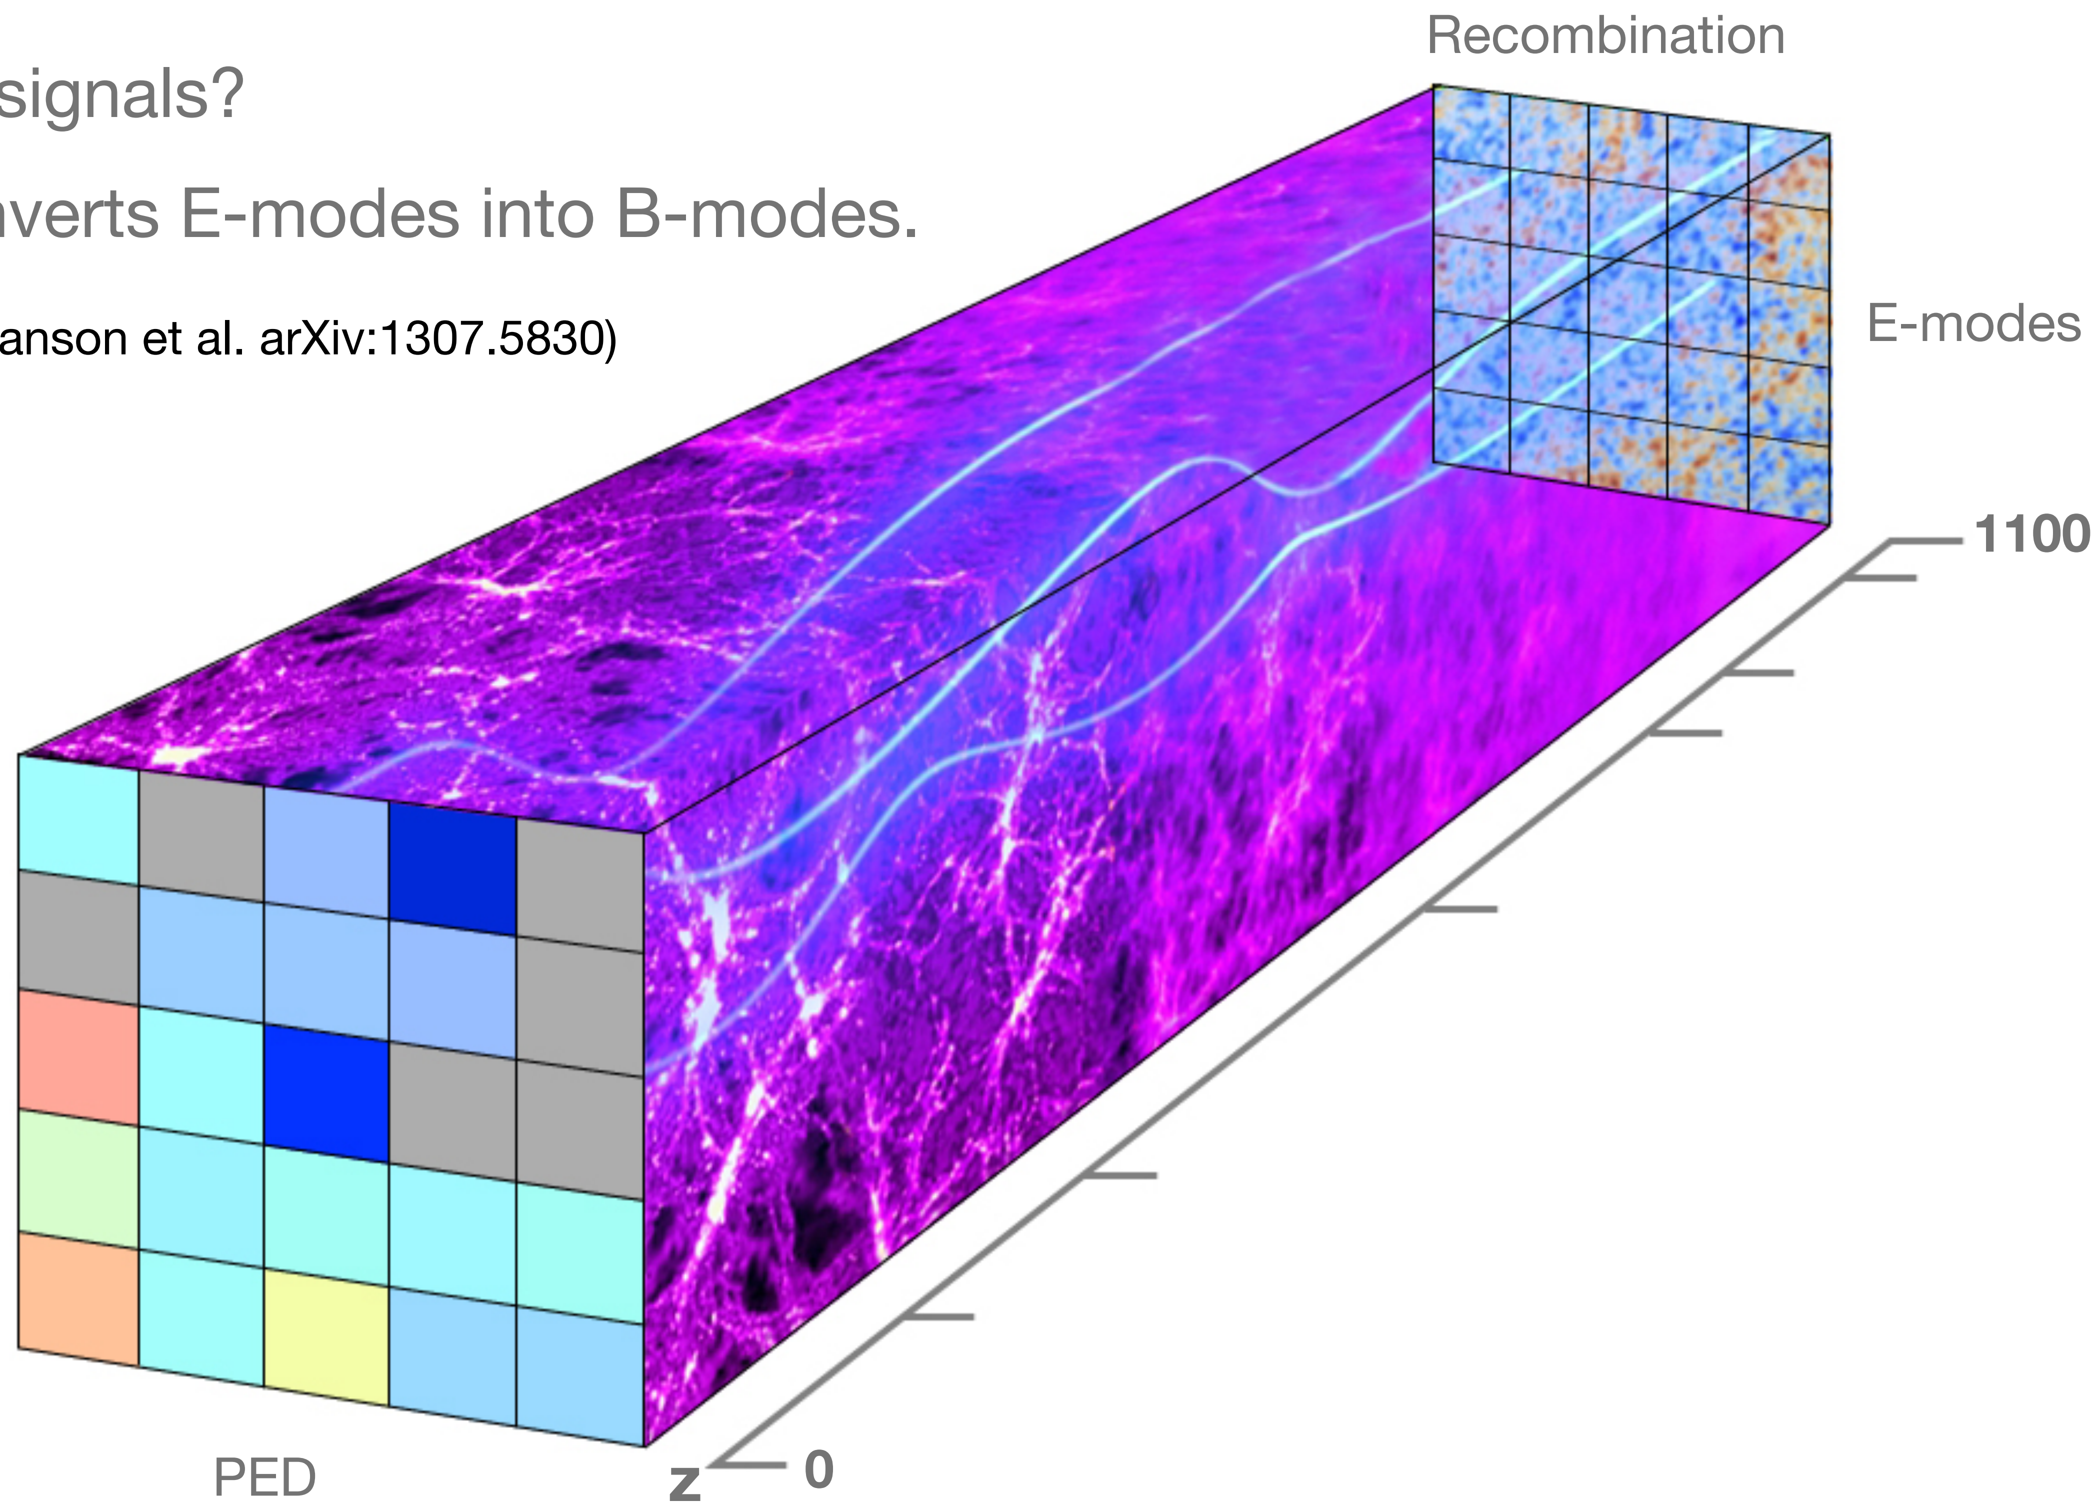

Exploration vs. Exploitation: B-mode Surveys

- Competing signals?
- Lensing converts E-modes into B-modes.
- Detected! (Hanson et al. arXiv:1307.5830)
- De-lensing?

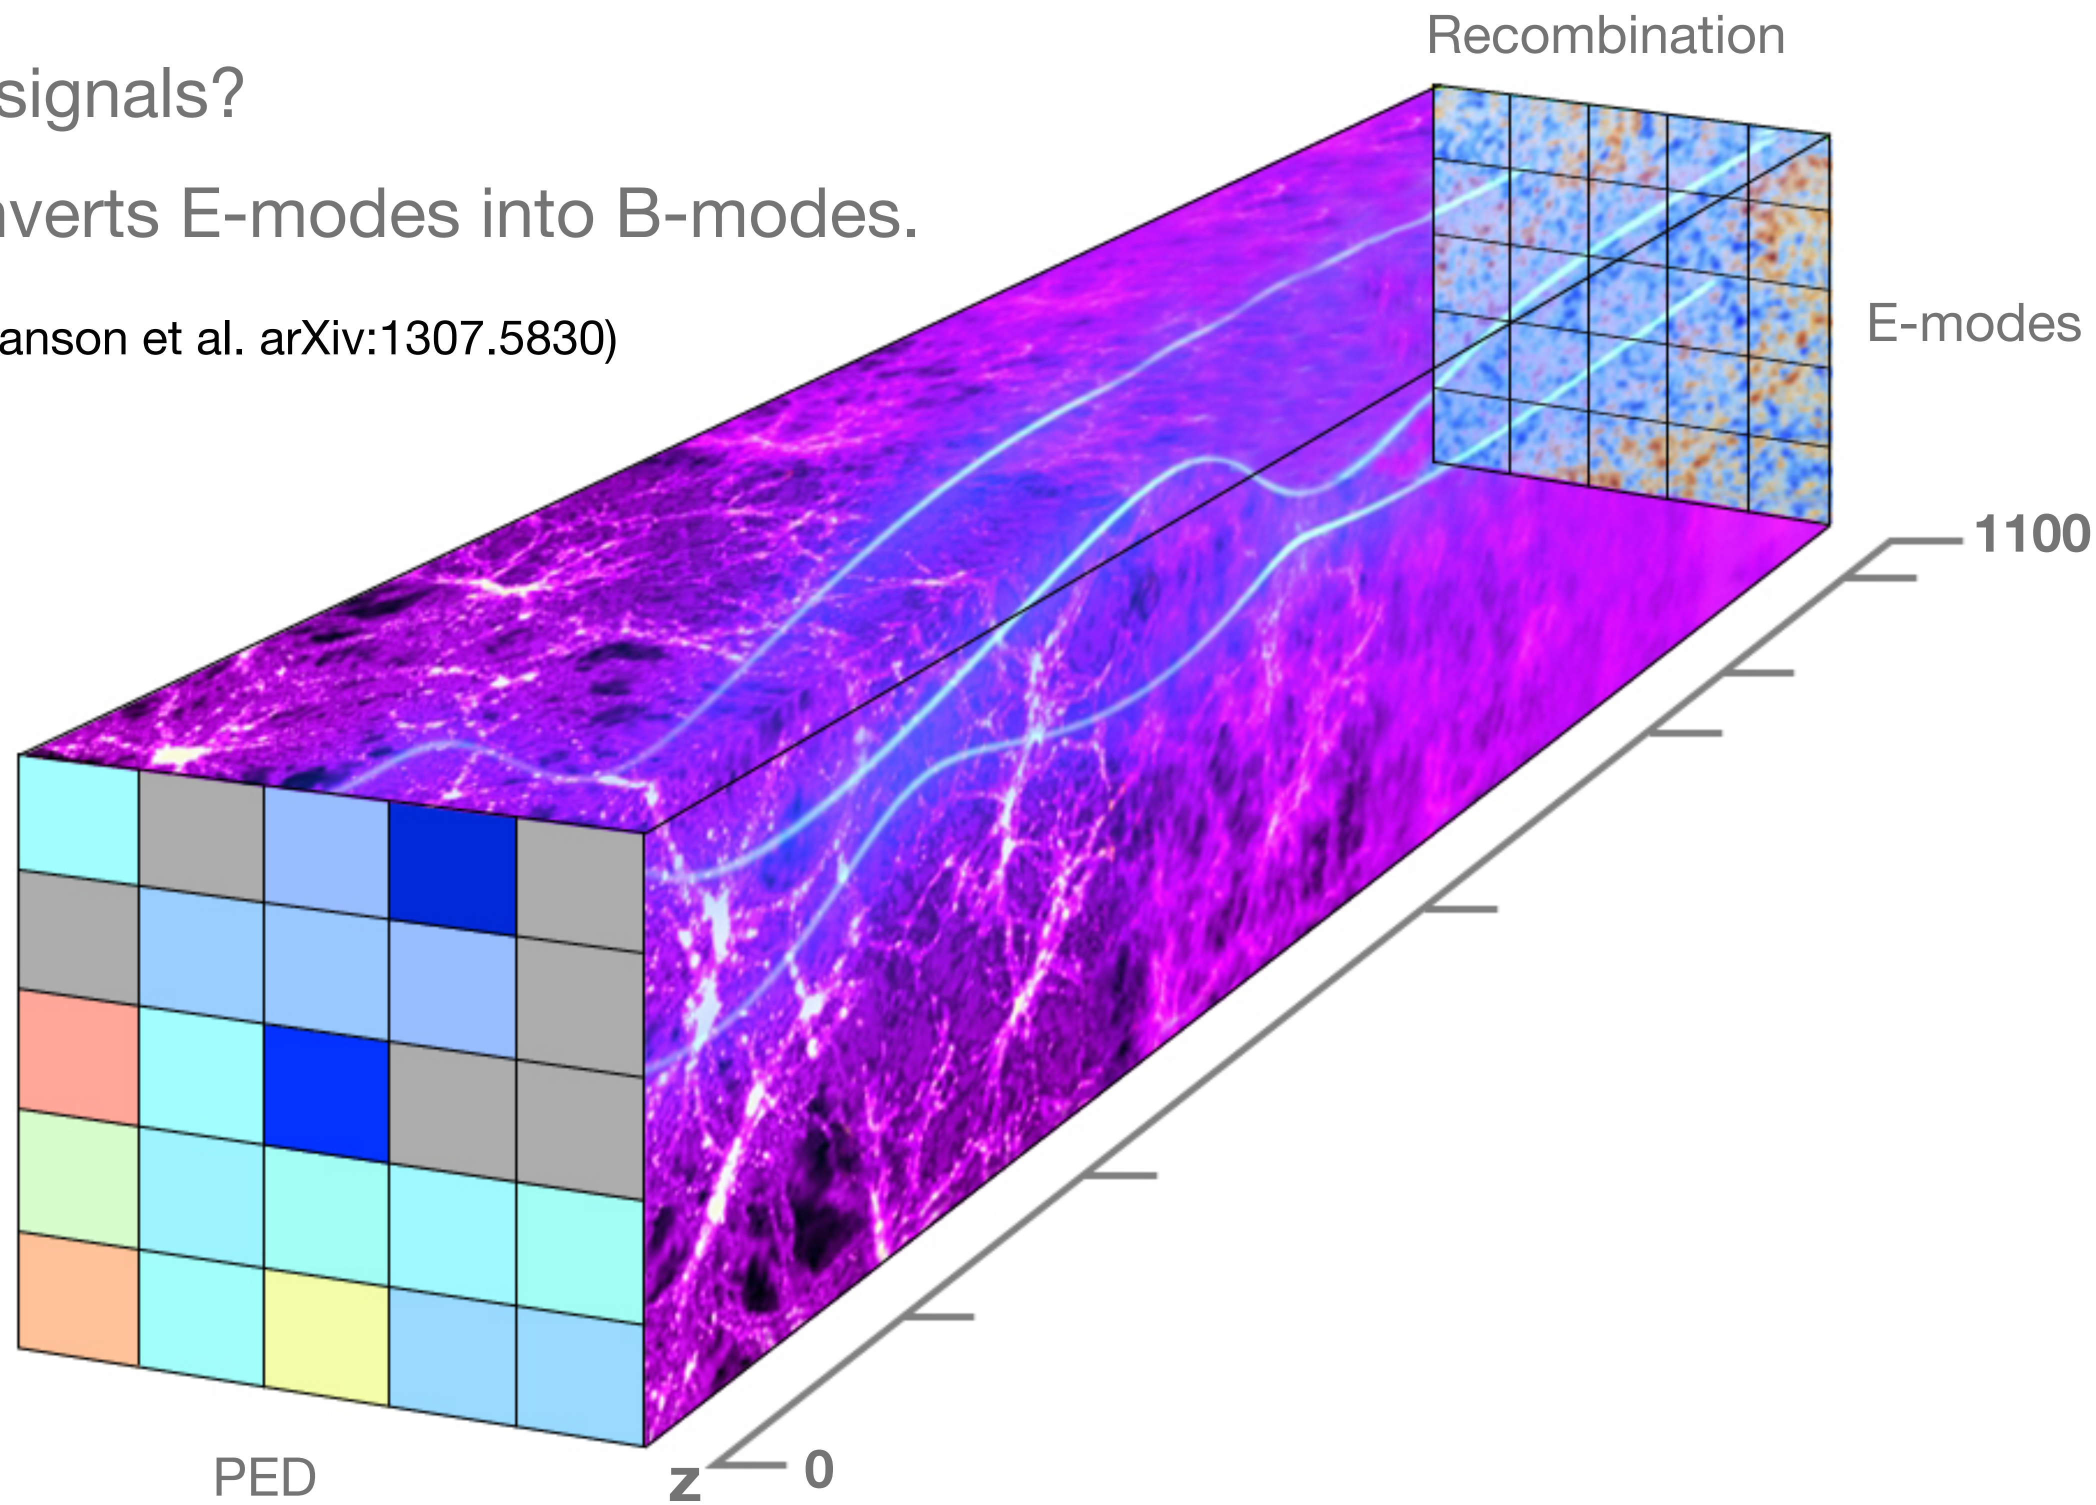

Outline

Outline

- The Multi-Armed-Bandit Problem

Outline

- The Multi-Armed-Bandit Problem

- Heuristic Solution Algorithms

Outline

- The Multi-Armed-Bandit Problem
- Heuristic Solution Algorithms
- Mitigating CMB B-mode Foregrounds

Outline

- The Multi-Armed-Bandit Problem
- Heuristic Solution Algorithms
- Mitigating CMB B-mode Foregrounds
- MAB Strategies Elsewhere

The Multi-Armed-Bandit Problem

The Multi-Armed-Bandit Problem

The Multi-Armed-Bandit Problem

- The goal:

    Facing slots with different odds, maximize winnings.

The Multi-Armed-Bandit Problem

- The goal:

  Facing slots with different odds, maximize winnings.

- With infinite funds, this is easy. You *learn* the odds.

The Multi-Armed-Bandit Problem

- The goal:

  Facing slots with different odds, maximize winnings.

- With infinite funds, this is easy. You *learn* the odds.
- With a finite number of plays, problem is unsolved.

The Multi-Armed-Bandit Problem

- The goal:

  Facing slots with different odds, maximize winnings.

- With infinite funds, this is easy. You *learn* the odds.
- With a finite number of plays, problem is unsolved.
- Heuristics have been developed and compared.

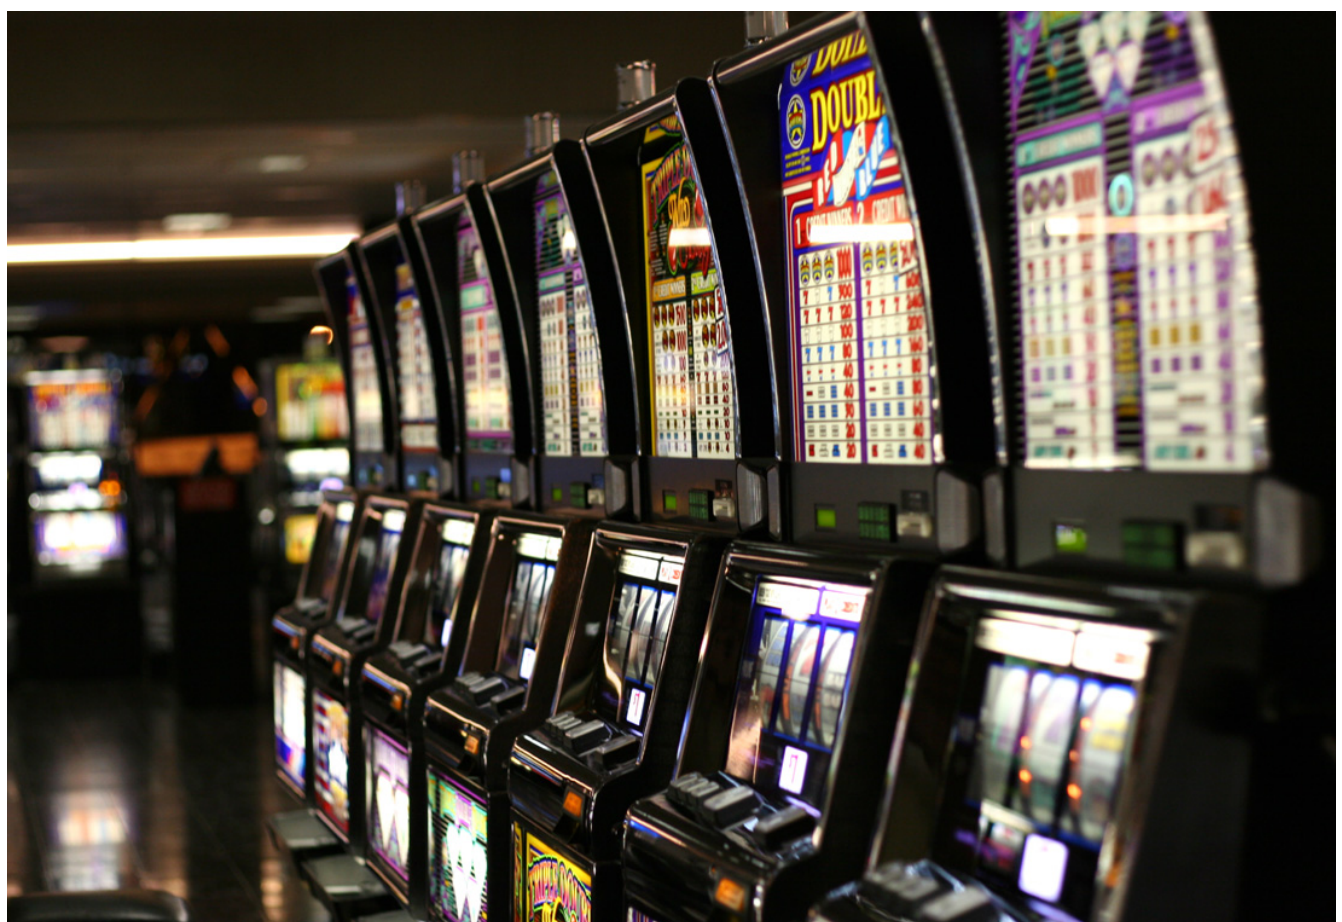

Multi-Armed-Bandit Strategies

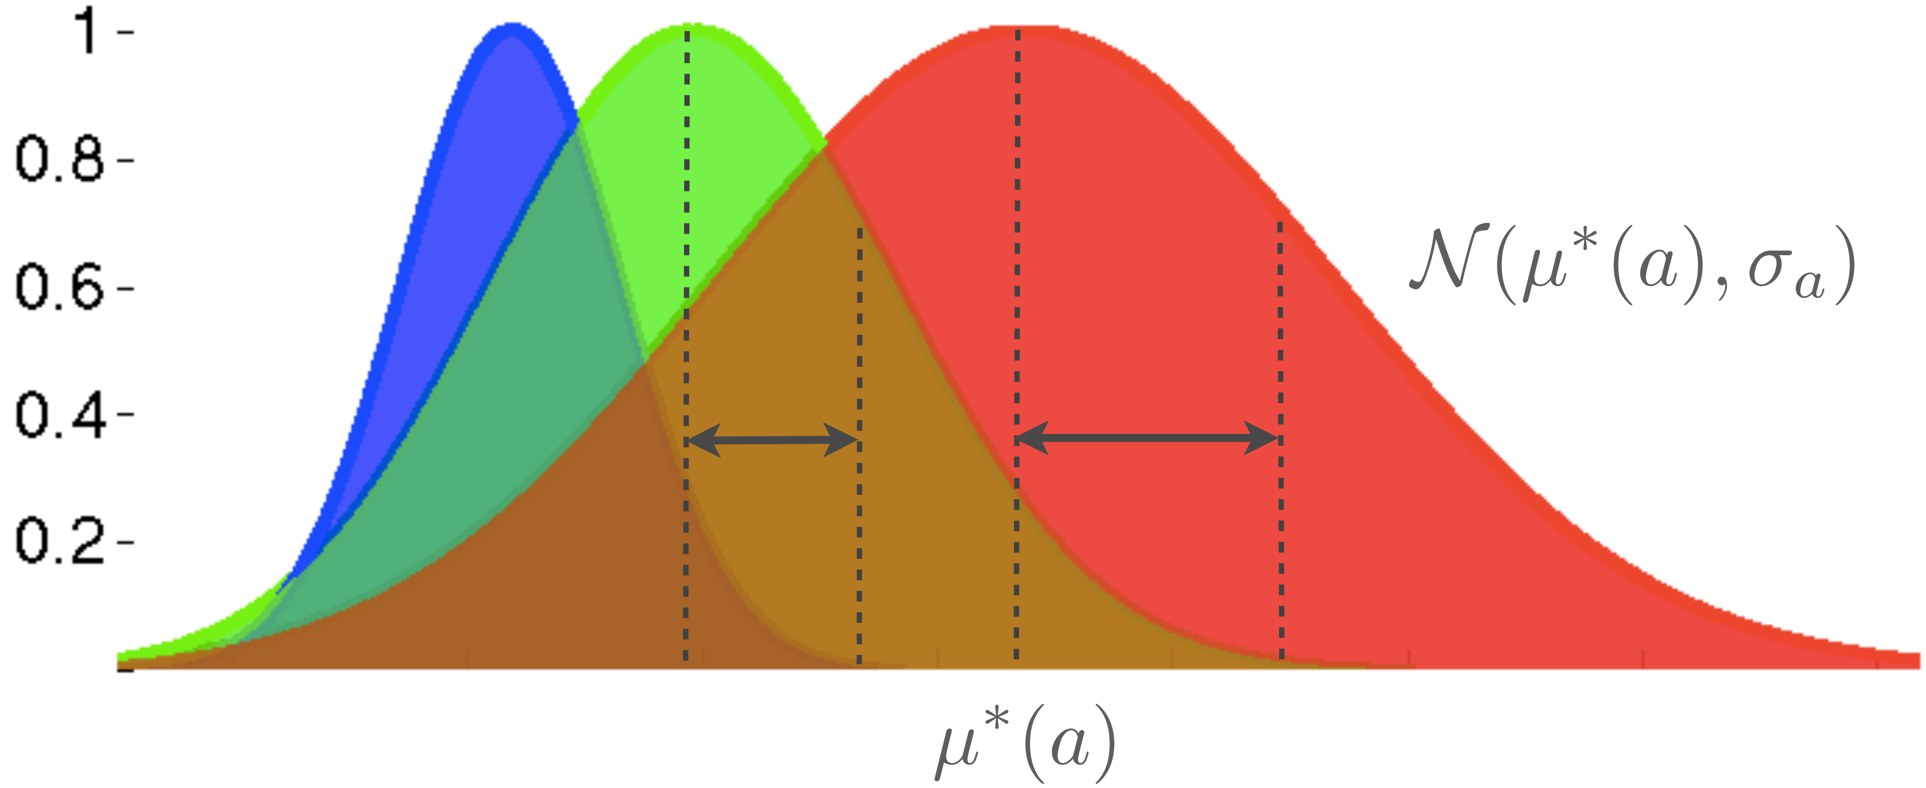

Multi-Armed-Bandit Strategies

An MAB strategy:

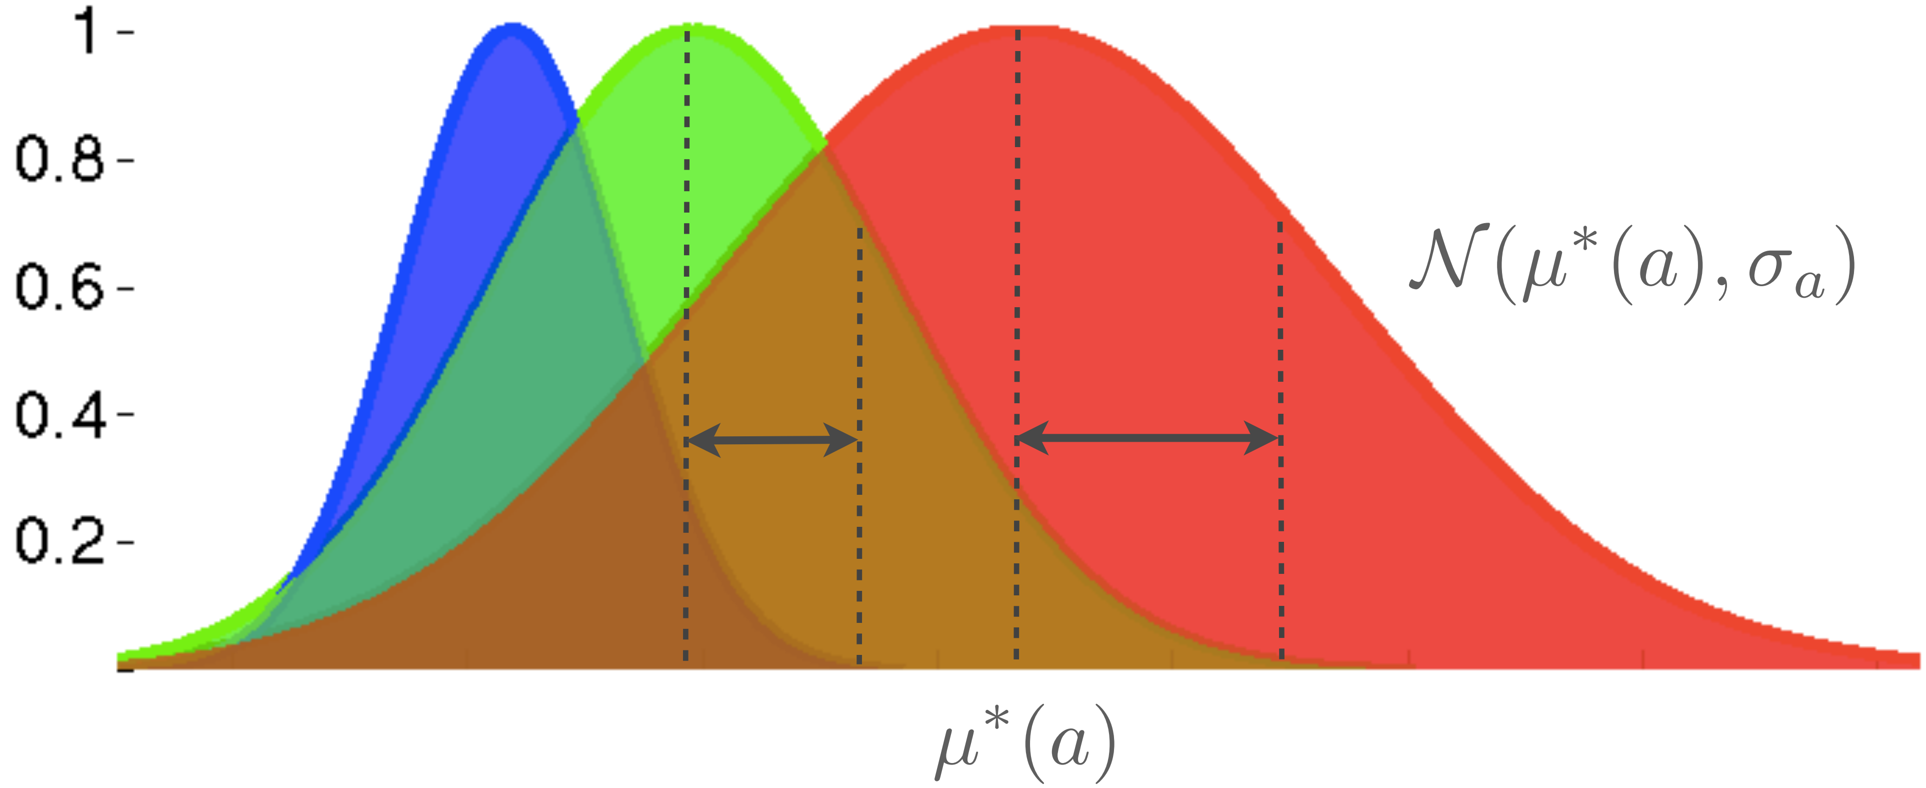

Multi-Armed-Bandit Strategies

An MAB strategy:

- Expected reward estimates are based on a sample-average of previous rewards:

$$\mu_t(a) = \frac{r_1 + r_2 + \cdots + r_{N_t(a)}}{N_t(a)}$$

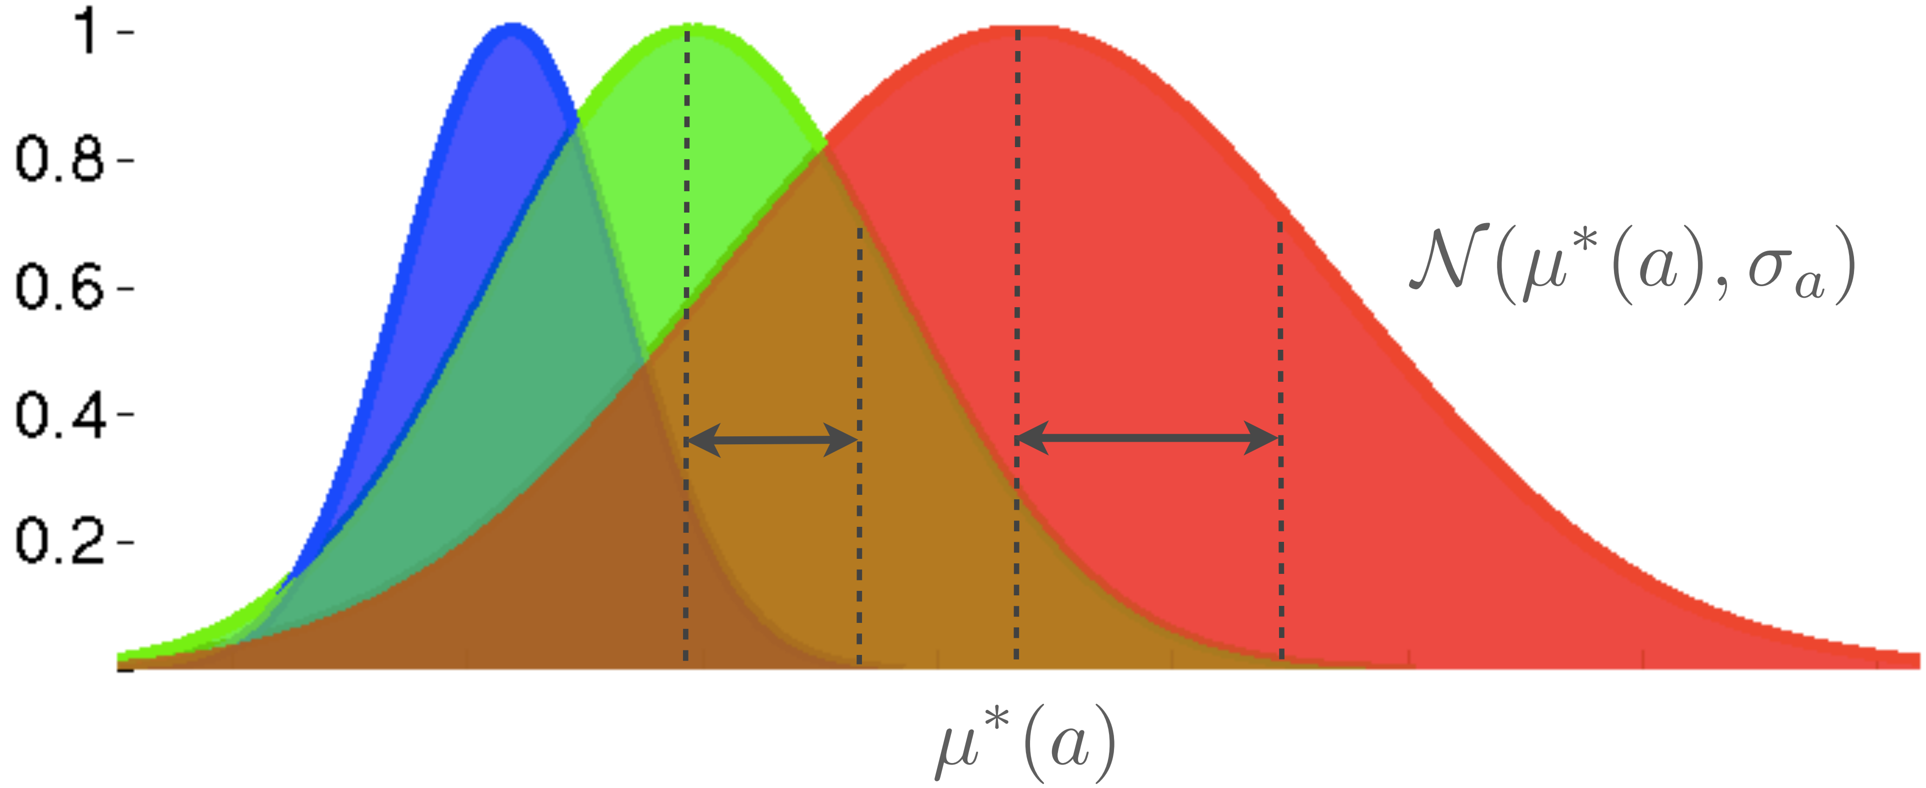

Multi-Armed-Bandit Strategies

An MAB strategy:

- Expected reward estimates are based on a sample-average of previous rewards:

$$\mu_t(a) = \frac{r_1 + r_2 + \cdots + r_{N_t(a)}}{N_t(a)} \qquad \mu_t(a) \xrightarrow[N_t(a) \to \infty]{} \mu^*(a)$$

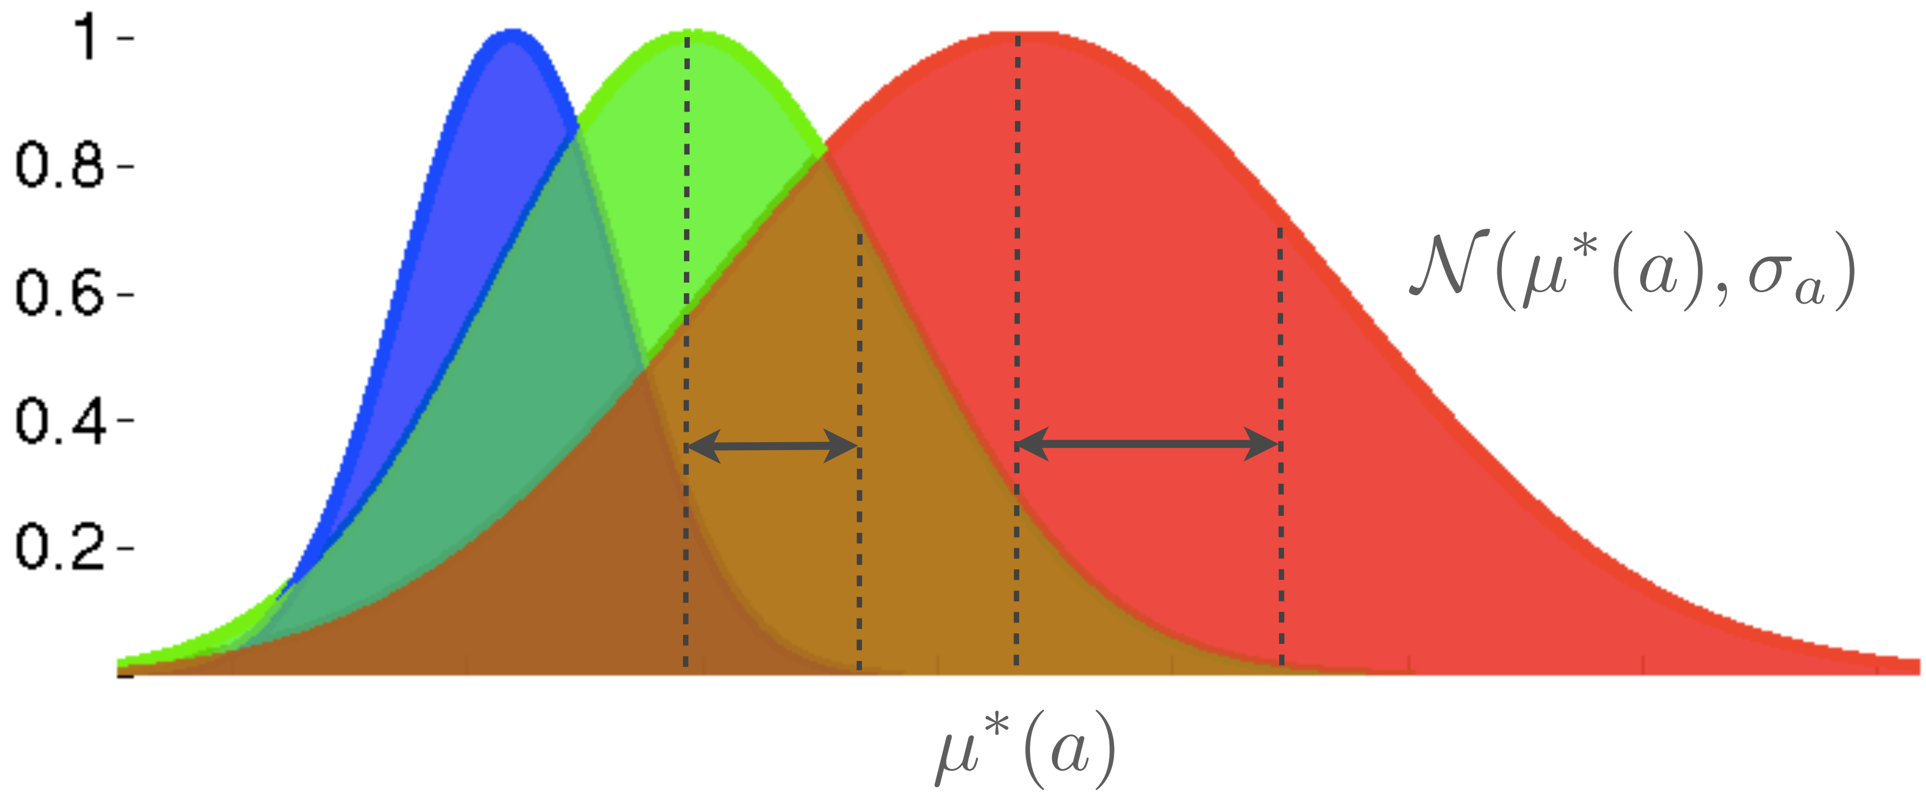

Multi-Armed-Bandit Strategies

An MAB strategy:

- Expected reward estimates are based on a sample-average of previous rewards:

$$\mu_t(a) = \frac{r_1 + r_2 + \cdots + r_{N_t(a)}}{N_t(a)} \qquad \mu_t(a) \xrightarrow[N_t(a) \to \infty]{} \mu^*(a)$$

- These estimates are then used in order to choose the action at $t+1$.

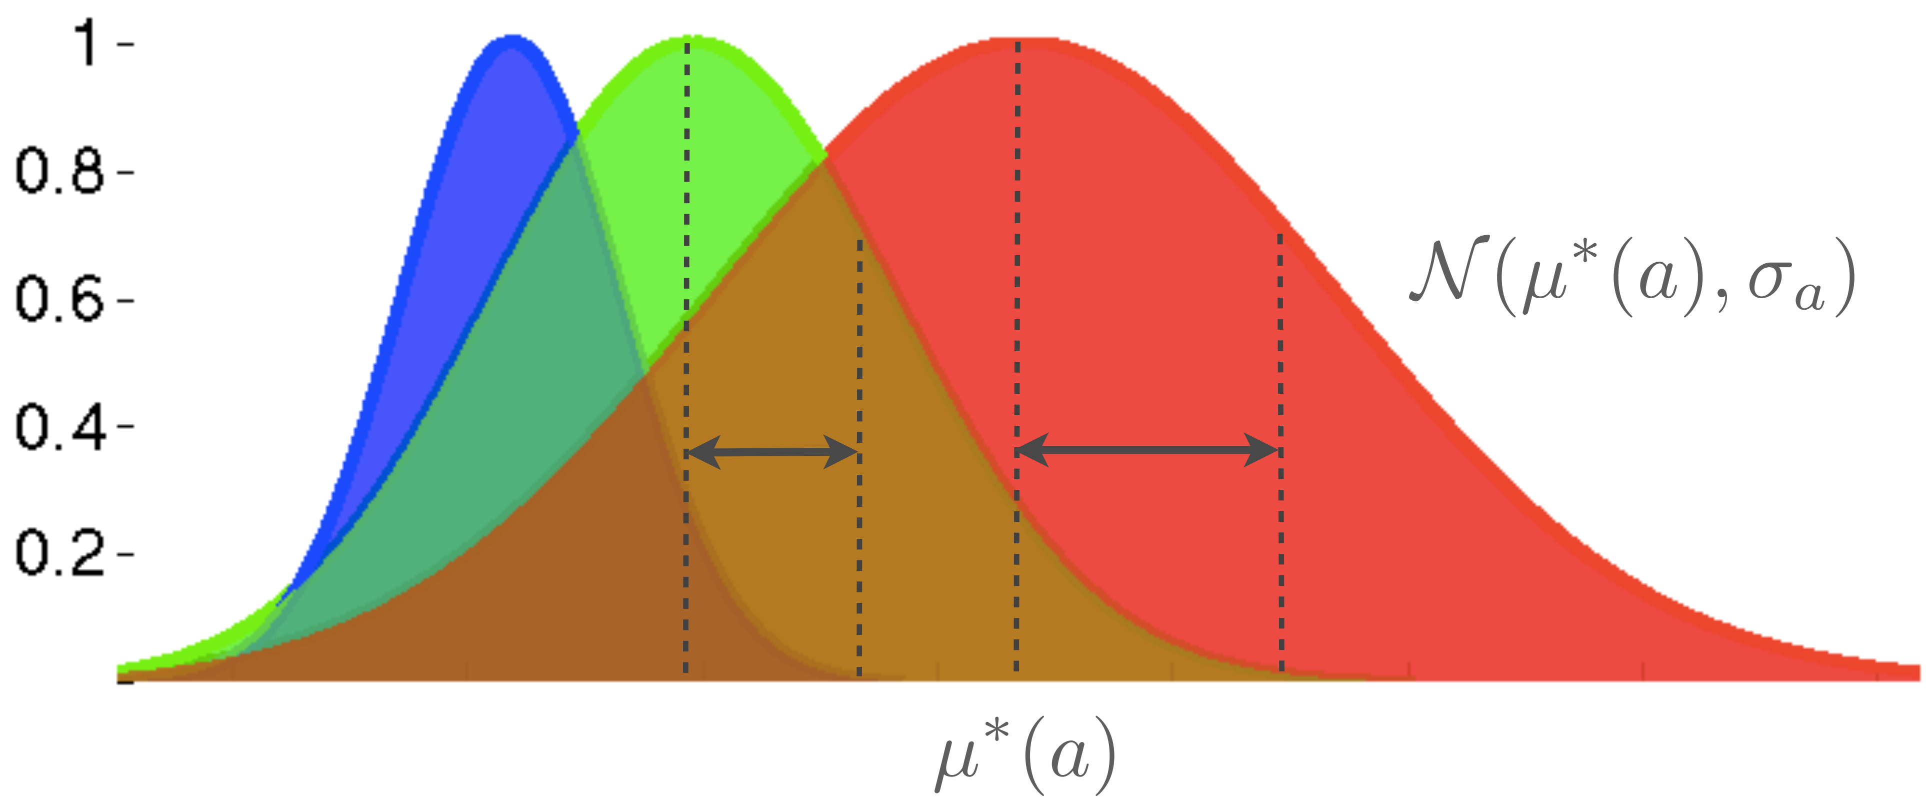

Multi-Armed-Bandit Strategies

An MAB strategy:

- Expected reward estimates are based on a sample-average of previous rewards:

$$\mu_t(a) = \frac{r_1 + r_2 + \cdots + r_{N_t(a)}}{N_t(a)} \qquad \mu_t(a) \xrightarrow[N_t(a) \to \infty]{} \mu^*(a)$$

- These estimates are then used in order to choose the action at $t+1$.
- Strategies can be compared to one another or to an asymptotic lower bound.

(Lai & Robbins, 1985)

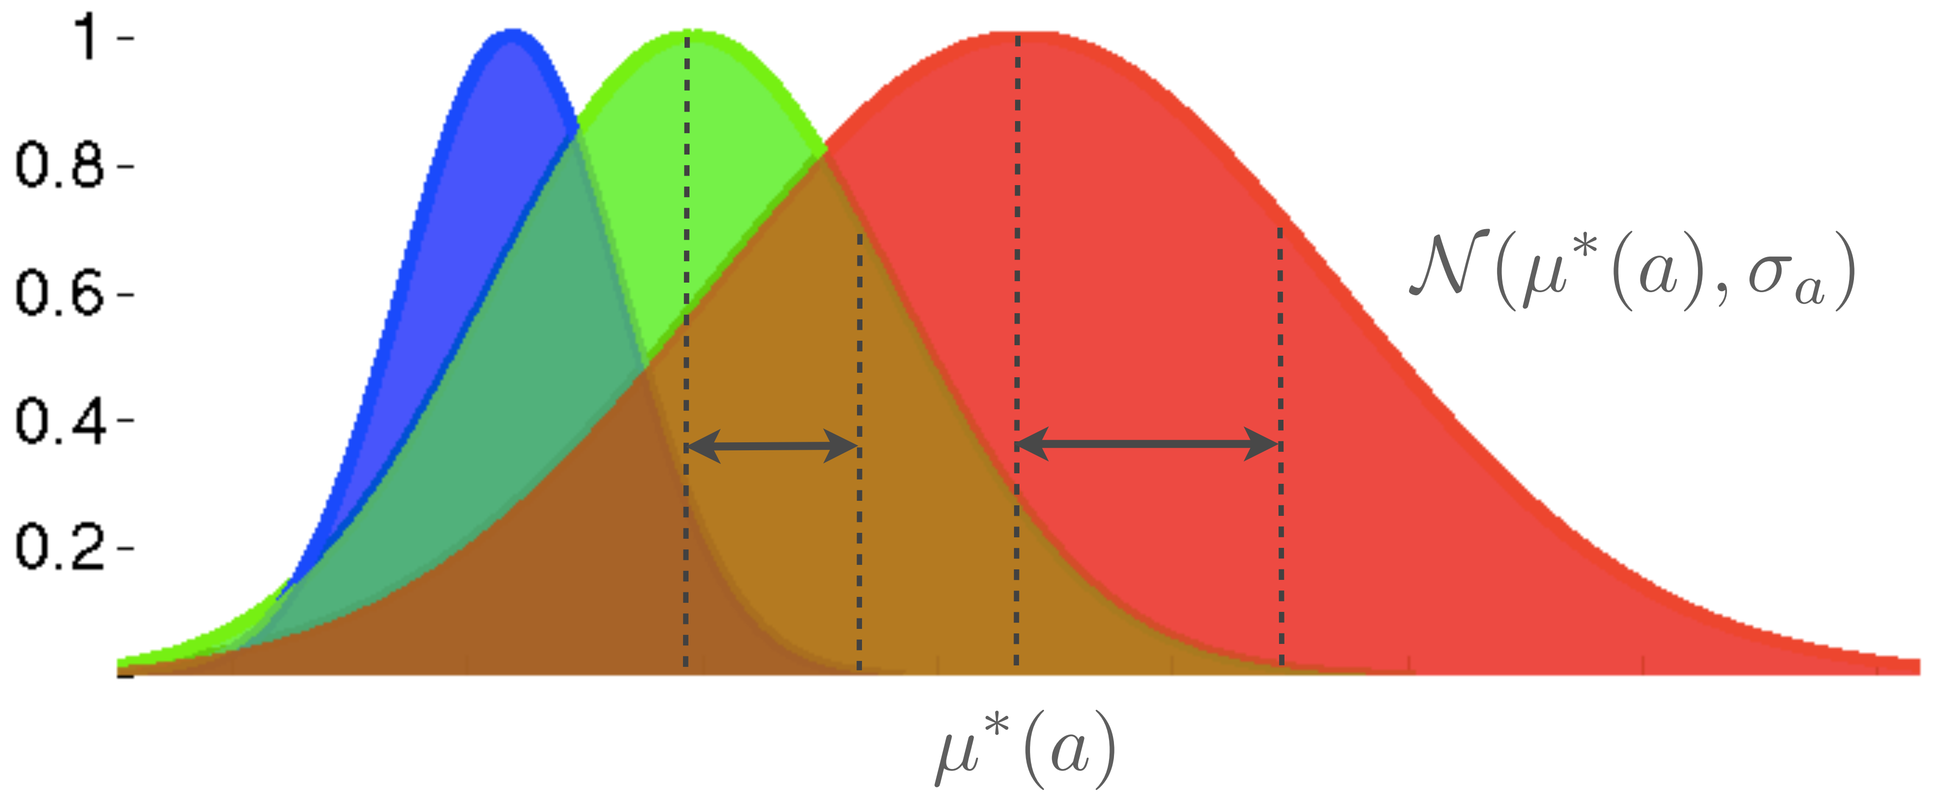

Outline

Outline

- The Multi-Armed-Bandit Problem
- Heuristic Solution Algorithms

Heuristic Solution Algorithms: Naive (rigid)

Heuristic Solution Algorithms: Naive (rigid)

- Uniformly random

  For n_a arms: $p_t(a) = 1/n_a$

Heuristic Solution Algorithms: Naive (rigid)

- Uniformly random

For n_a arms: $p_t(a) = 1/n_a$

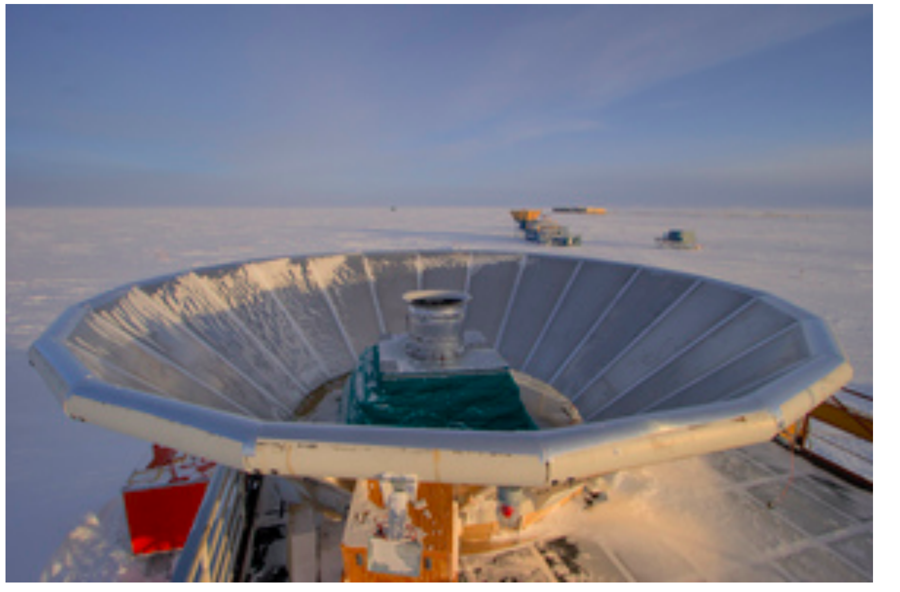

BICEP (larger $f_{\rm sky}$)

Heuristic Solution Algorithms: Naive (rigid)

- Uniformly random

For n_a arms: $p_t(a) = 1/n_a$

Never explores --> has linear total regret.

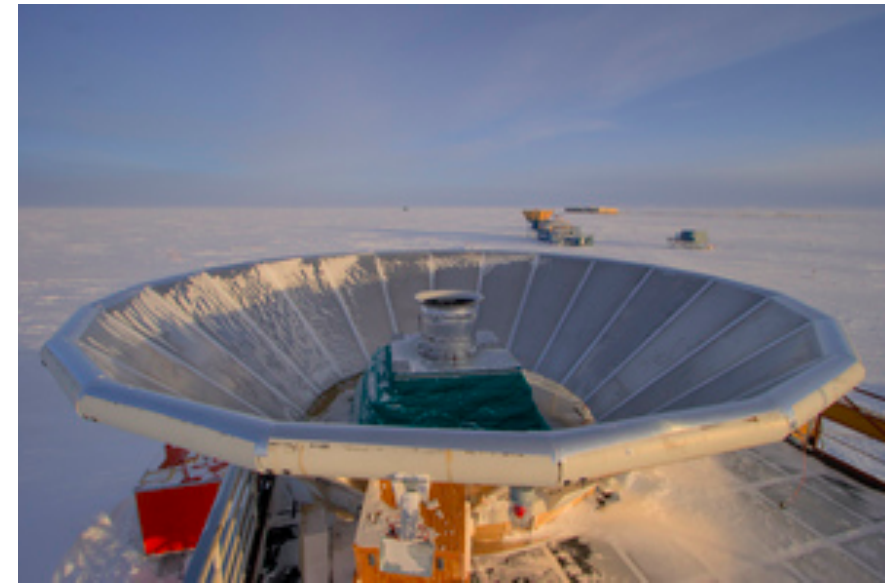

BICEP (larger f_{sky})

Heuristic Solution Algorithms: Naive (rigid)

- Uniformly random

  For n_a arms: $p_t(a) = 1/n_a$

  Never explores --> has linear total regret.

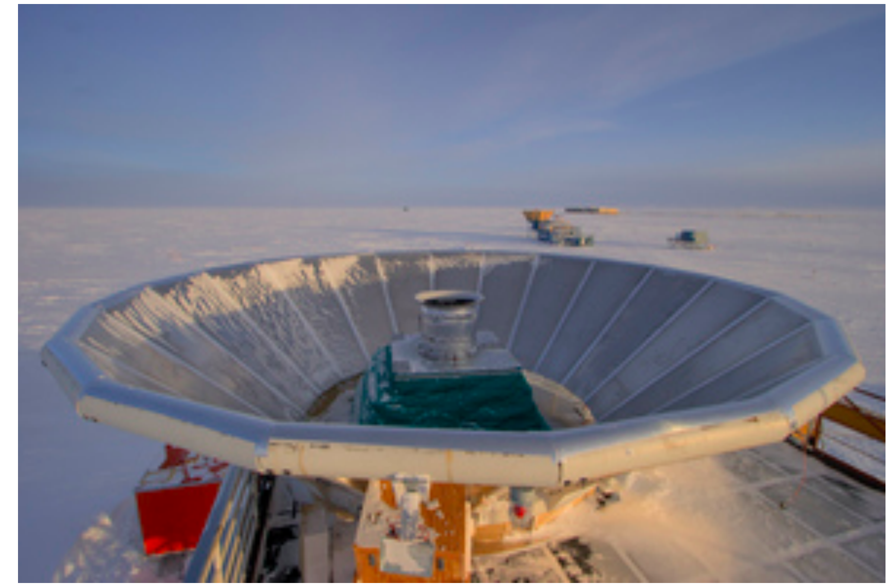

BICEP (larger $f_{\rm sky}$)

- Greedy

$$p_t(a) = 1 \quad \text{if a} = \underset{a'}{\text{argmax}}\ \mu_t(a')$$

Heuristic Solution Algorithms: Naive (rigid)

- Uniformly random

  For n_a arms: $p_t(a) = 1/n_a$

  Never explores --> has linear total regret.

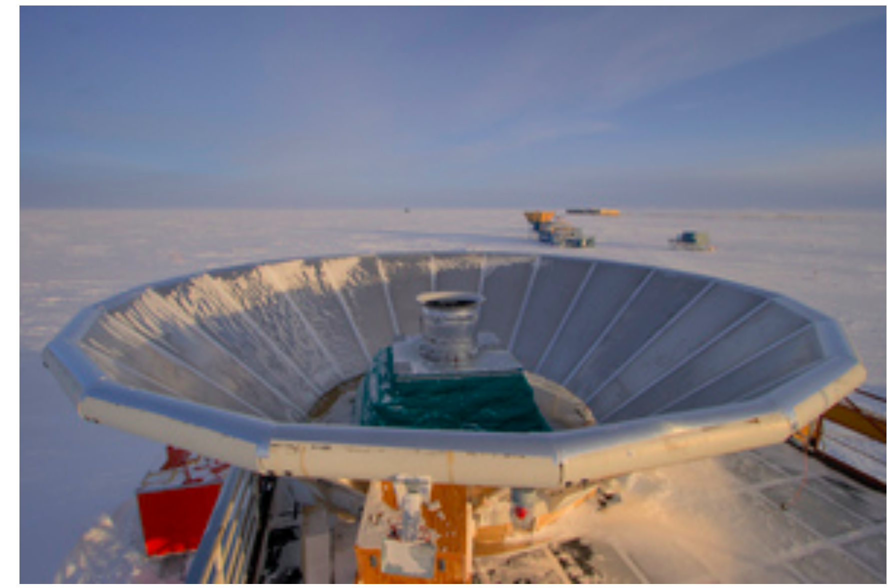

BICEP (larger $f_{\rm sky}$)

- Greedy

$$p_t(a) = 1 \quad \text{if a} = \underset{a'}{\text{argmax}}\ \mu_t(a')$$

POLARBEAR (small $f_{\rm sky}$)

Heuristic Solution Algorithms: Naive (rigid)

- Uniformly random

For n_a arms: $p_t(a) = 1/n_a$

Never explores --> has linear total regret.

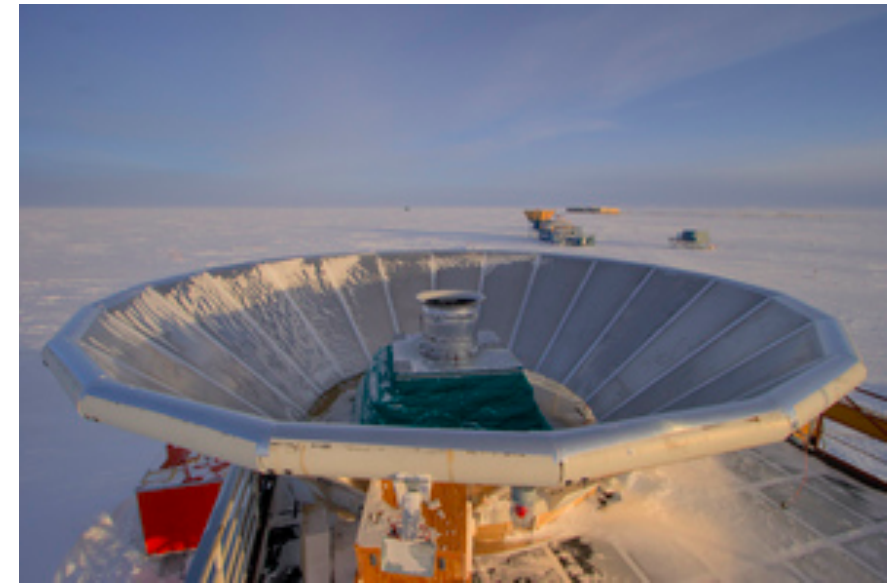

BICEP (larger $f_{\rm sky}$)

- Greedy

$$p_t(a) = 1 \quad \text{if a} = \underset{a'}{\operatorname{argmax}}\, \mu_t(a')$$

Never explores --> has linear total regret.

POLARBEAR (small $f_{\rm sky}$)

Heuristic Solution Algorithms: Forced Exploration

Heuristic Solution Algorithms: Forced Exploration

- ϵ - greedy

$$\begin{cases} p_t = 1 - \epsilon & \text{Greedy Arm} \\ p_t = \quad \epsilon & \text{Uniformly Random} \end{cases}$$

Heuristic Solution Algorithms: Forced Exploration

- ϵ - greedy $$\begin{cases} p_t = 1 - \epsilon & \text{Greedy Arm} \\ p_t = \quad \epsilon & \text{Uniformly Random} \end{cases}$$

  Always explores --> has linear total regret.

Heuristic Solution Algorithms: Forced Exploration

- ϵ - greedy
$$\begin{cases} p_t = 1 - \epsilon & \text{Greedy Arm} \\ p_t = \quad \epsilon & \text{Uniformly Random} \end{cases}$$
  Always explores --> has linear total regret.

- ϵ_t - greedy with a decaying strategy: $\epsilon_t \propto 1/t^{\alpha}$

Heuristic Solution Algorithms: Forced Exploration

- ϵ - greedy $\begin{cases} p_t = 1 - \epsilon & \text{Greedy Arm} \\ p_t = \quad \epsilon & \text{Uniformly Random} \end{cases}$

  Always explores --> has linear total regret.

- ϵ_t - greedy with a decaying strategy: $\epsilon_t \propto 1/t^\alpha$

  --> can achieve logarithmic regret!

Heuristic Solution Algorithms: Forced Exploration

- ϵ - greedy
$$\begin{cases} p_t = 1 - \epsilon & \text{Greedy Arm} \\ p_t = \quad \epsilon & \text{Uniformly Random} \end{cases}$$

  Always explores --> has linear total regret.

- ϵ_t - greedy with a decaying strategy: $\epsilon_t \propto 1/t^\alpha$

  --> can achieve logarithmic regret!

- Probability matching

  (Boltzmann)

$$p_t(a) = \frac{e^{\mu_t(a)/\tau}}{\sum_{a'} e^{\mu_t(a')/\tau}}$$

(W.R. Thompson, 1933)

Heuristic Solution Algorithms: Forced Exploration

- ϵ - greedy

$$\begin{cases} p_t = 1 - \epsilon & \text{Greedy Arm} \\ p_t = \quad \epsilon & \text{Uniformly Random} \end{cases}$$

  Always explores --> has linear total regret.

- ϵ_t - greedy with a decaying strategy: $\epsilon_t \propto 1/t^\alpha$

  --> can achieve logarithmic regret!

- Probability matching

  (Boltzmann)

$$p_t(a) = \frac{e^{\mu_t(a)/\tau}}{\sum_{a'} e^{\mu_t(a')/\tau}}$$

(W.R. Thompson, 1933)

  Limits: $\tau \to 0$ (greedy) $\tau \to \infty$ (uniformly random)

Heuristic Solution Algorithms: Forced Exploration

- ϵ - greedy

$$\begin{cases} p_t = 1 - \epsilon & \text{Greedy Arm} \\ p_t = \quad \epsilon & \text{Uniformly Random} \end{cases}$$

  Always explores --> has linear total regret.

- ϵ_t - greedy with a decaying strategy: $\epsilon_t \propto 1/t^\alpha$

  --> can achieve logarithmic regret!

- Probability matching (Boltzmann)

$$p_t(a) = \frac{e^{\mu_t(a)/\tau}}{\sum_{a'} e^{\mu_t(a')/\tau}}$$

(W.R. Thompson, 1933)

  Limits: $\tau \to 0$ (greedy) $\tau \to \infty$ (uniformly random)

- Upper Confidence Bound (UCB)

(Auer, Cesa-Bianchi, and Fischer, 2002)

Heuristic Solution Algorithms: Forced Exploration

- ϵ - greedy $\begin{cases} p_t = 1 - \epsilon & \text{Greedy Arm} \\ p_t = \quad \epsilon & \text{Uniformly Random} \end{cases}$

  Always explores --> has linear total regret.

- ϵ_t - greedy with a decaying strategy: $\epsilon_t \propto 1/t^\alpha$

  --> can achieve logarithmic regret!

- Probability matching (Boltzmann) $$p_t(a) = \frac{e^{\mu_t(a)/\tau}}{\sum_{a'} e^{\mu_t(a')/\tau}}$$ (W.R. Thompson, 1933)

  Limits: $\tau \to 0$ (greedy) $\tau \to \infty$ (uniformly random)

- Upper Confidence Bound (UCB) (Auer, Cesa-Bianchi, and Fischer, 2002)

  define: $\mu^*(a) \leq \mu_t(a) + U_t(a)$ and choose: $a_t = \underset{a}{\operatorname{argmax}} \{\mu_t(a) + U_t(a)\}$

Heuristic Solution Algorithms: Forced Exploration

- ϵ - greedy $\begin{cases} p_t = 1 - \epsilon & \text{Greedy Arm} \\ p_t = \quad \epsilon & \text{Uniformly Random} \end{cases}$

  Always explores --> has linear total regret.

- ϵ_t - greedy with a decaying strategy: $\epsilon_t \propto 1/t^\alpha$

  --> can achieve logarithmic regret!

- Probability matching (Boltzmann)

$$p_t(a) = \frac{e^{\mu_t(a)/\tau}}{\sum_{a'} e^{\mu_t(a')/\tau}}$$

(W.R. Thompson, 1933)

  Limits: $\tau \to 0$ (greedy) $\tau \to \infty$ (uniformly random)

- Upper Confidence Bound (UCB)

(Auer, Cesa-Bianchi, and Fischer, 2002)

  define: $\mu^*(a) \leq \mu_t(a) + U_t(a)$ and choose: $a_t = \underset{a}{\operatorname{argmax}} \{\mu_t(a) + U_t(a)\}$

  for normal distributions: $U_t(a) \propto \sigma_a/\sqrt{N_t(a)}$

Heuristic Solution Algorithms: Simulation

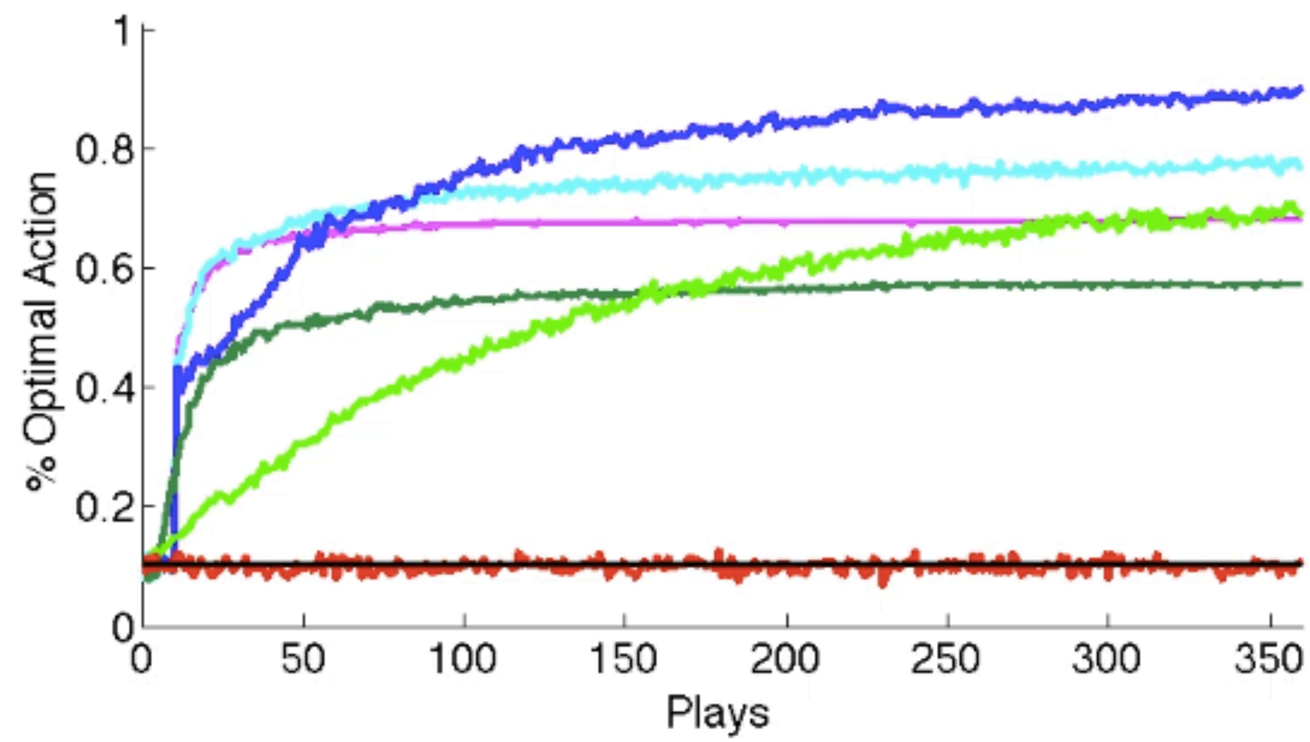

Heuristic Solution Algorithms: Simulation

- MAB Simulation: $\langle \ldots \rangle_{1,000}$

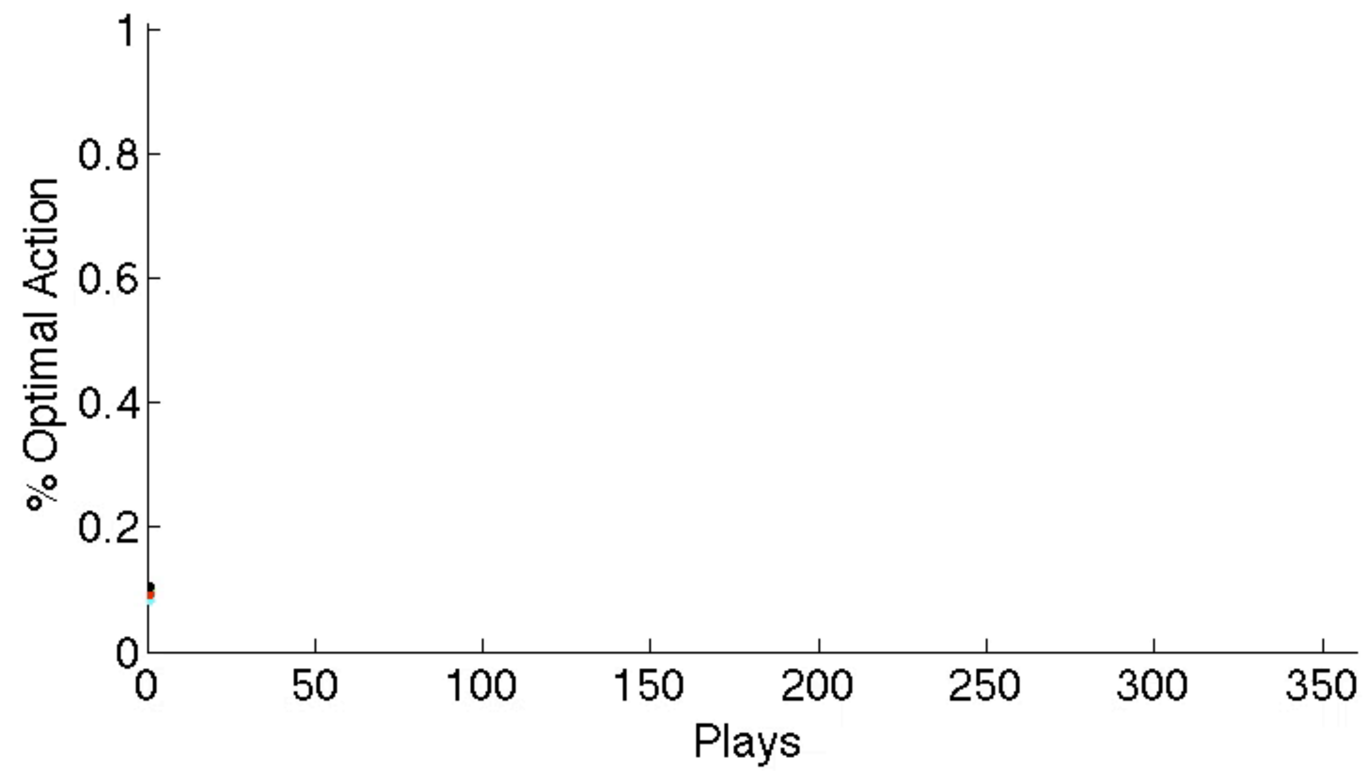

Heuristic Solution Algorithms: Simulation

- MAB Simulation: $\langle \ldots \rangle_{1,000}$

$n_p = 10;\ \forall a\colon\ \mathcal{N}(\mu^*(a), \sigma_a)$

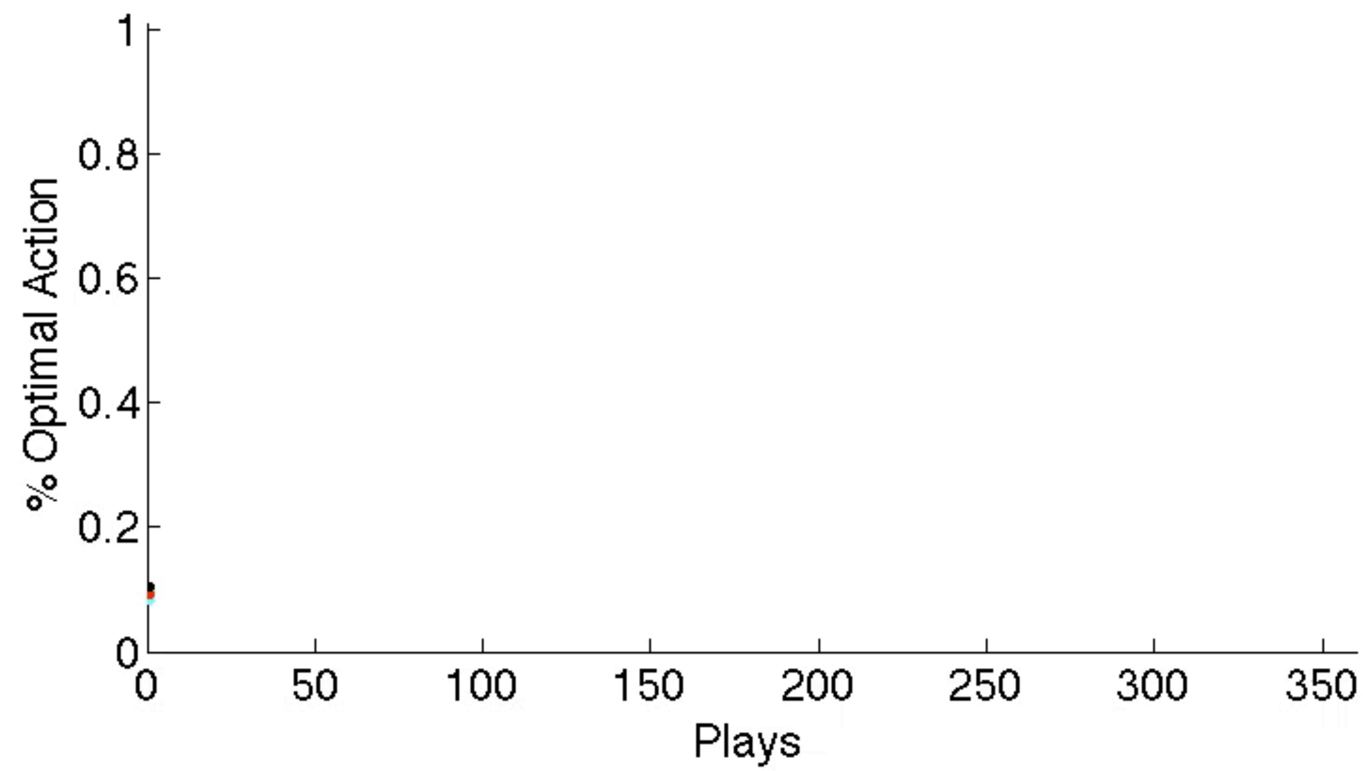

Heuristic Solution Algorithms: Simulation

- MAB Simulation: $\langle \ldots \rangle_{1,000}$

$n_p = 10;\ \forall a\colon\ \mathcal{N}(\mu^*(a), \sigma_a)$

- greedy
- greedy $\epsilon = 1$
- greedy $\epsilon = 0.1$
- greedy decay-ϵ
- UCB
- Boltzmann τ
- opt + greedy

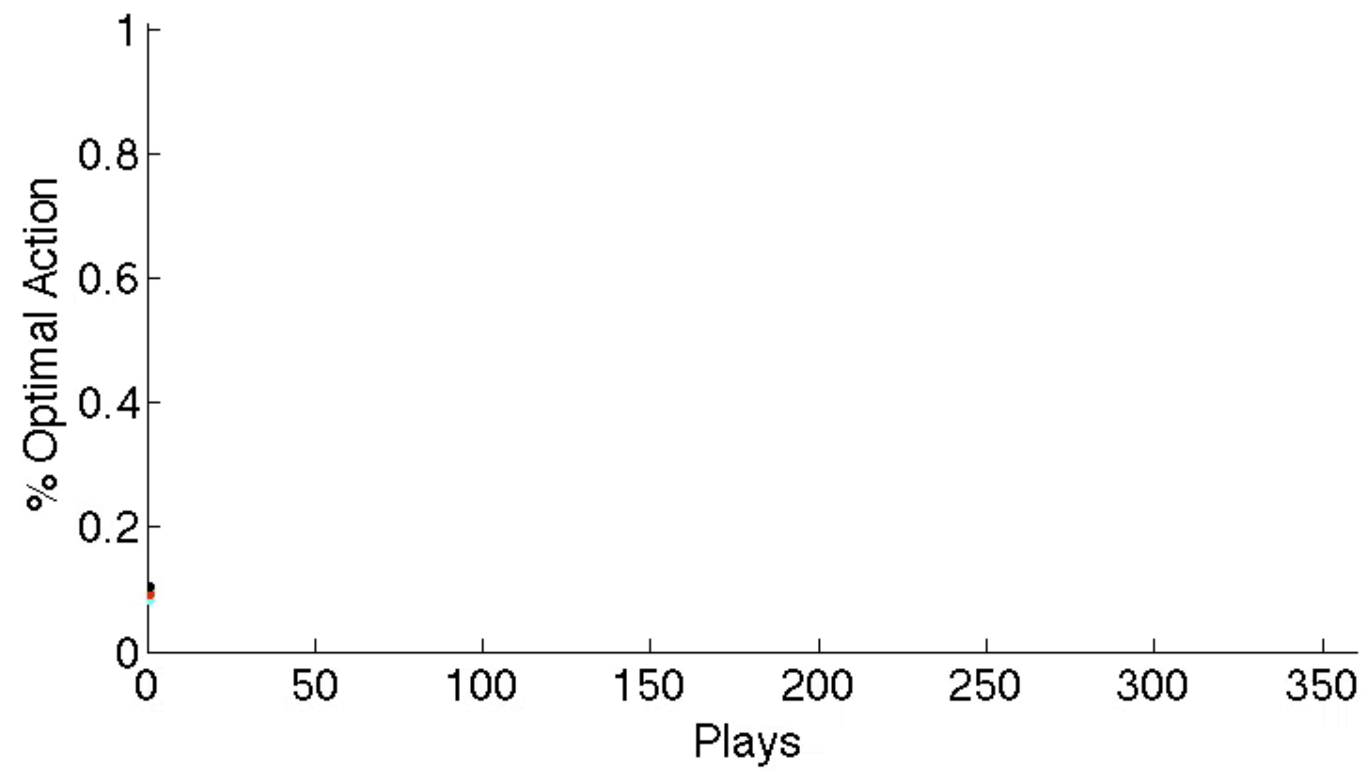

Heuristic Solution Algorithms: Simulation

- MAB Simulation: $\langle \ldots \rangle_{1,000}$

$n_p = 10;\ \forall a:\ \mathcal{N}(\mu^*(a), \sigma_a)$

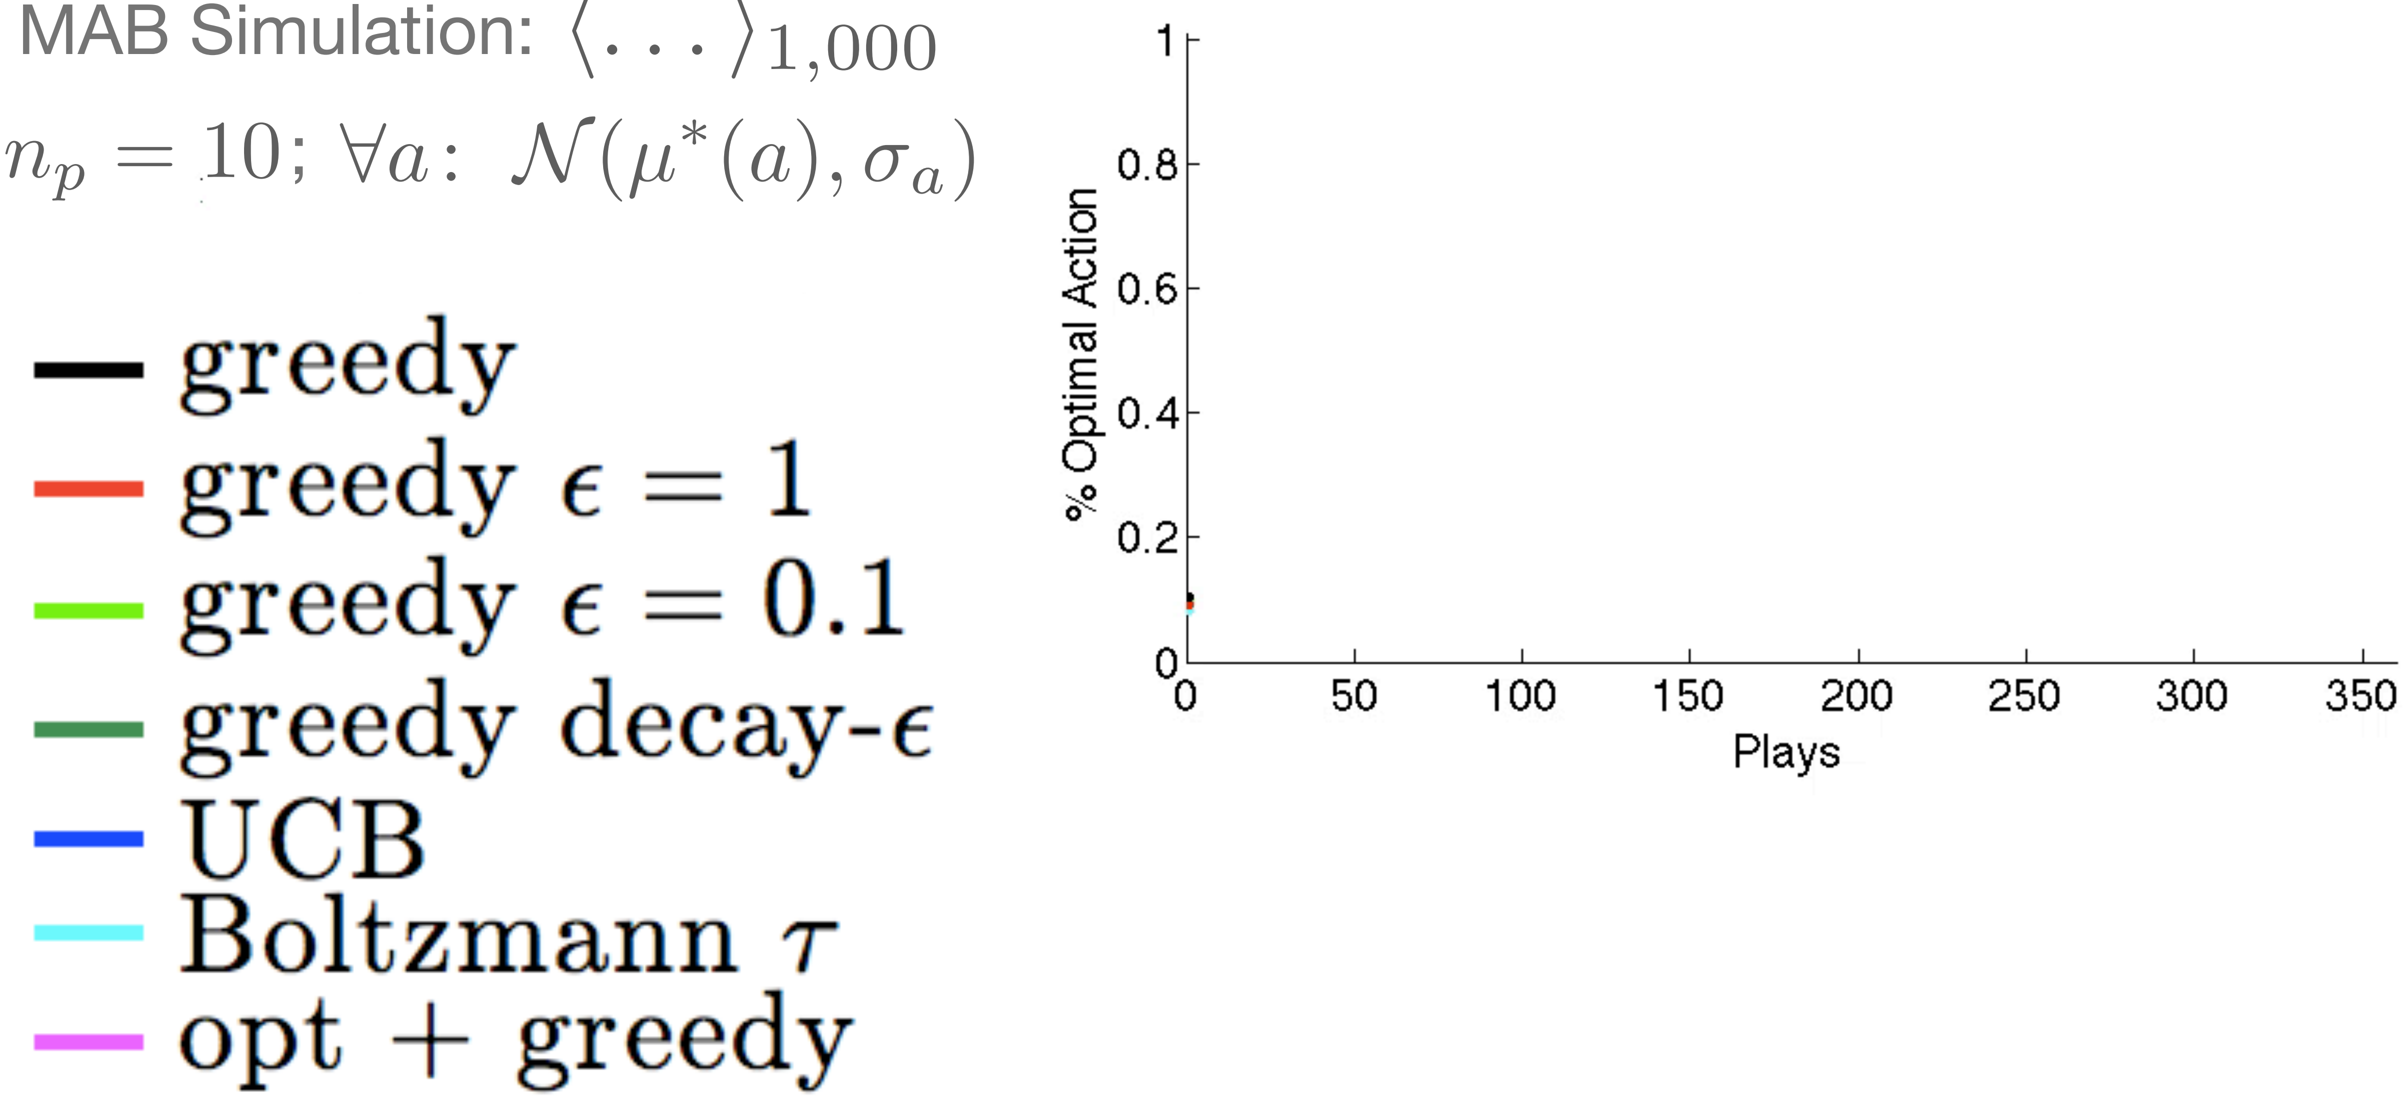

Outline

Outline

- The Multi-Armed-Bandit Problem
- Heuristic Solution Algorithms
- Mitigating CMB B-mode Foregrounds

B-modes: Scale Dependence

B-modes: Scale Dependence

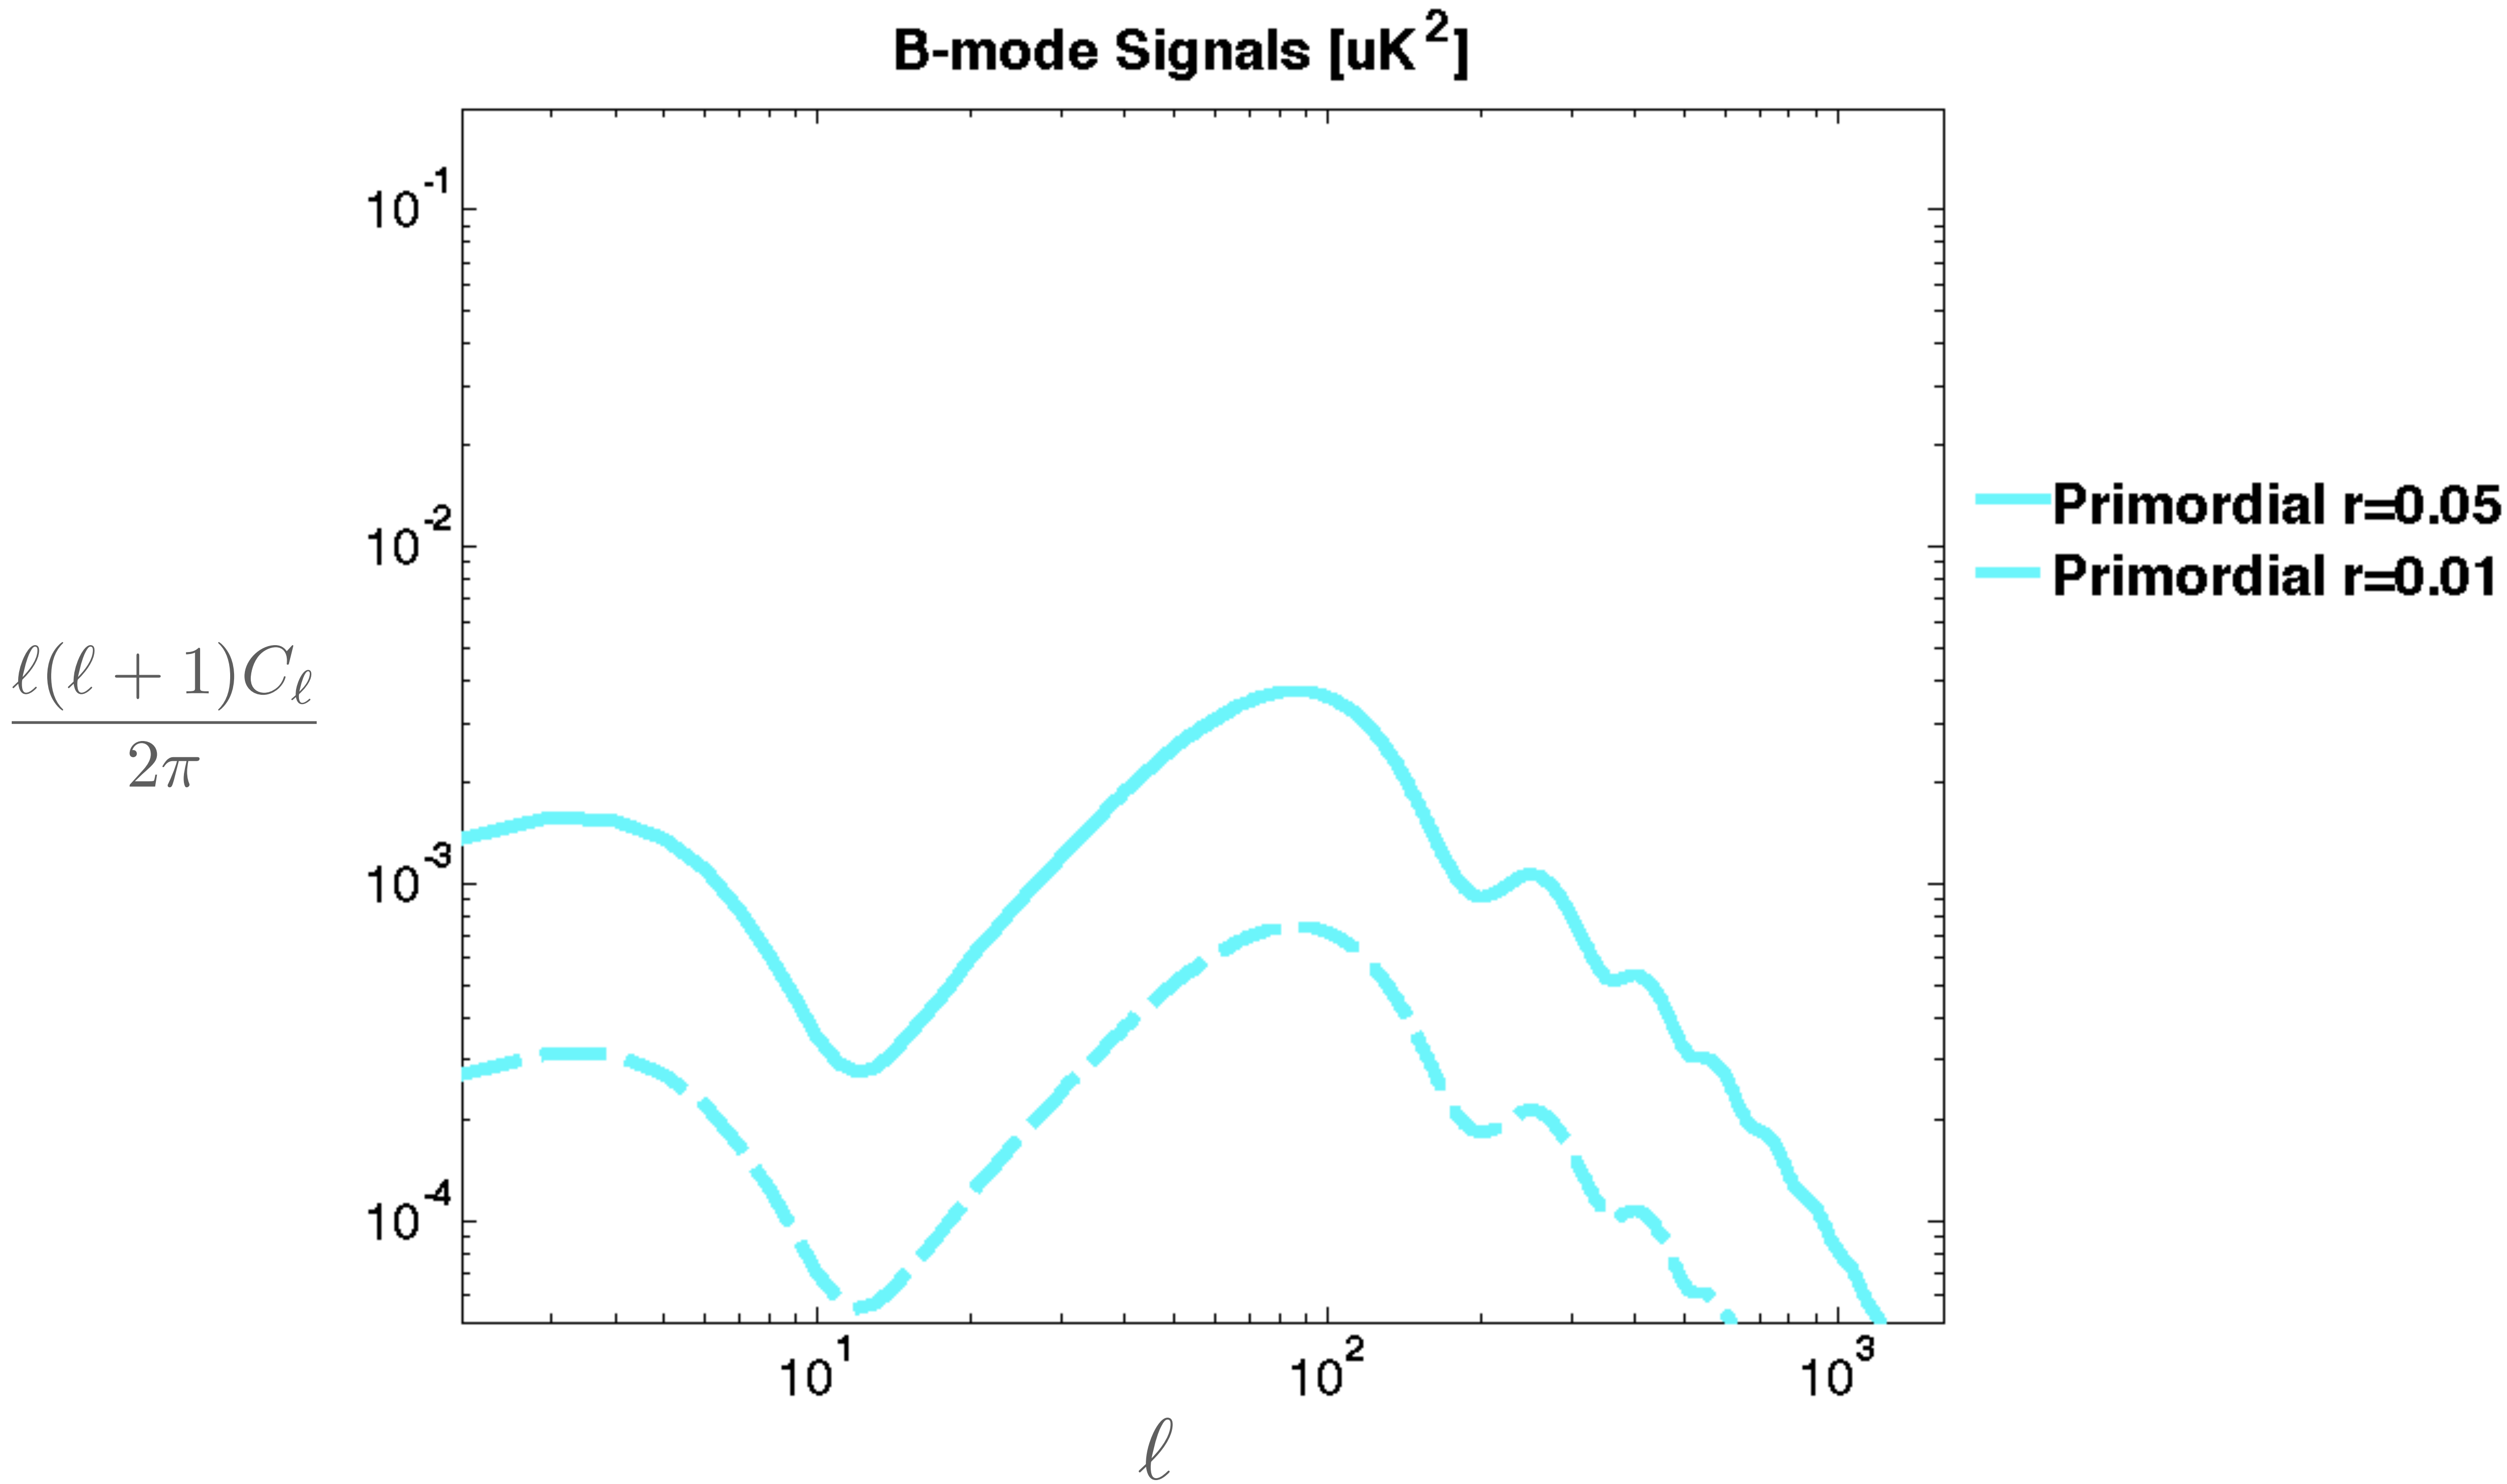

B-modes: Scale Dependence

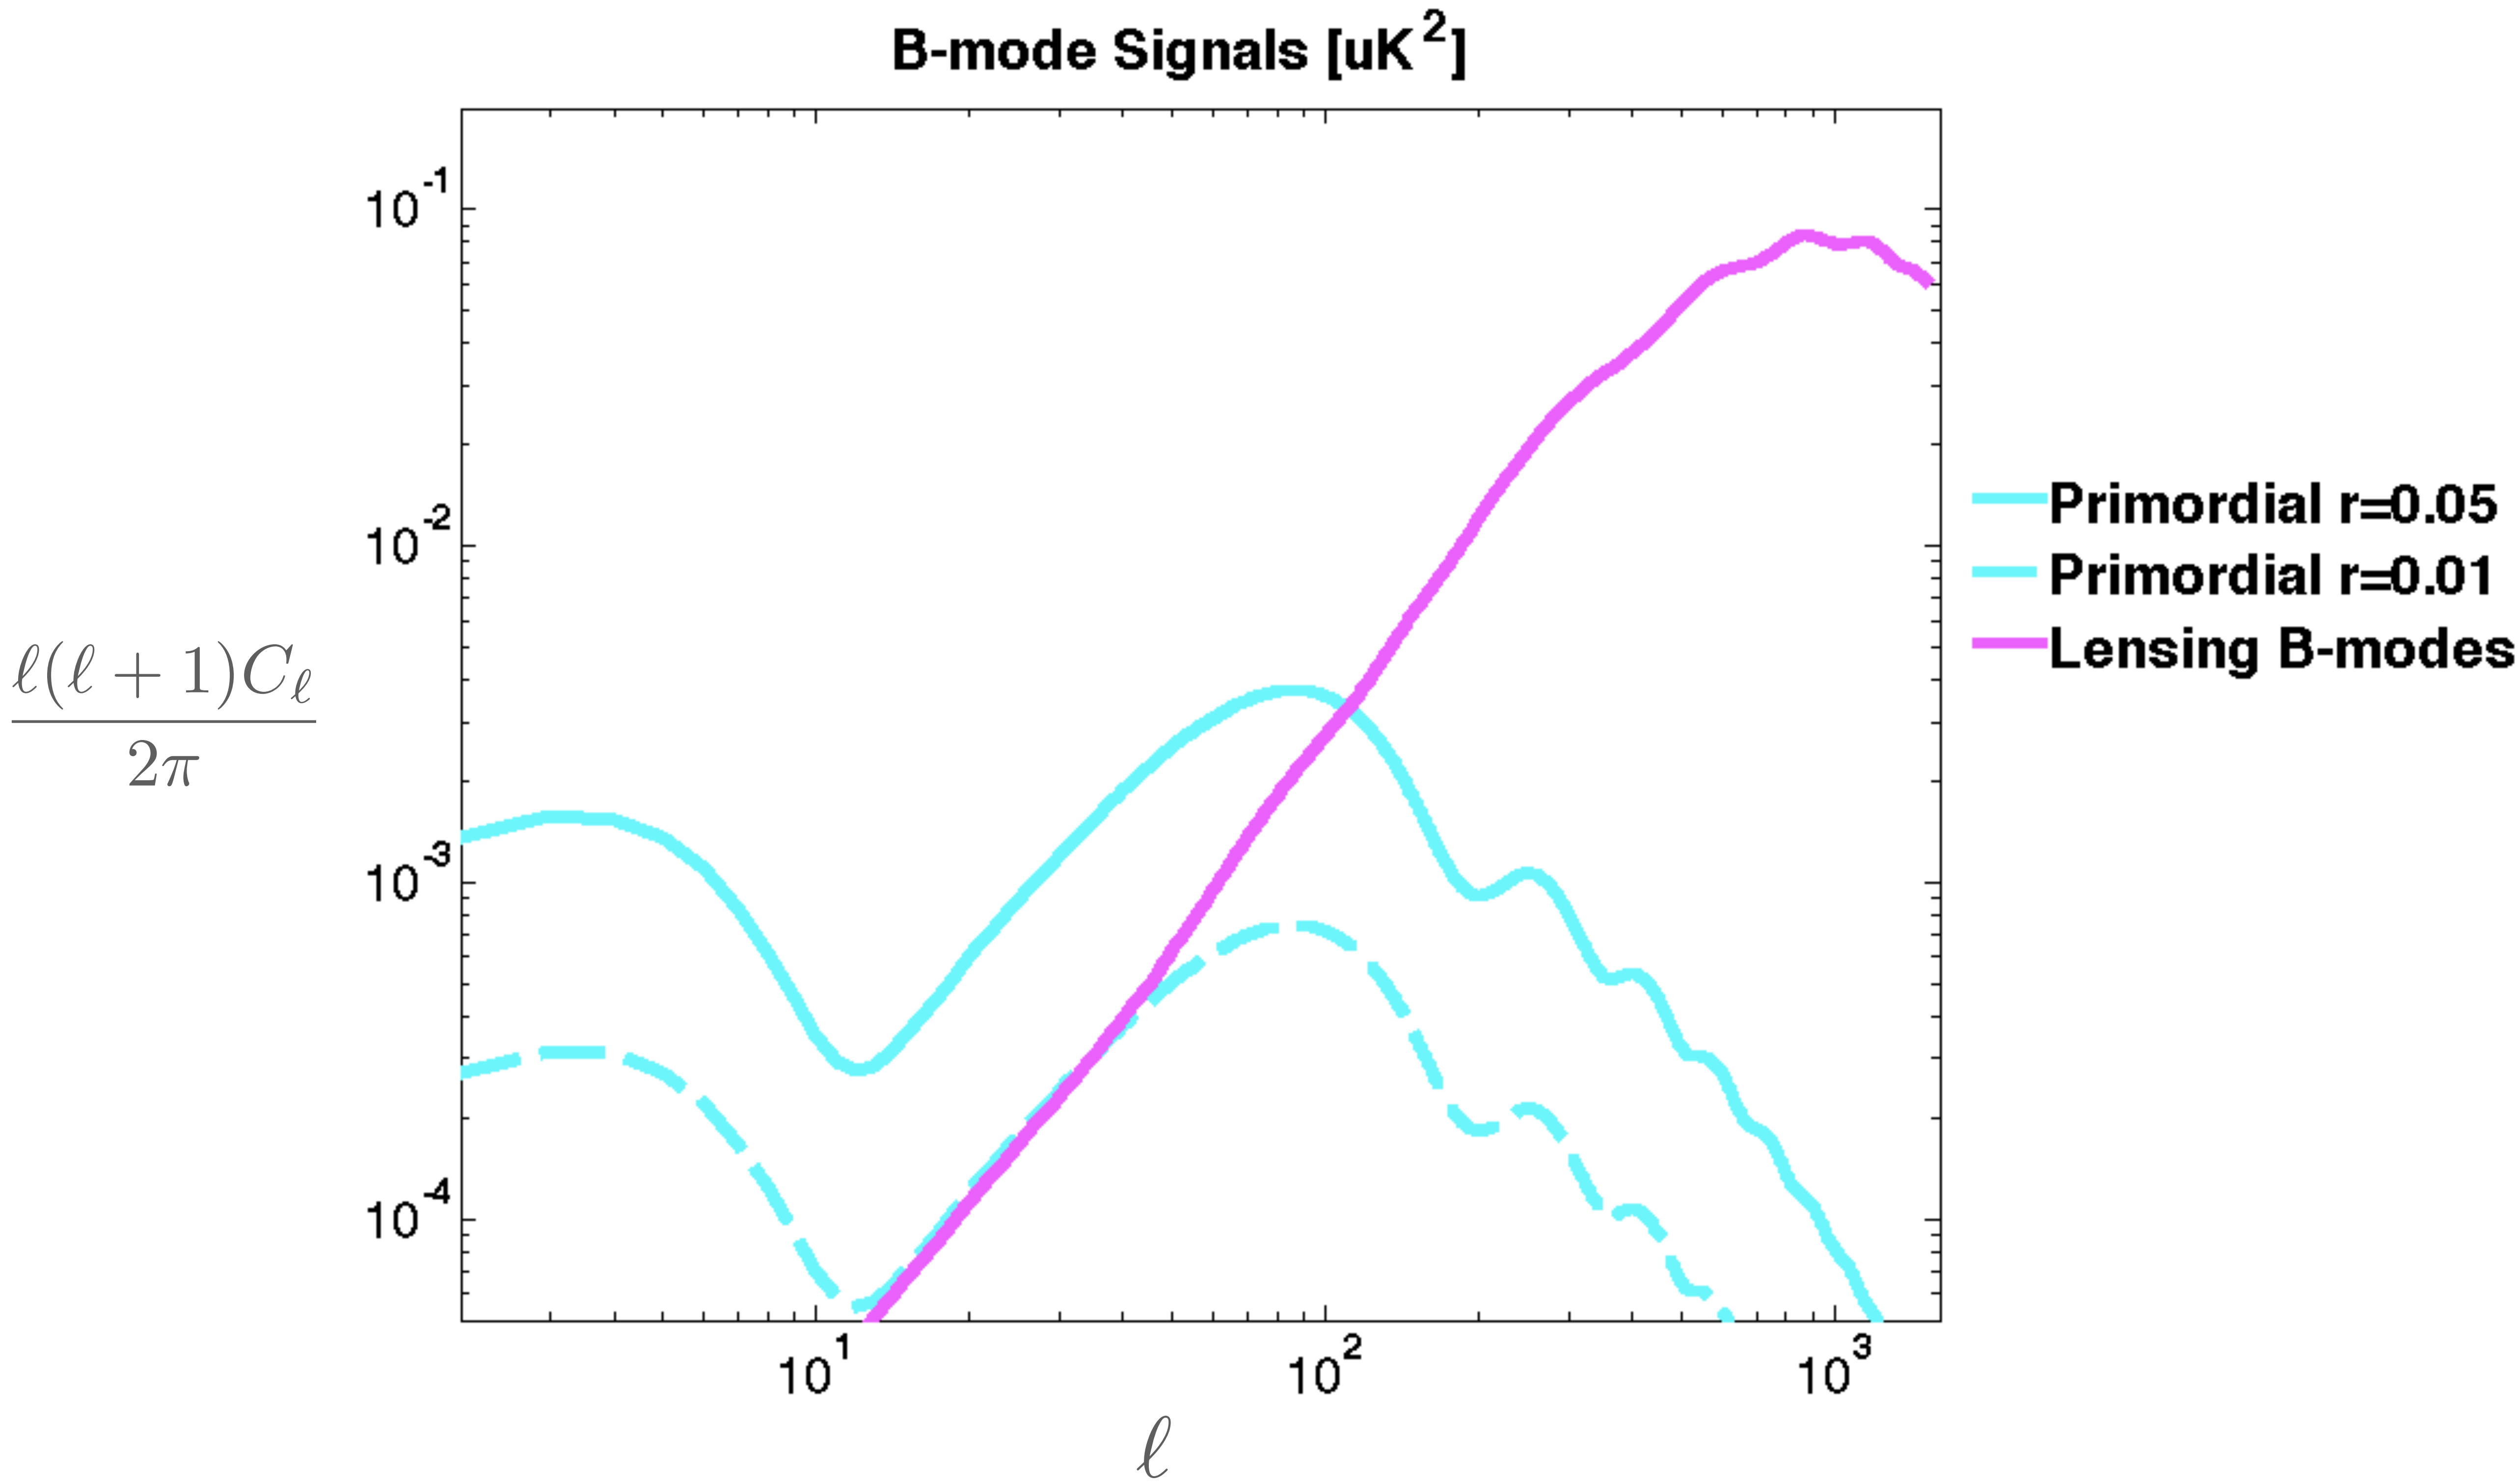

B-modes: Scale Dependence

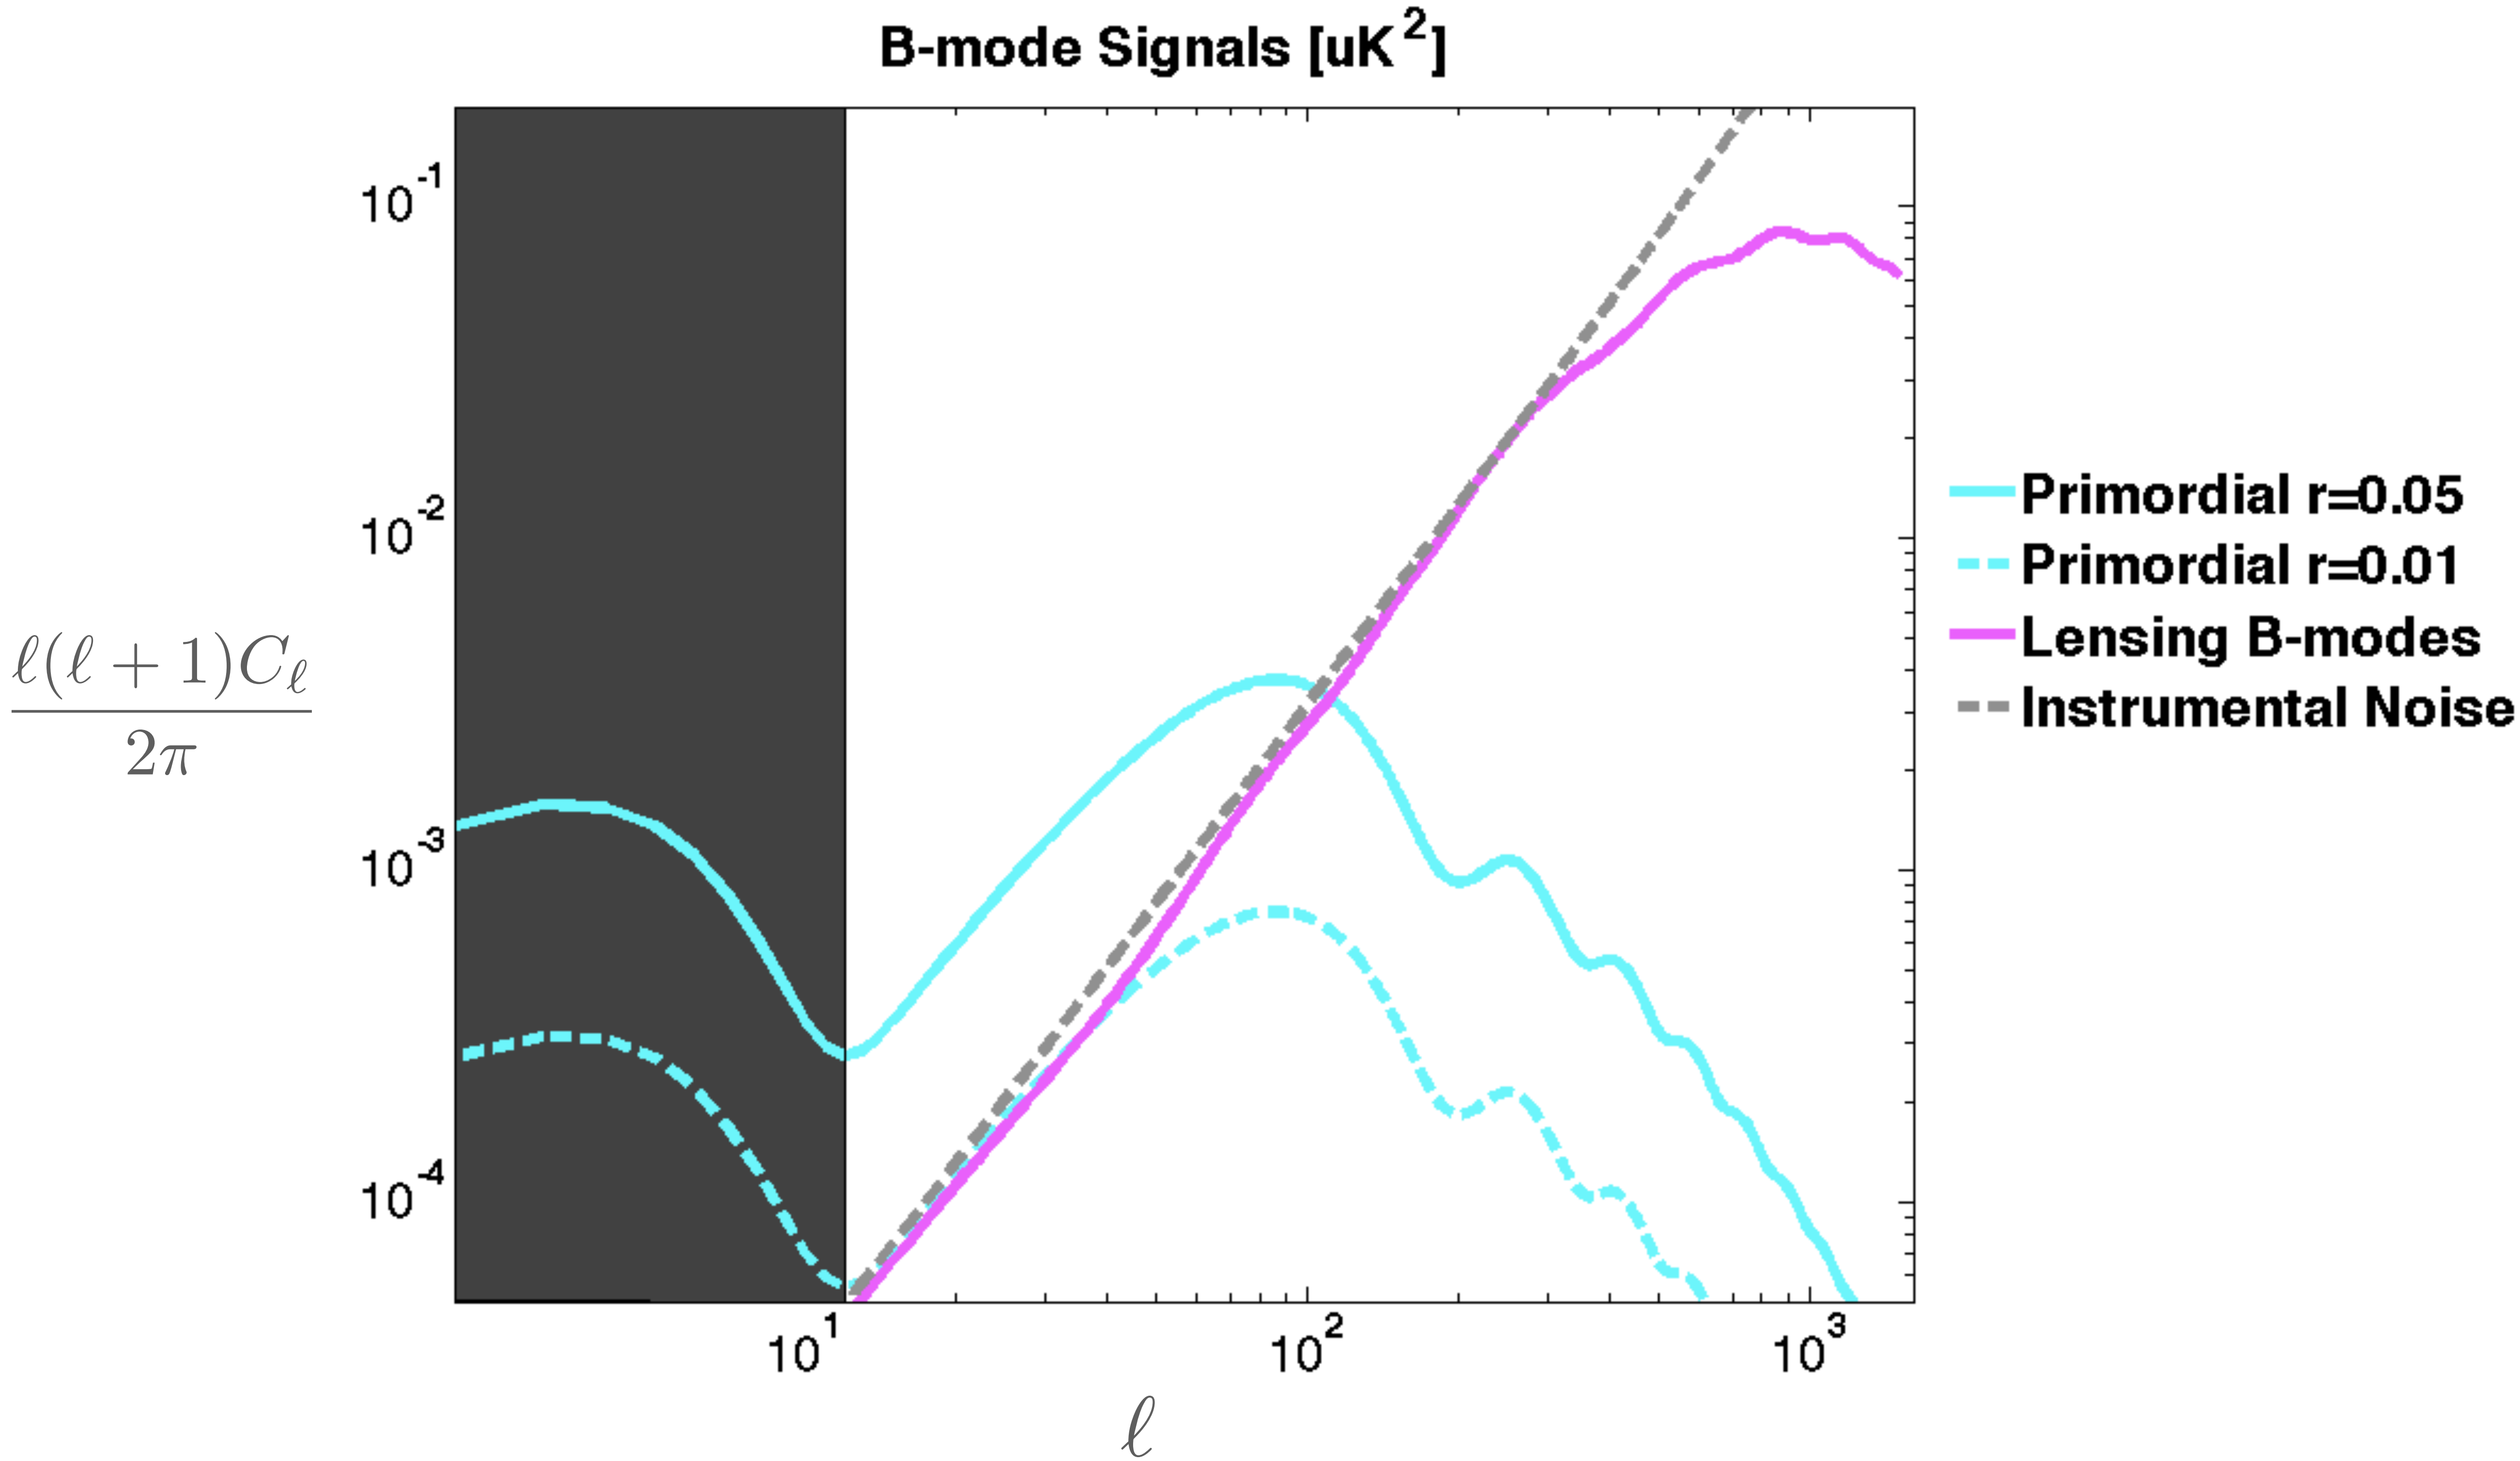

B-modes: Scale Dependence

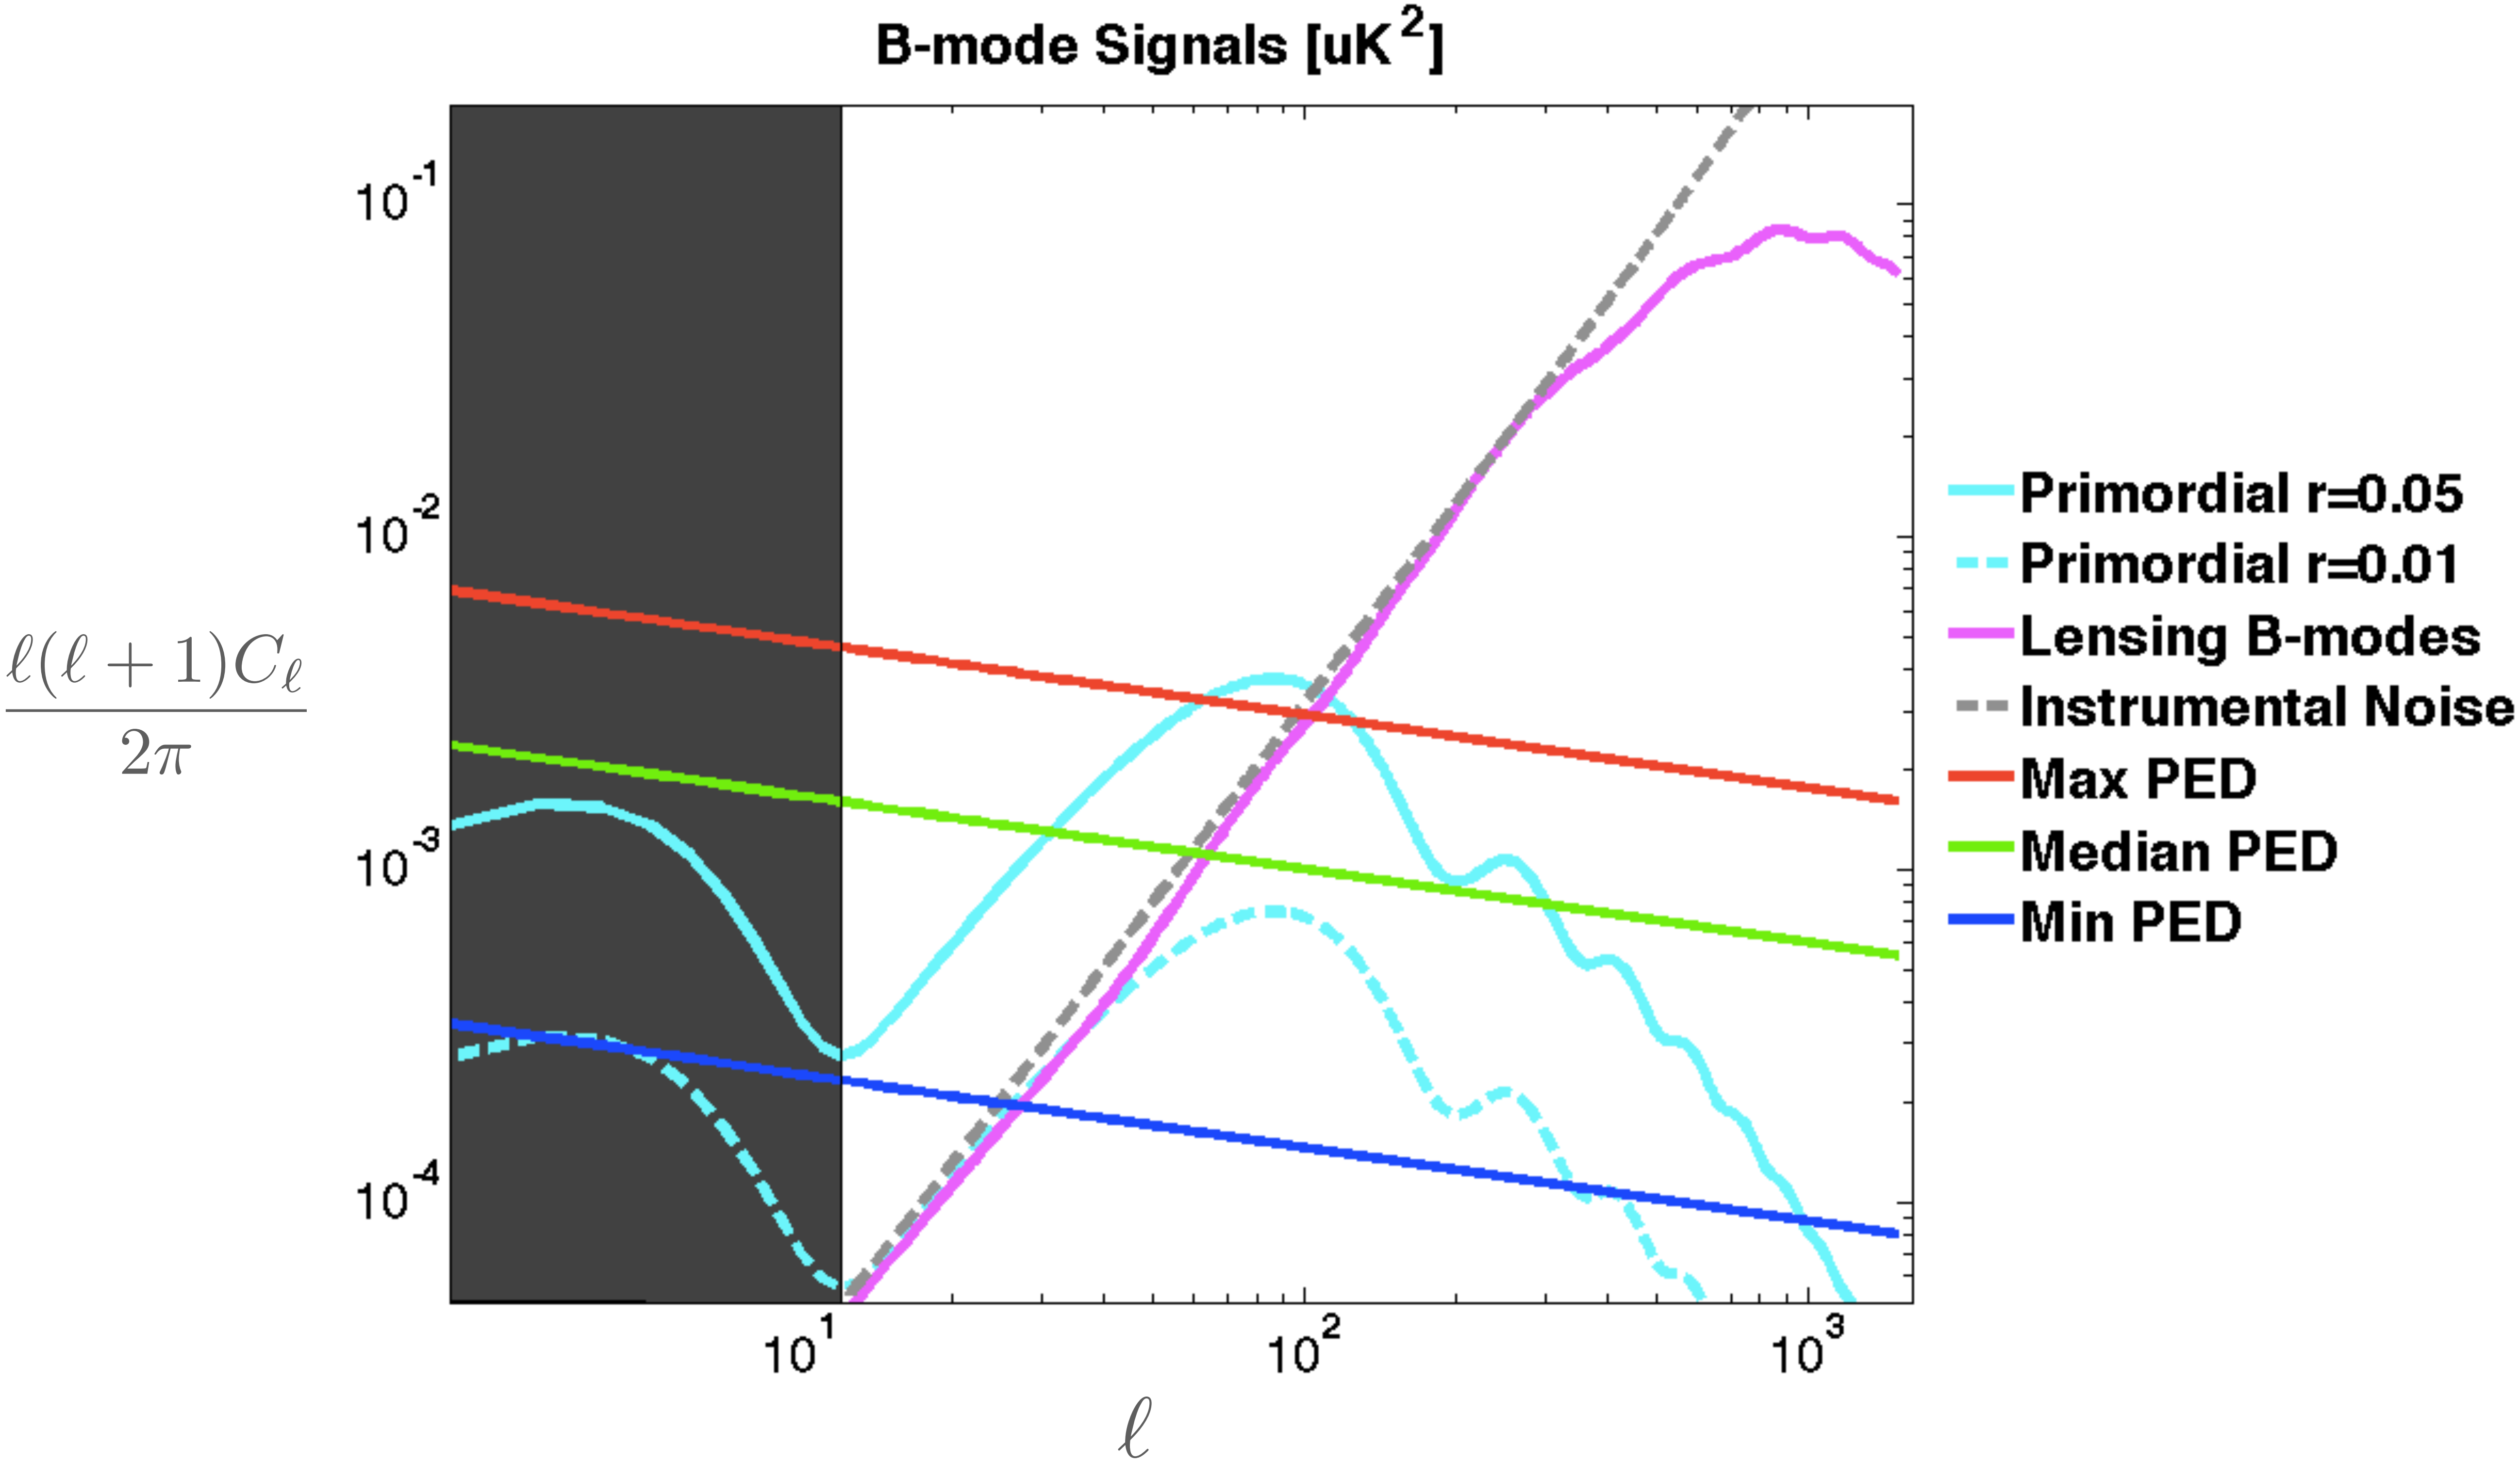

B-modes: Scale Dependence

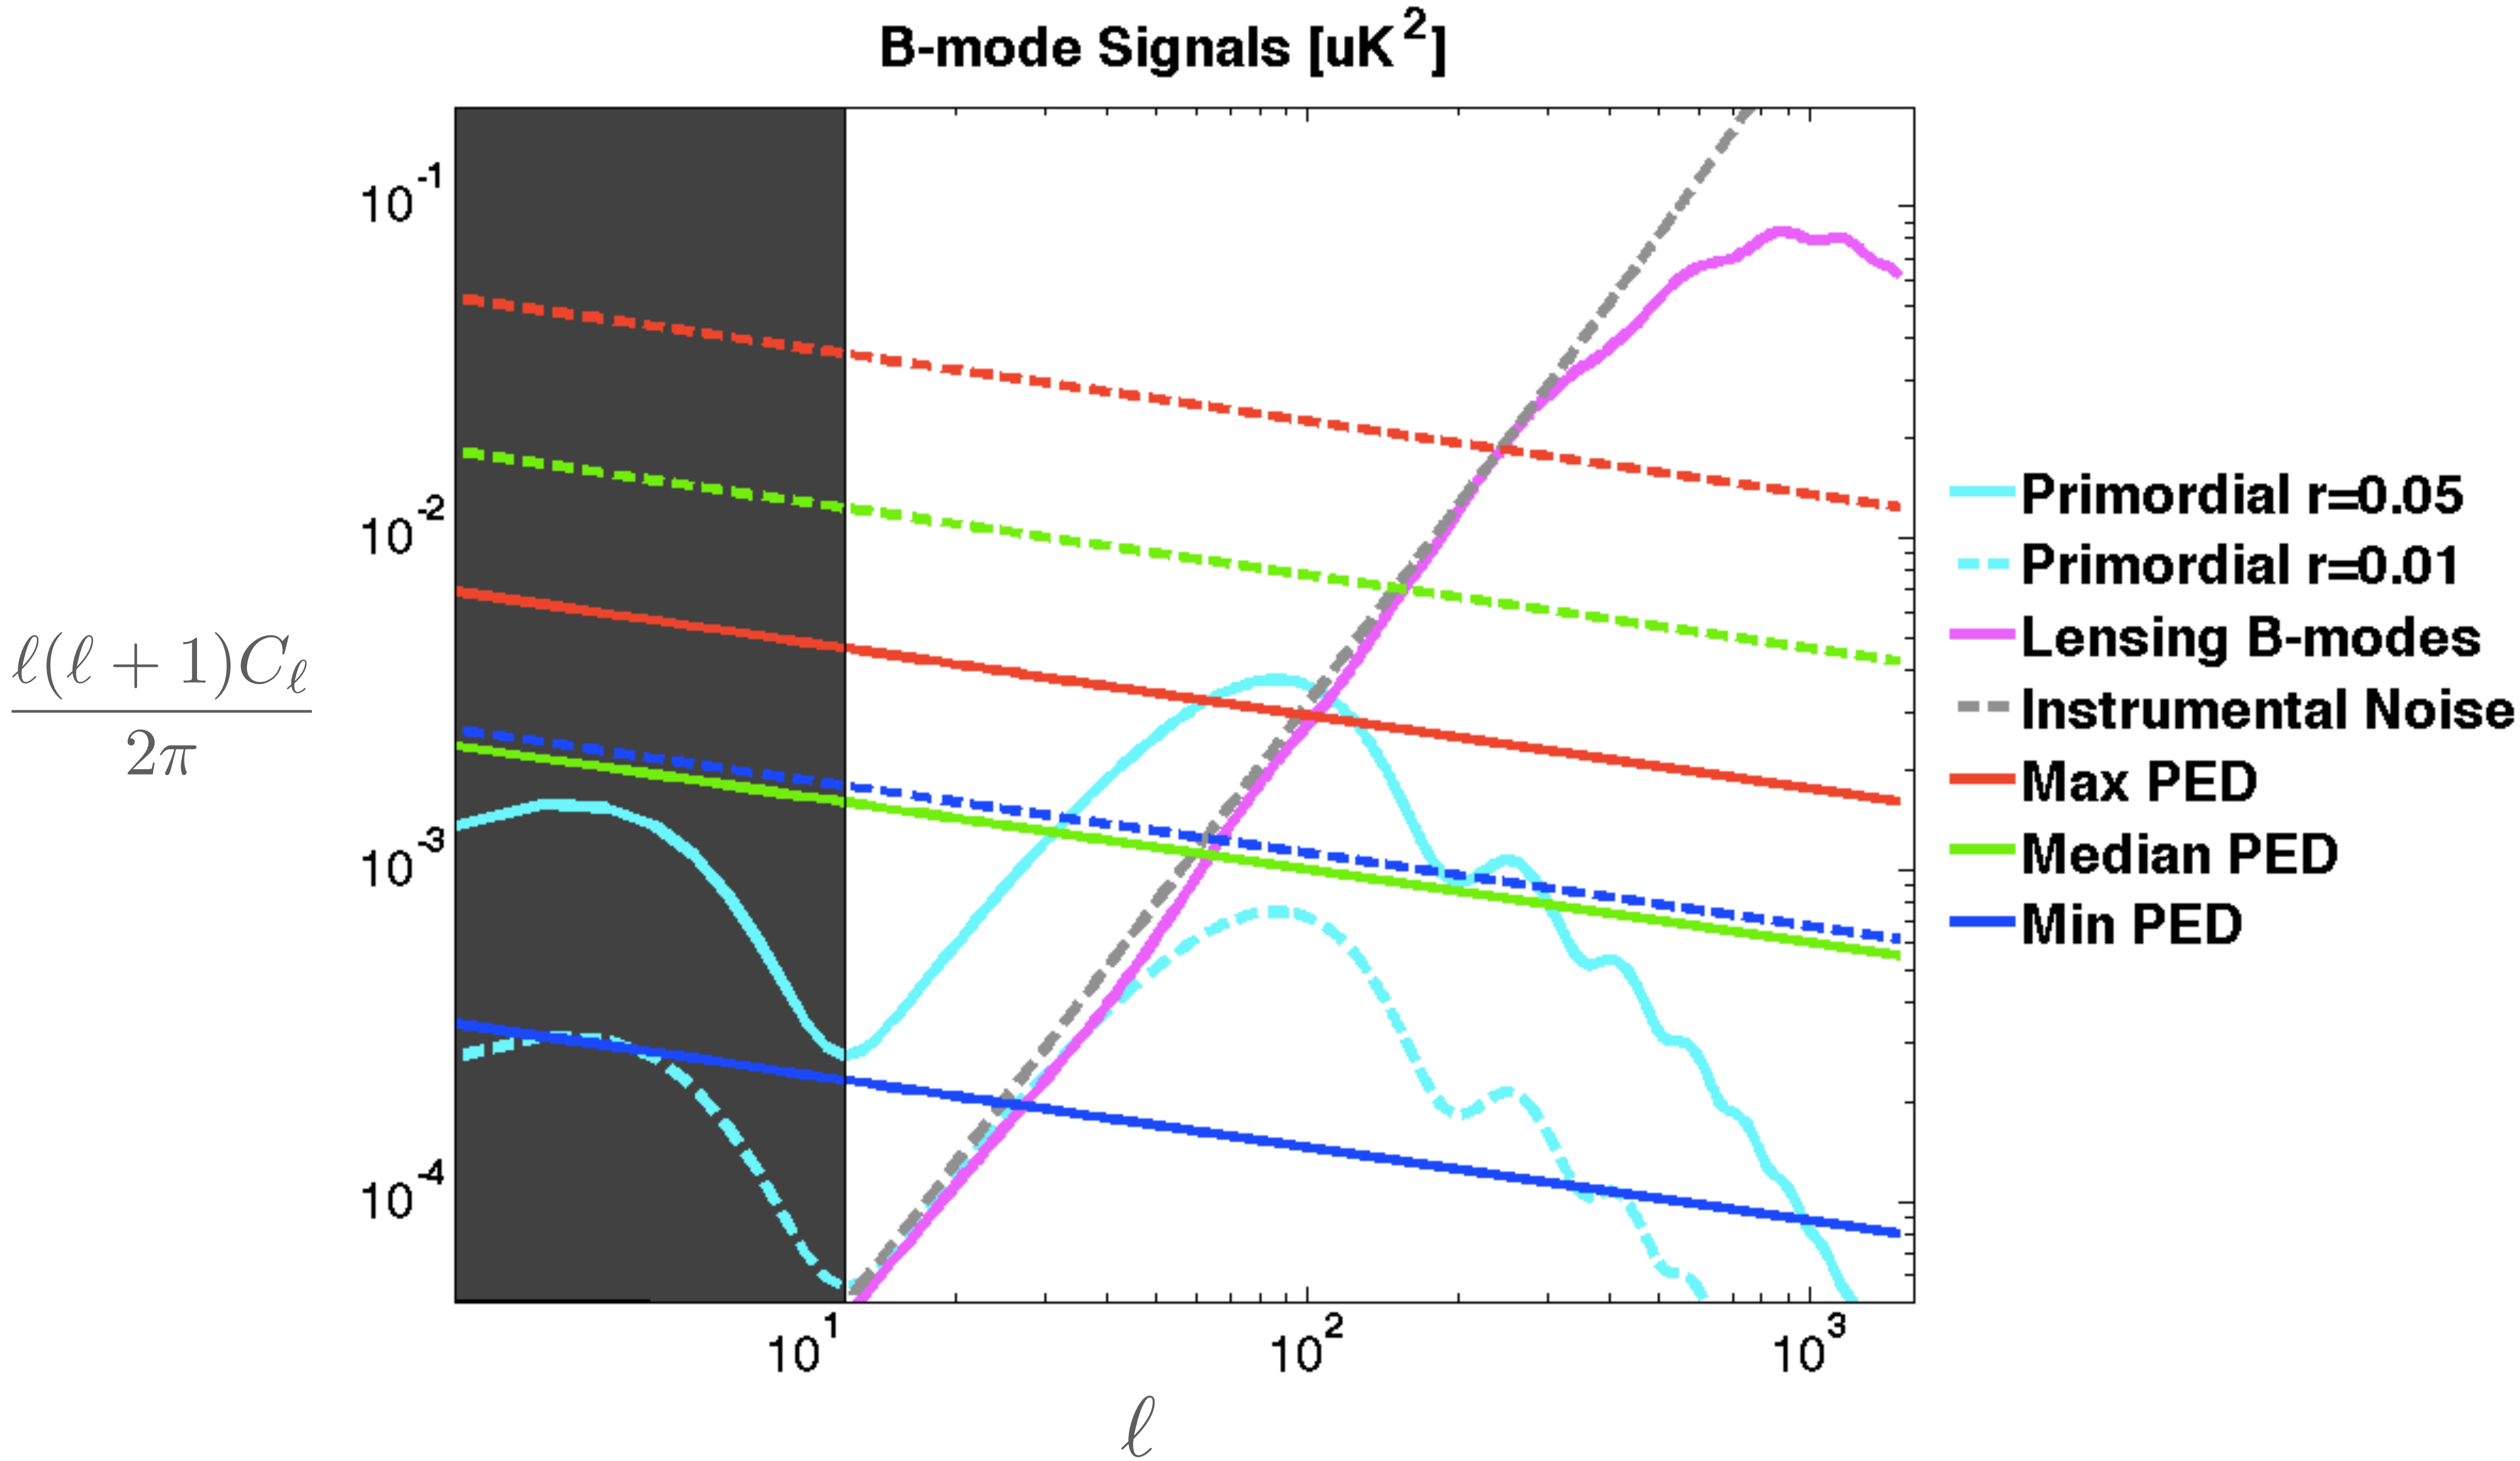

B-mode Surveys: Bandit Simulations

B-mode Surveys: Bandit Simulations

What do we measure?

B-mode Surveys: Bandit Simulations

What do we measure?

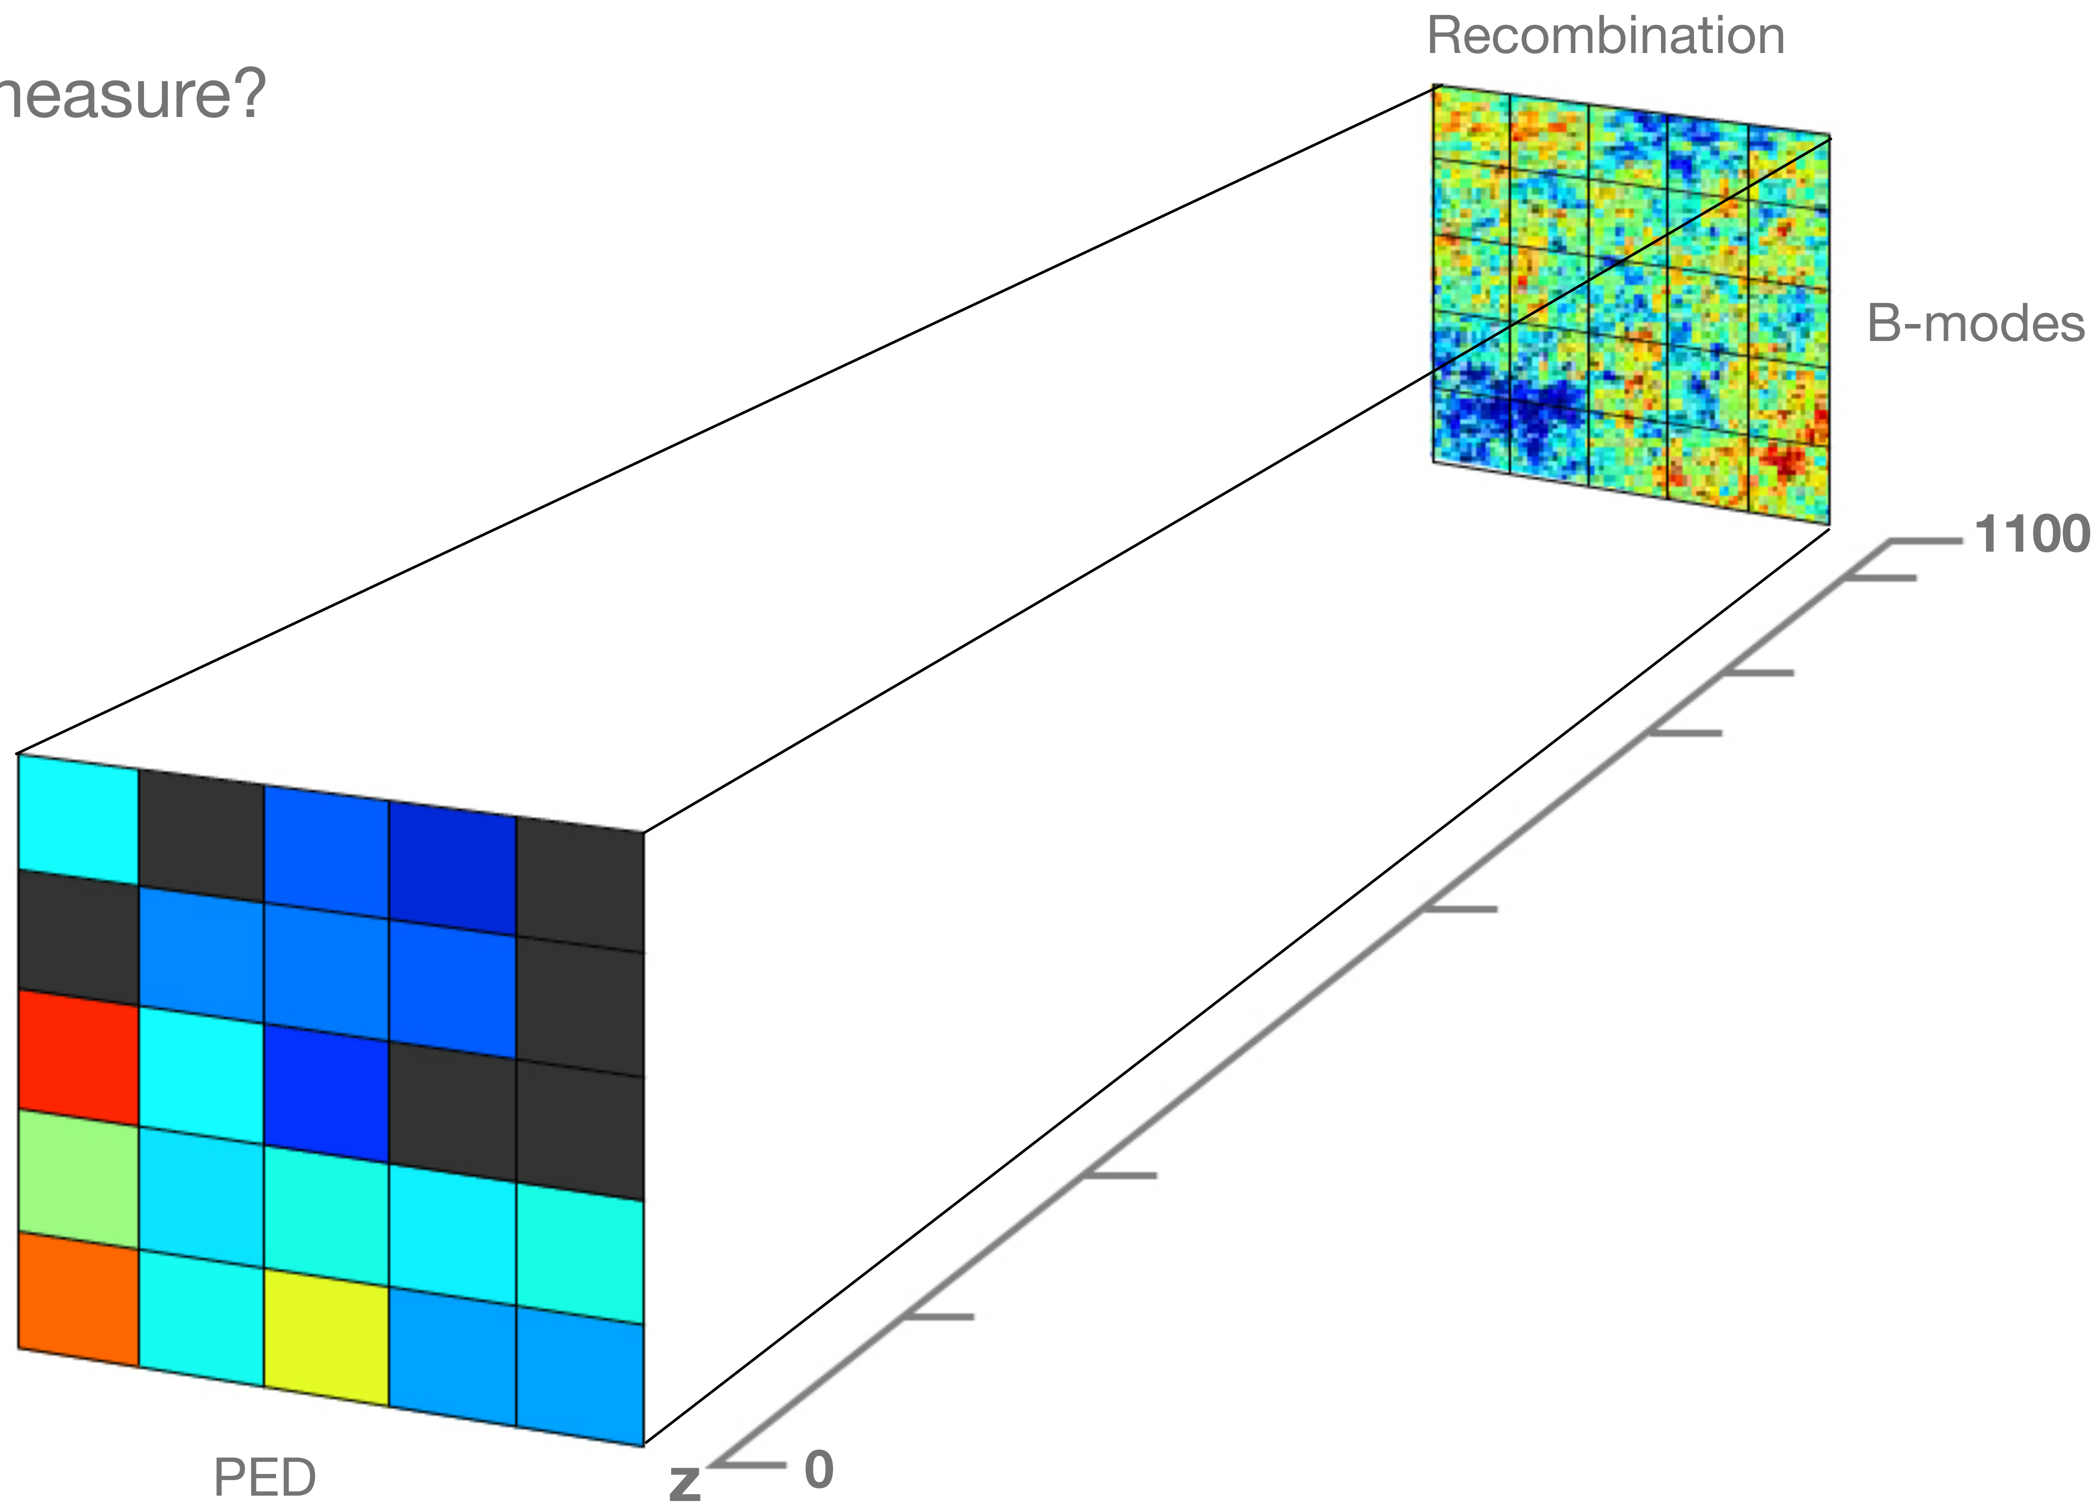

B-mode Surveys: Bandit Simulations

What do we measure?

- Initially: the PED amplitudes: A_p

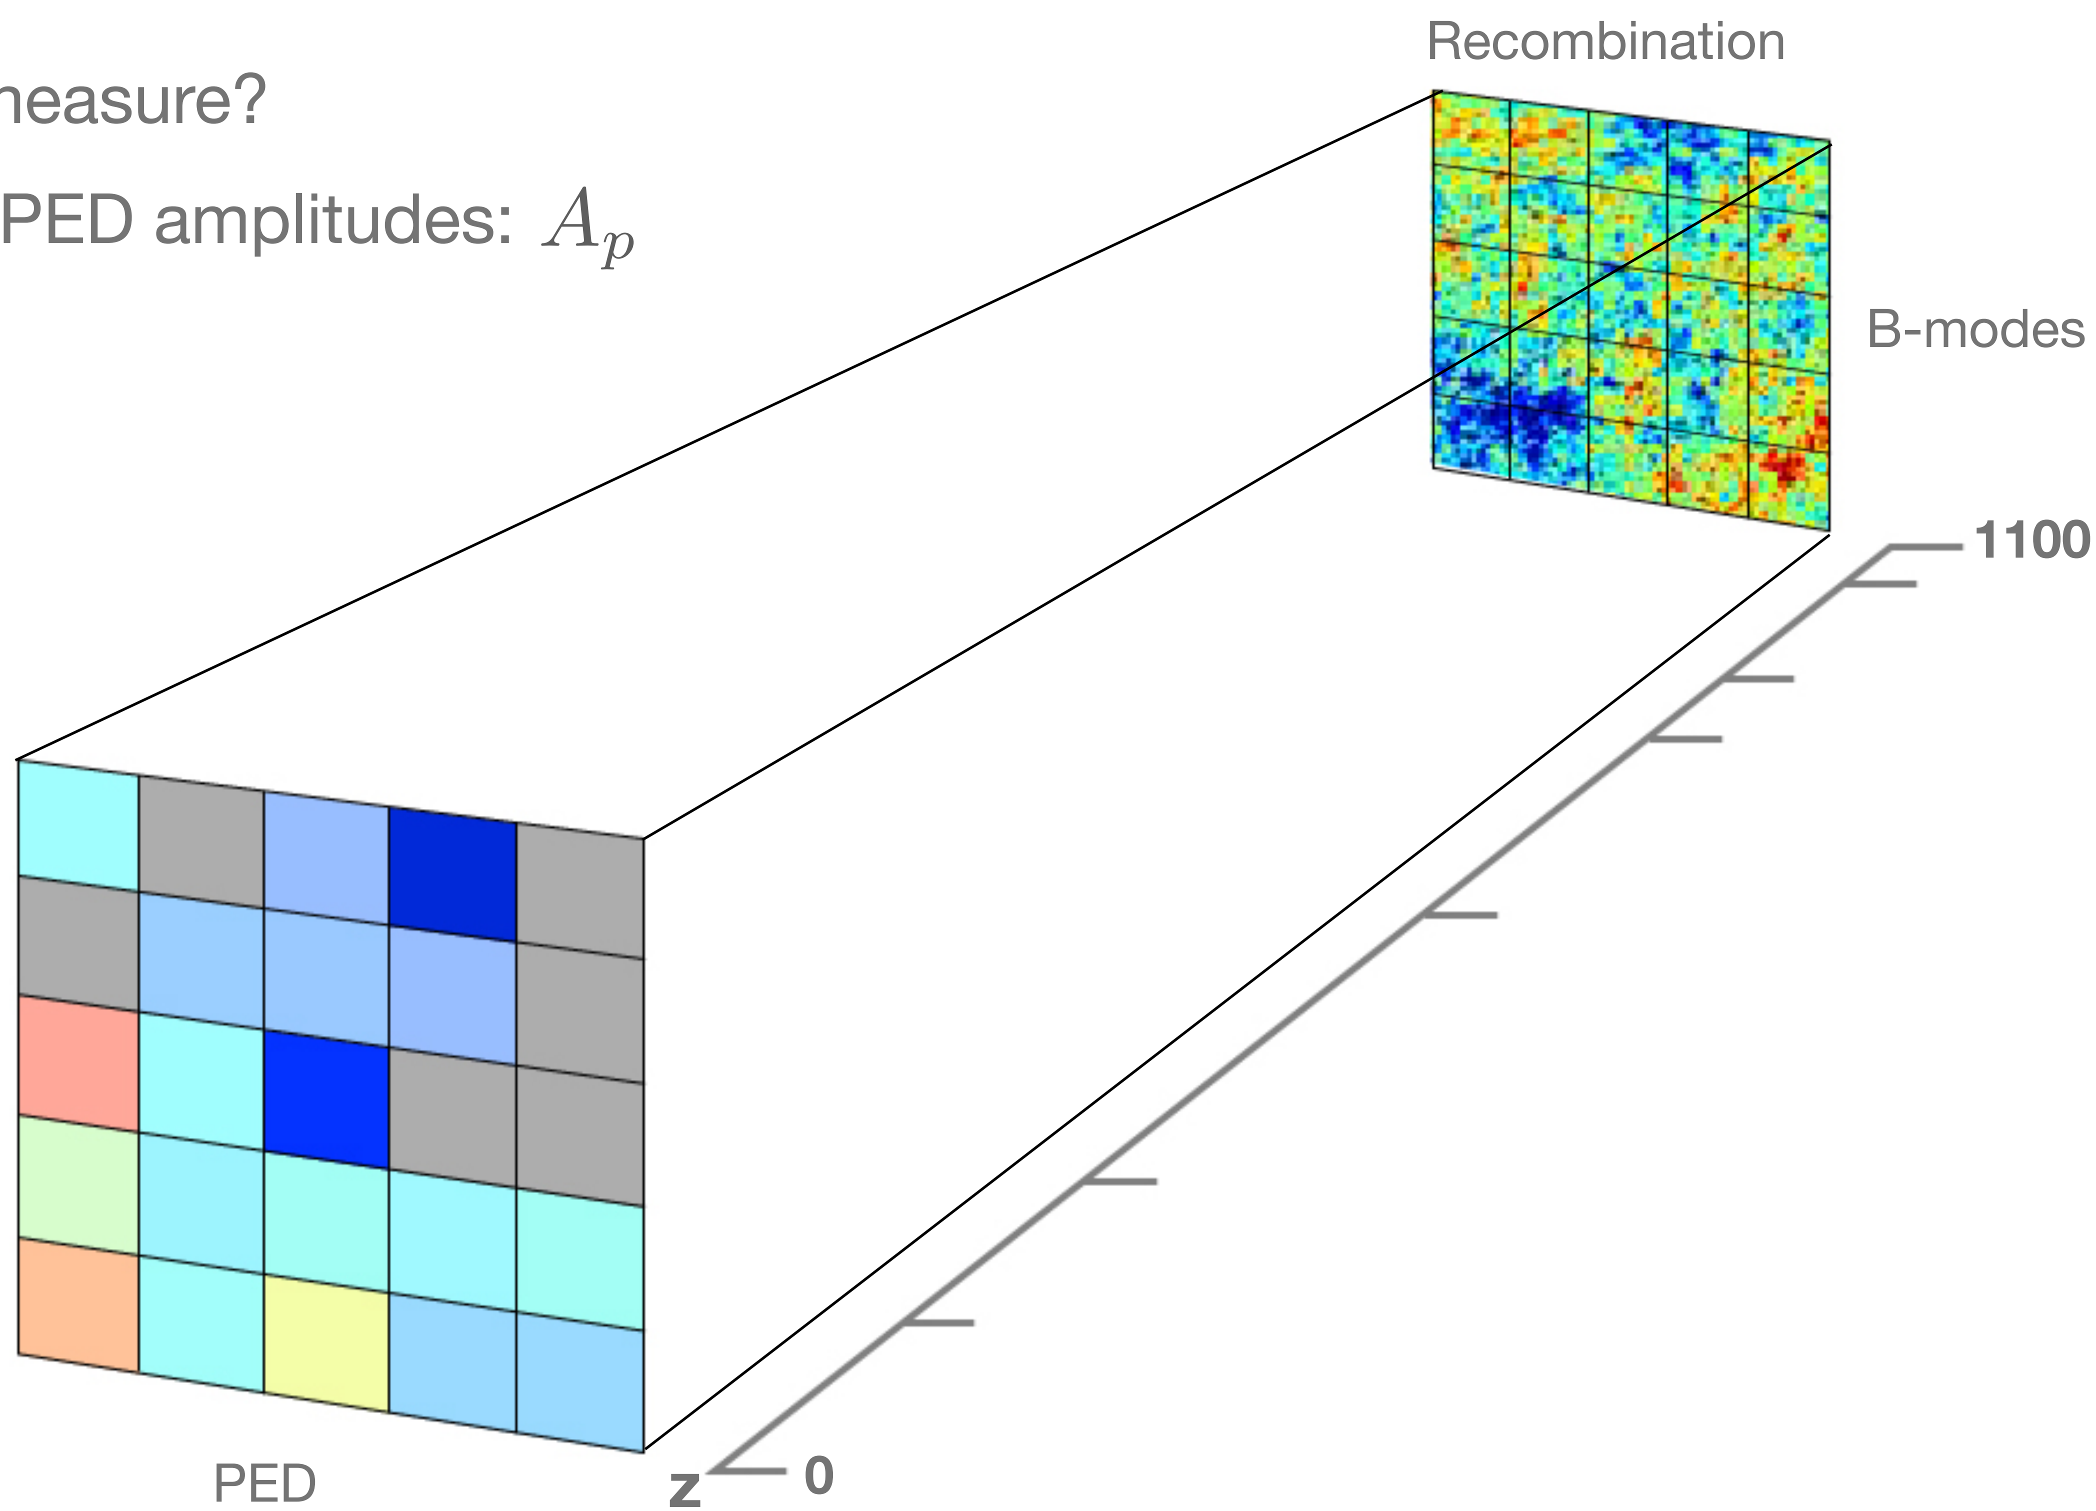

B-mode Surveys: Bandit Simulations

B-mode Surveys: Bandit Simulations

- PED ``1-sigma'' uncertainty:

$$\sigma_\ell^{\widehat{A}_p} = \sqrt{\frac{2}{f_{\text{sky}}(2\ell+1)}}\left(A_p \tilde{C}_\ell^D + \alpha C_\ell^L + f_{\text{sky}} w^{-1}(T) e^{\ell^2 \sigma_b^2}\right)$$

B-mode Surveys: Bandit Simulations

- PED ``1-sigma'' uncertainty:

$$\sigma_\ell^{\widehat{A}_p} = \sqrt{\frac{2}{f_{\rm sky}(2\ell+1)}}\left(A_p \underbrace{\tilde{C}_\ell^D}_{\text{foregrounds}} + \alpha C_\ell^L + f_{\rm sky} w^{-1}(T) e^{\ell^2\sigma_b^2}\right)$$

B-mode Surveys: Bandit Simulations

- PED ``1-sigma'' uncertainty:

foregrounds lensing

$$\sigma_\ell^{\widehat{A}_p} = \sqrt{\frac{2}{f_{\mathrm{sky}}(2\ell+1)}} \left(A_p \tilde{C}_\ell^D + \alpha C_\ell^L + f_{\mathrm{sky}} w^{-1}(T) e^{\ell^2 \sigma_b^2} \right)$$

B-mode Surveys: Bandit Simulations

- PED ``1-sigma'' uncertainty:

$$\sigma_\ell^{\widehat{A}_p} = \sqrt{\frac{2}{f_{\rm sky}(2\ell+1)}} \left(\underbrace{A_p \tilde{C}_\ell^D}_{\text{foregrounds}} + \underbrace{\alpha C_\ell^L}_{\text{lensing}} + \underbrace{f_{\rm sky} w^{-1}(T) e^{\ell^2 \sigma_b^2}}_{\text{instrumental noise}} \right)$$

B-mode Surveys: Bandit Simulations

- PED "1-sigma" uncertainty:

foregrounds lensing instrumental noise

$$\sigma_\ell^{\widehat{A}_p} = \sqrt{\frac{2}{f_{\rm sky}(2\ell+1)}} \left(A_p \tilde{C}_\ell^D + \alpha C_\ell^L + f_{\rm sky} w^{-1}(T) e^{\ell^2 \sigma_b^2} \right)$$

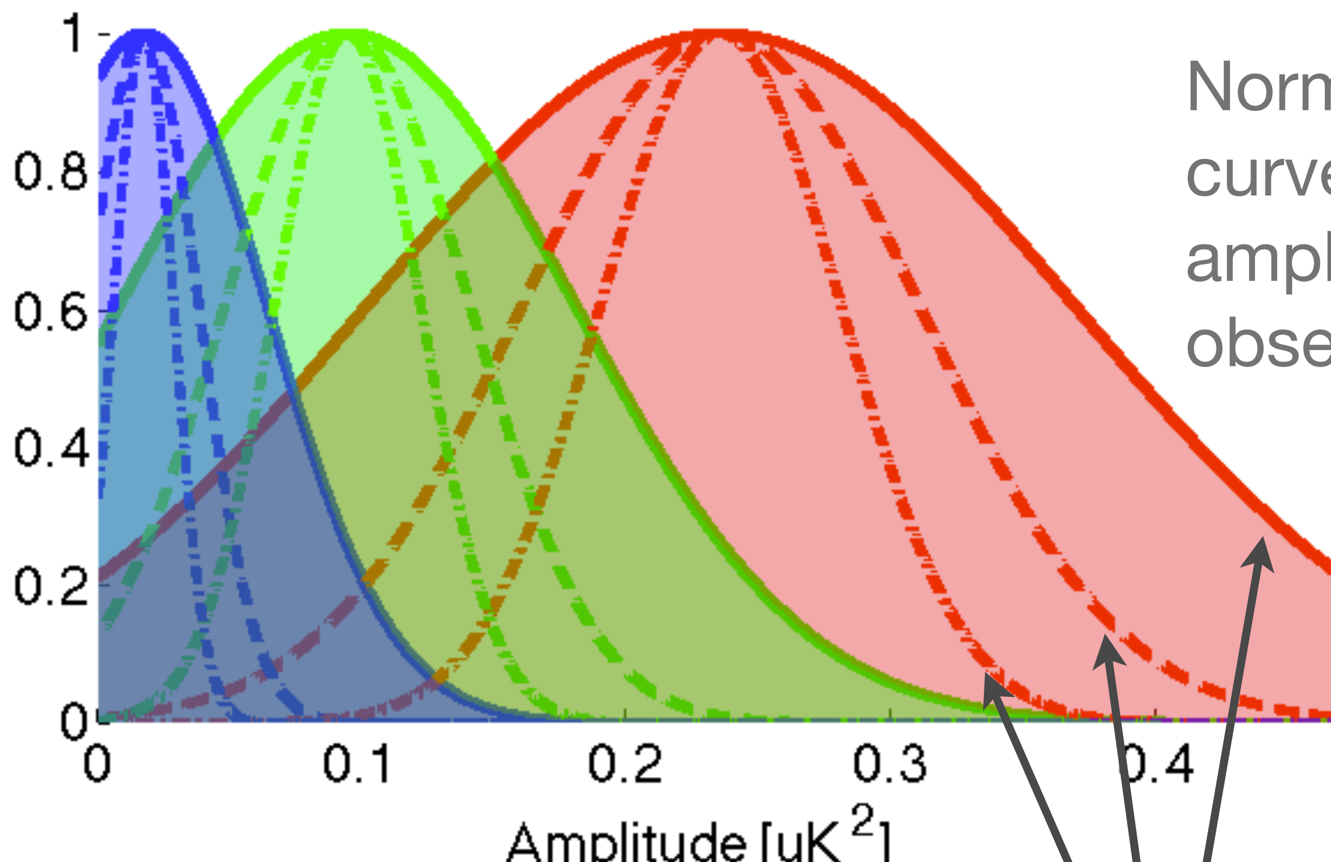

Normalized likelihood curves for measured amplitudes after 3, 10, 30 observation days.

$$\sigma^{\widehat{A}_p} = \left[\frac{f_{\rm sky}}{2} \sum_{\ell_{\rm min}}^{\ell_{\rm max}} \frac{(2\ell+1)(\tilde{C}_\ell^D)^2}{\left(A_p \tilde{C}_\ell^D + \alpha C_\ell^L + f_{\rm sky} w^{-1}(t_p) e^{\ell^2 \sigma_b^2} \right)^2} \right]^{-\frac{1}{2}}$$

B-mode Surveys: Bandit Simulations

B-mode Surveys: Bandit Simulations

What do we measure?

- Initially: the PED amplitudes: A_p

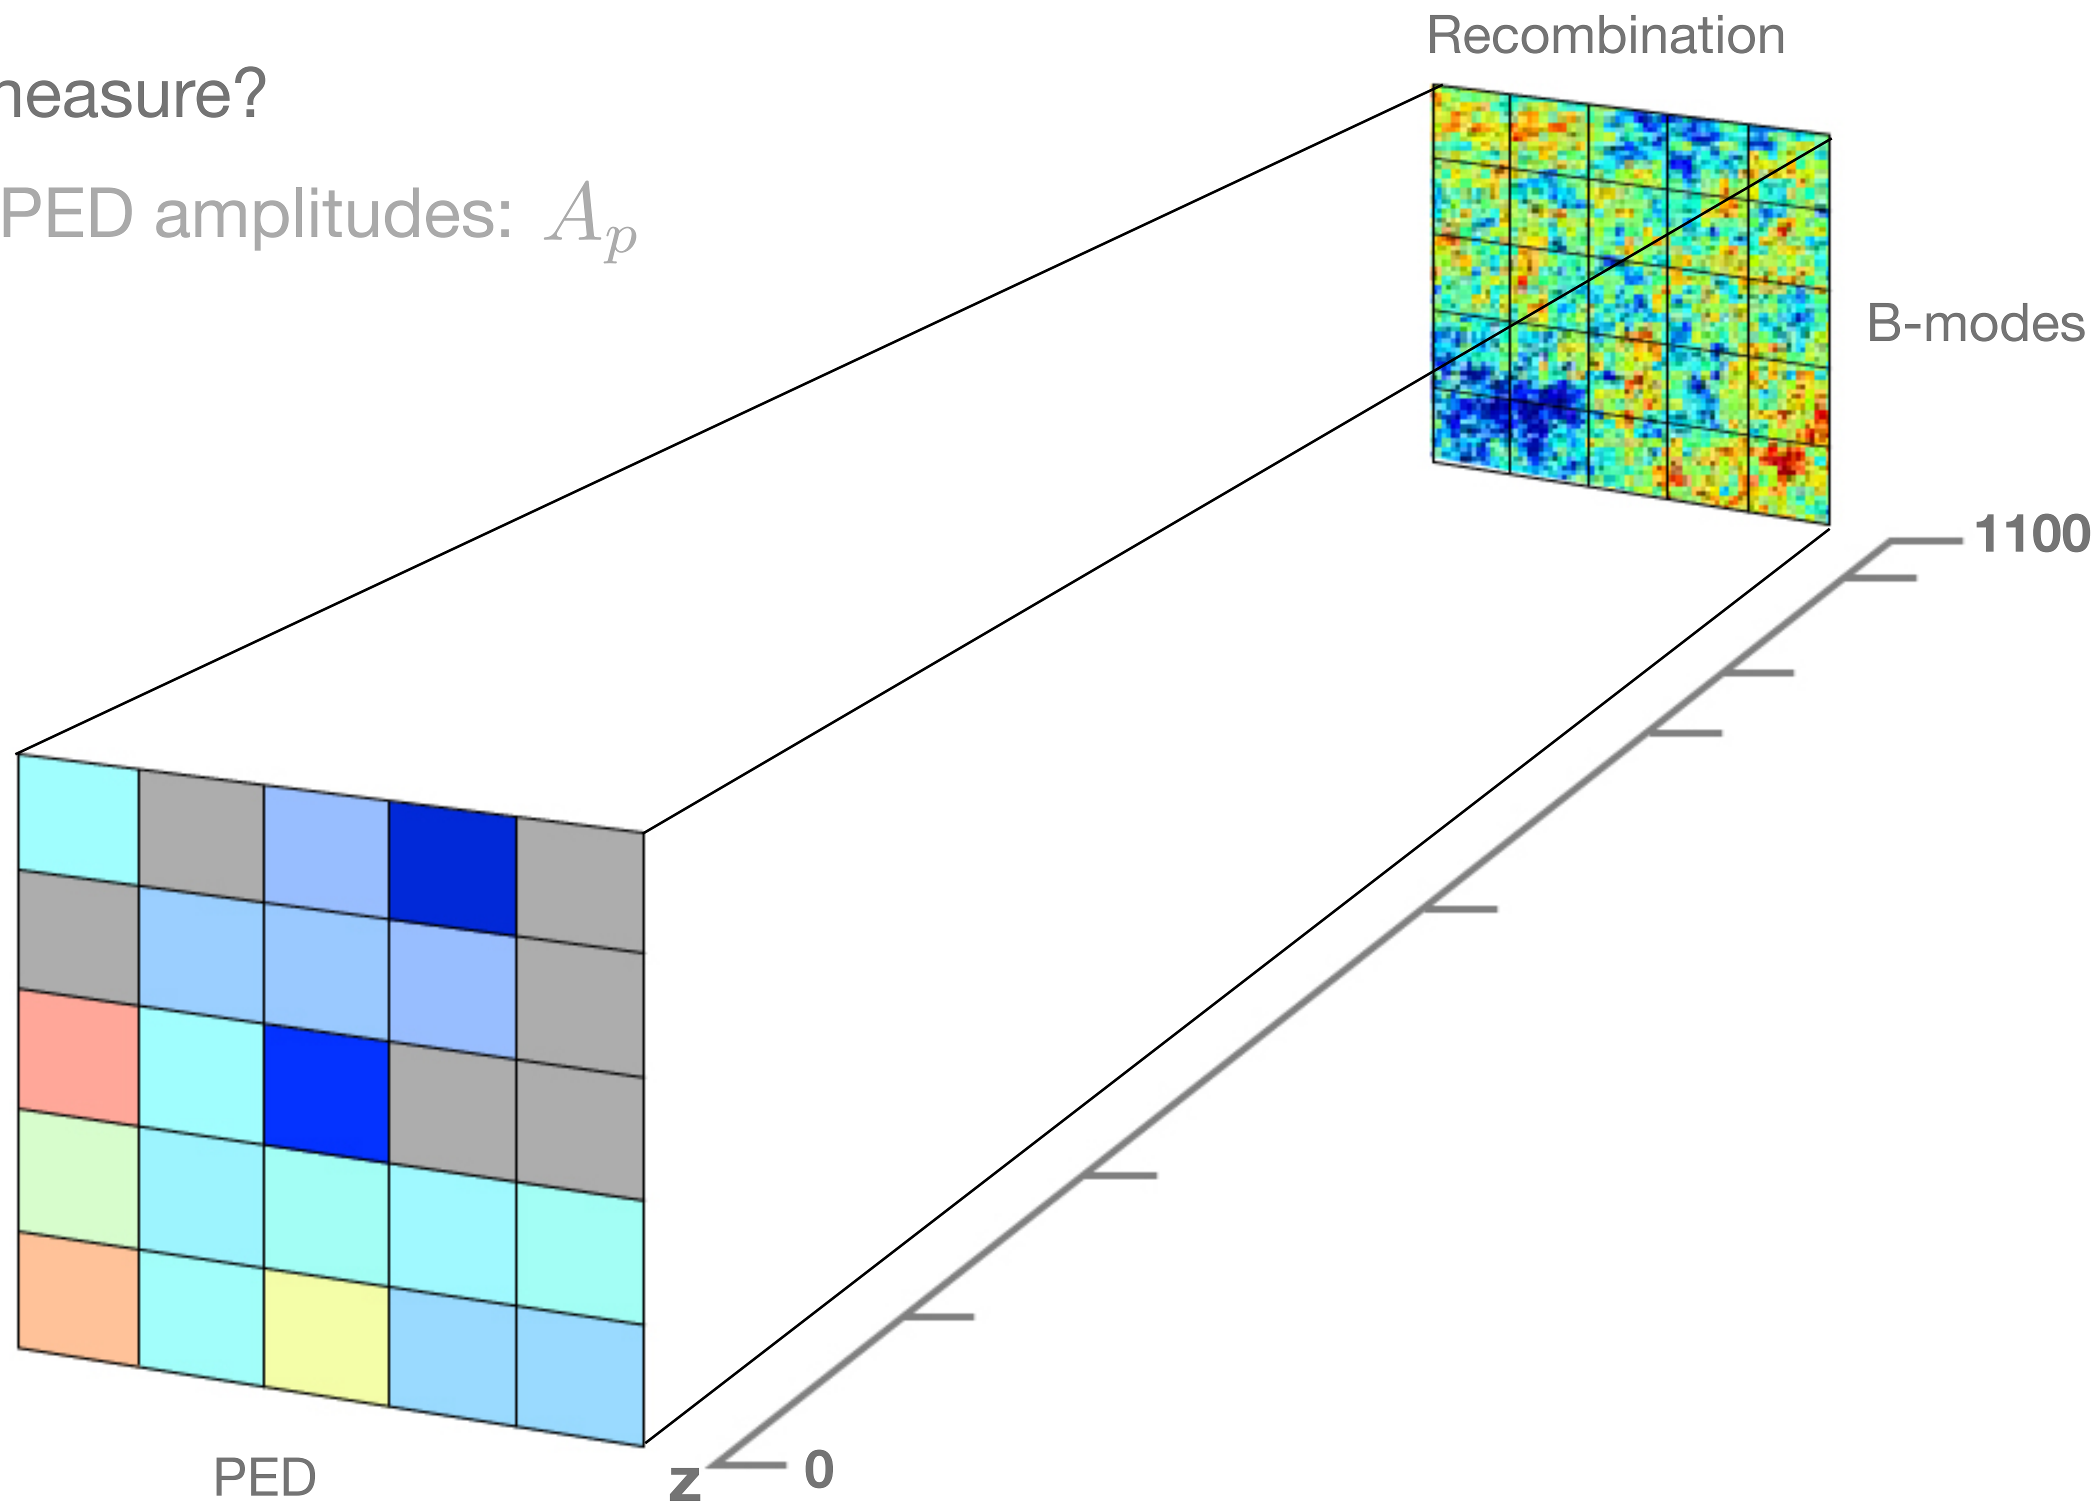

B-mode Surveys: Bandit Simulations

What do we measure?

- Initially: the PED amplitudes: A_p
- Ultimately: the CMB B-modes.

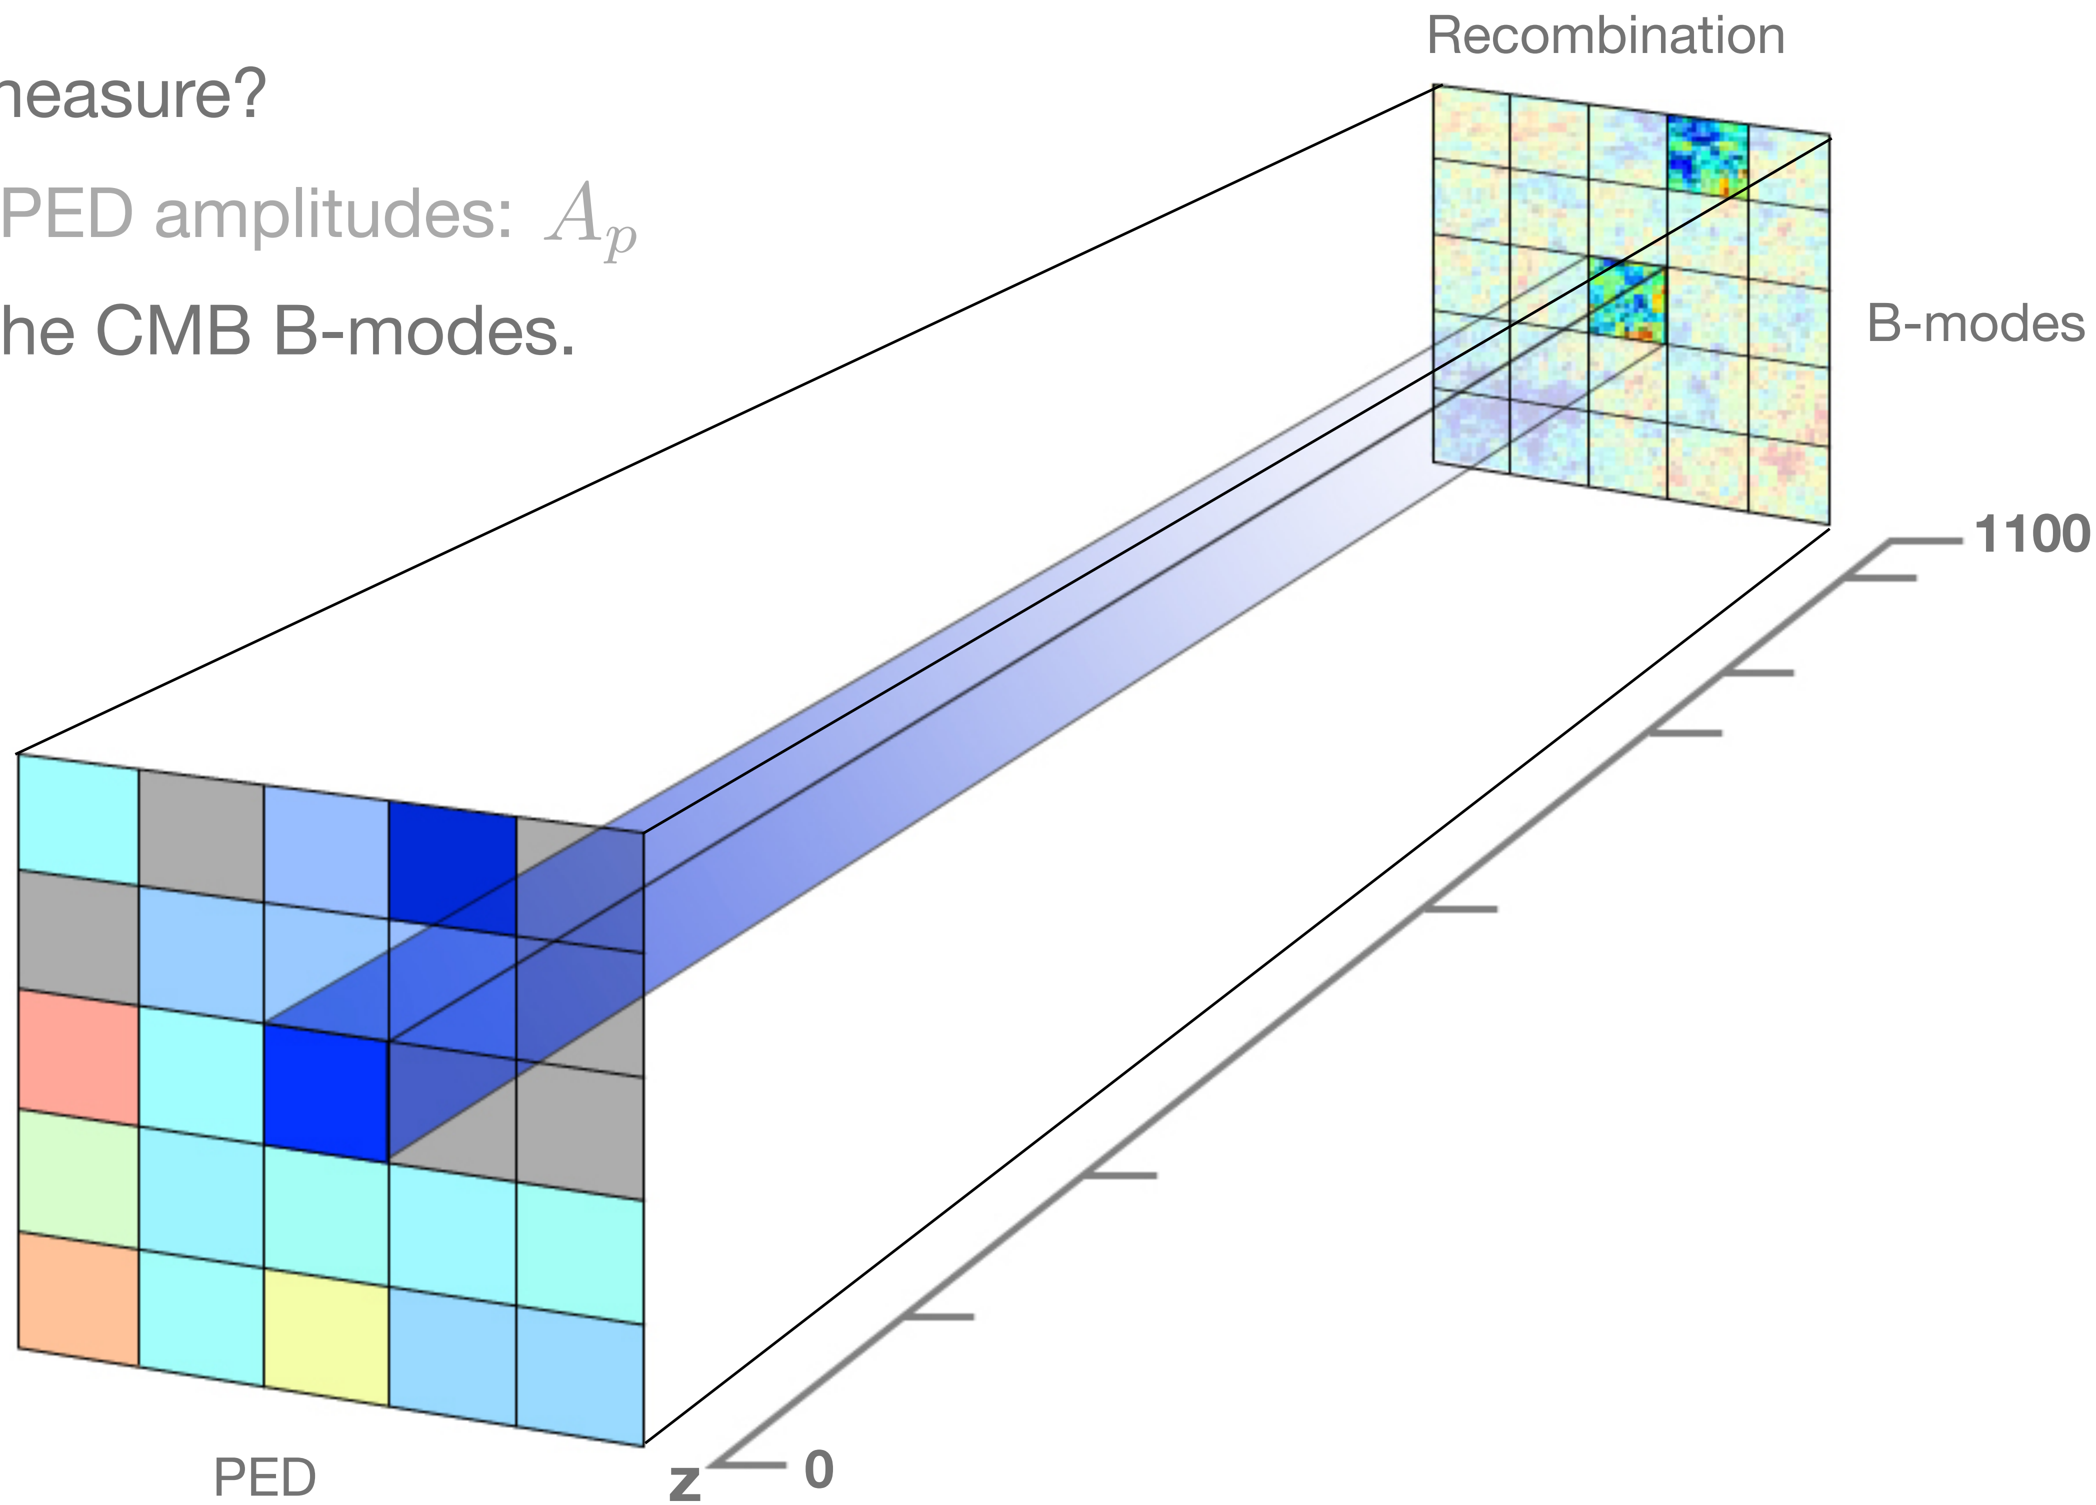

B-mode Surveys: Bandit Simulations

What do we measure?

- Initially: the PED amplitudes: A_p
- Ultimately: the CMB B-modes.

Figure of merit: σ^{r}

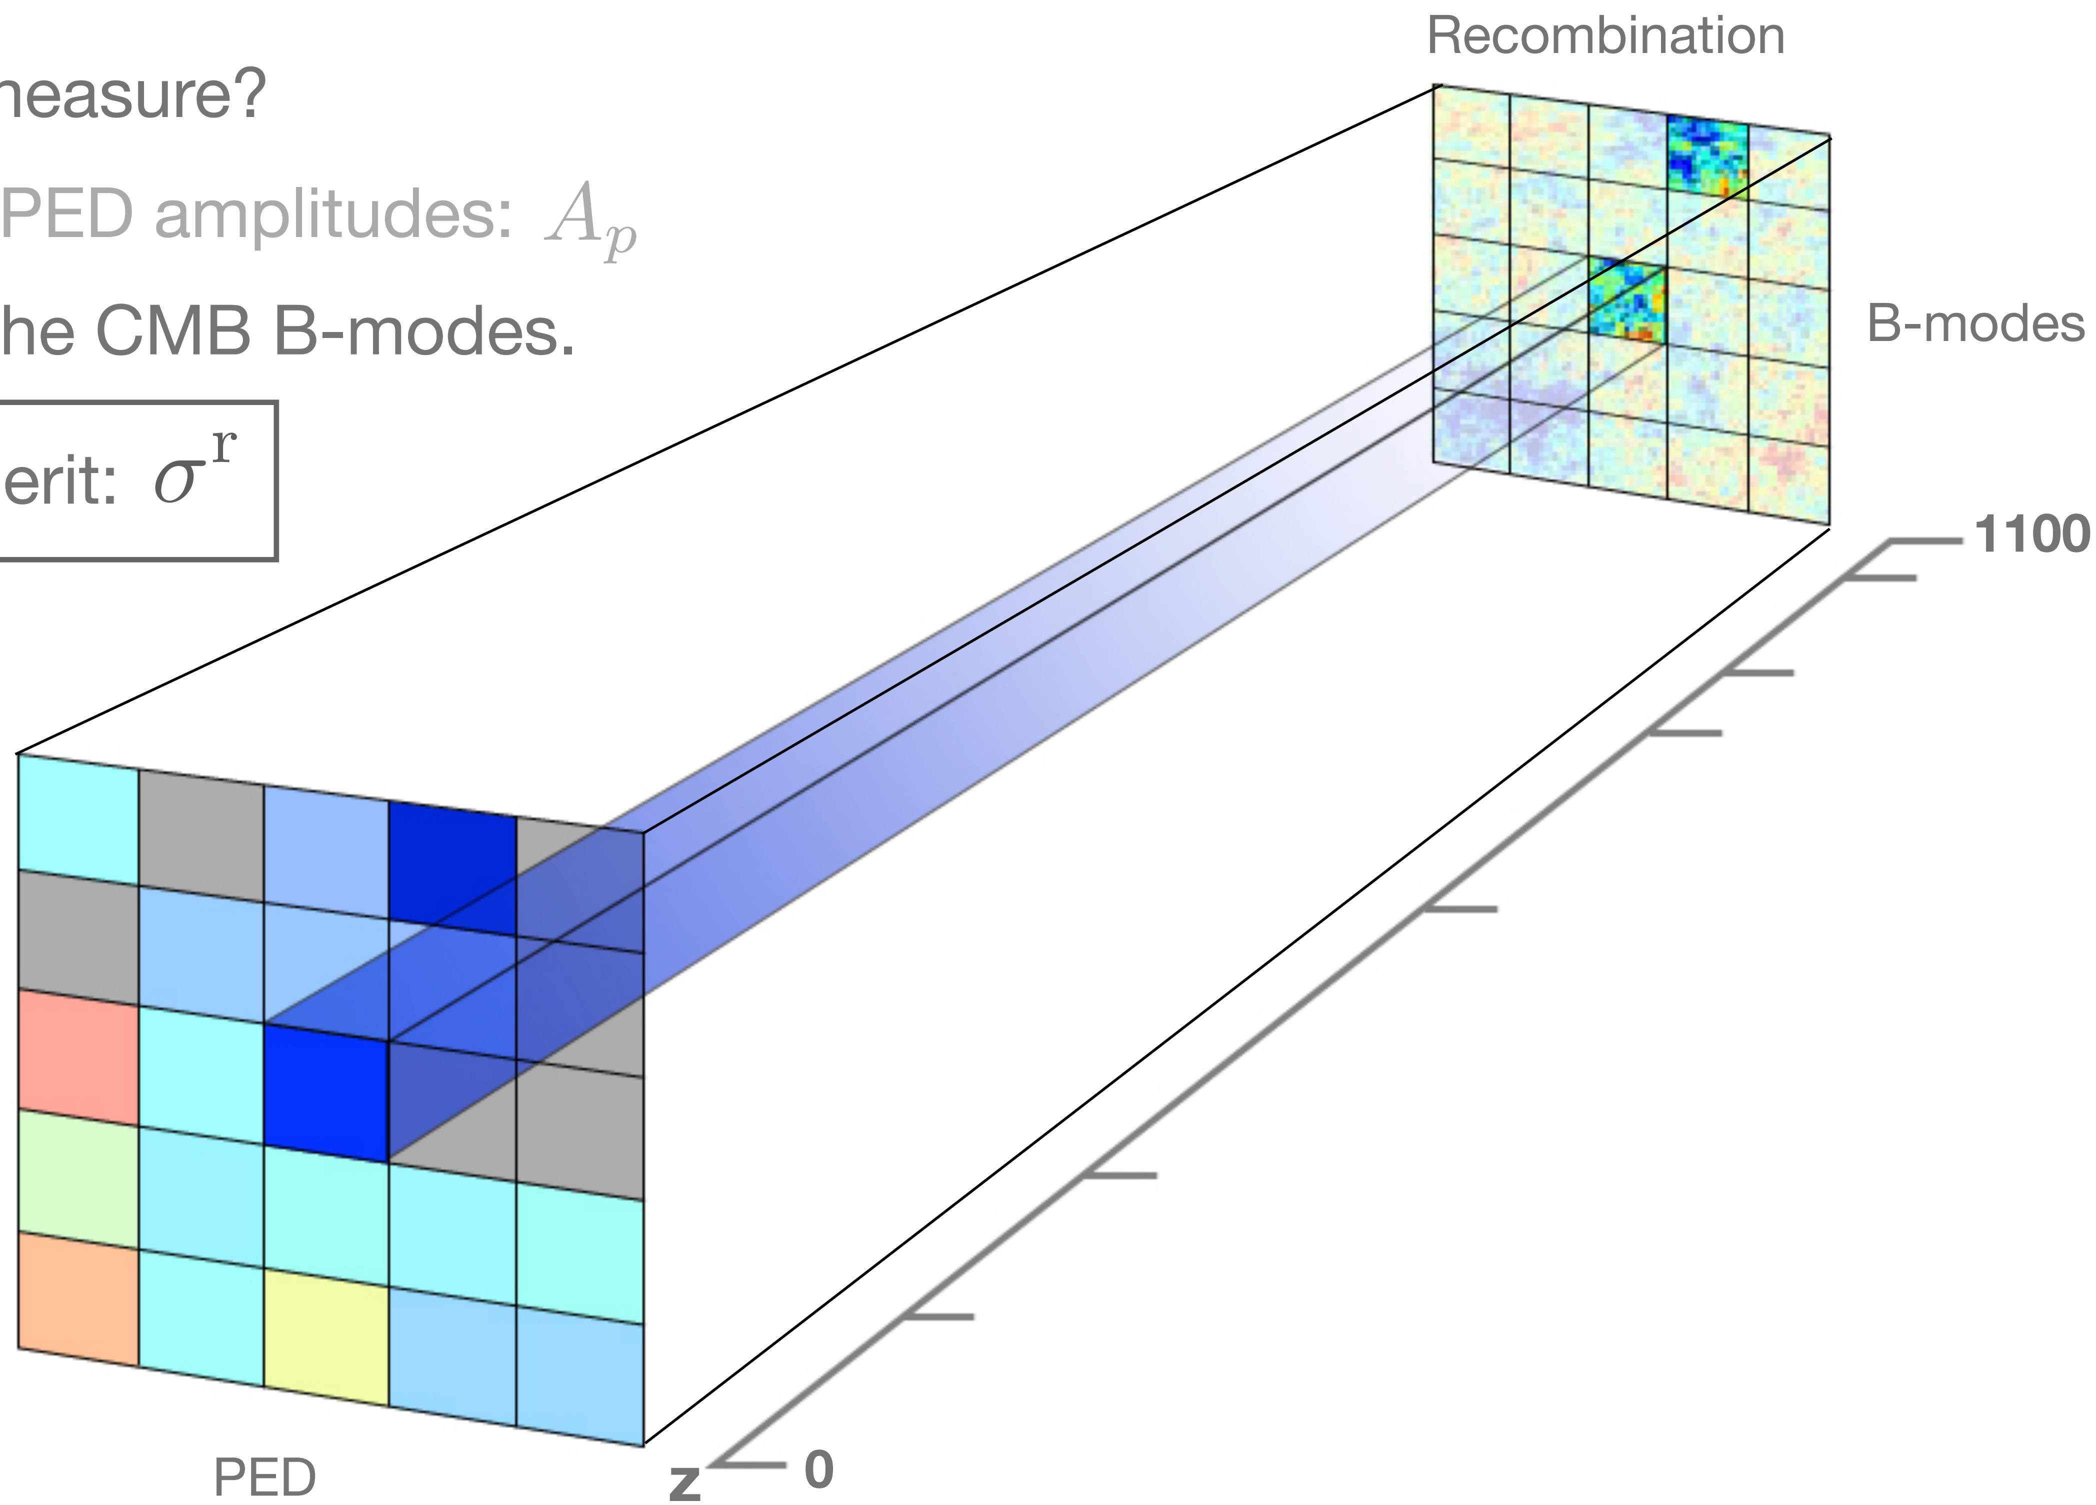

Results

Results

- In a *pessimistic* scenario - high dust fraction, no de-lensing:

Results

- In a *pessimistic* scenario - high dust fraction, no de-lensing:

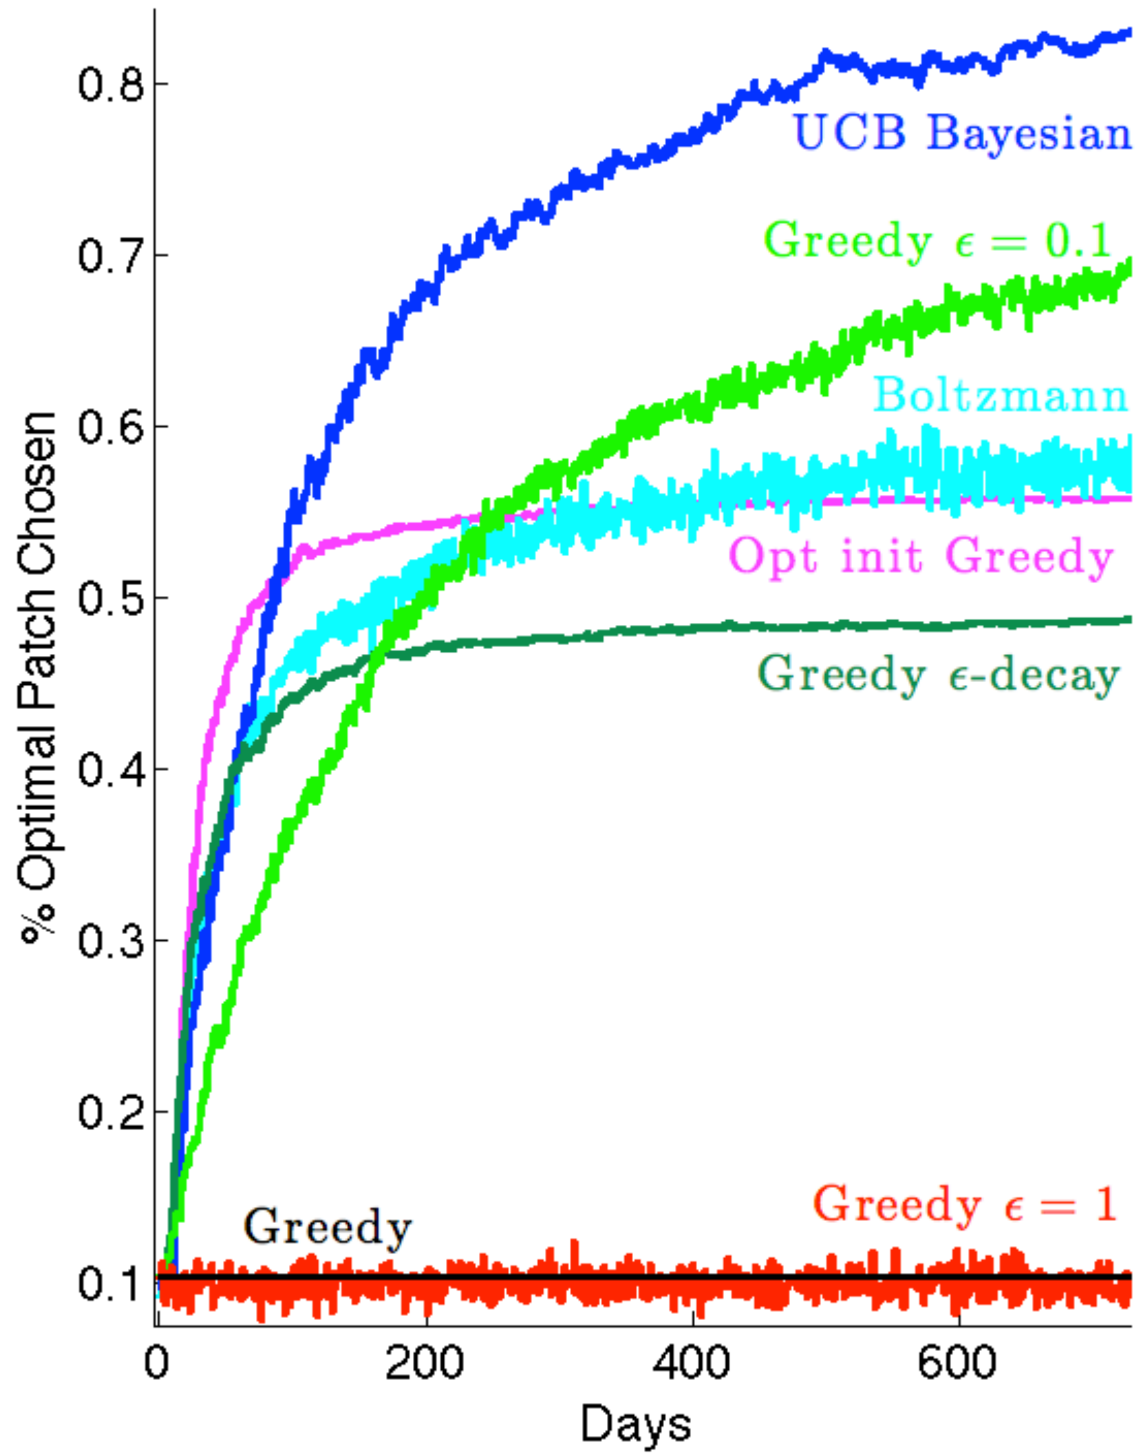

Results

Results

- In an *optimistic* scenario - low dust fraction, 80% de-lensing:

Results

- In an *optimistic* scenario - low dust fraction, 80% de-lensing:

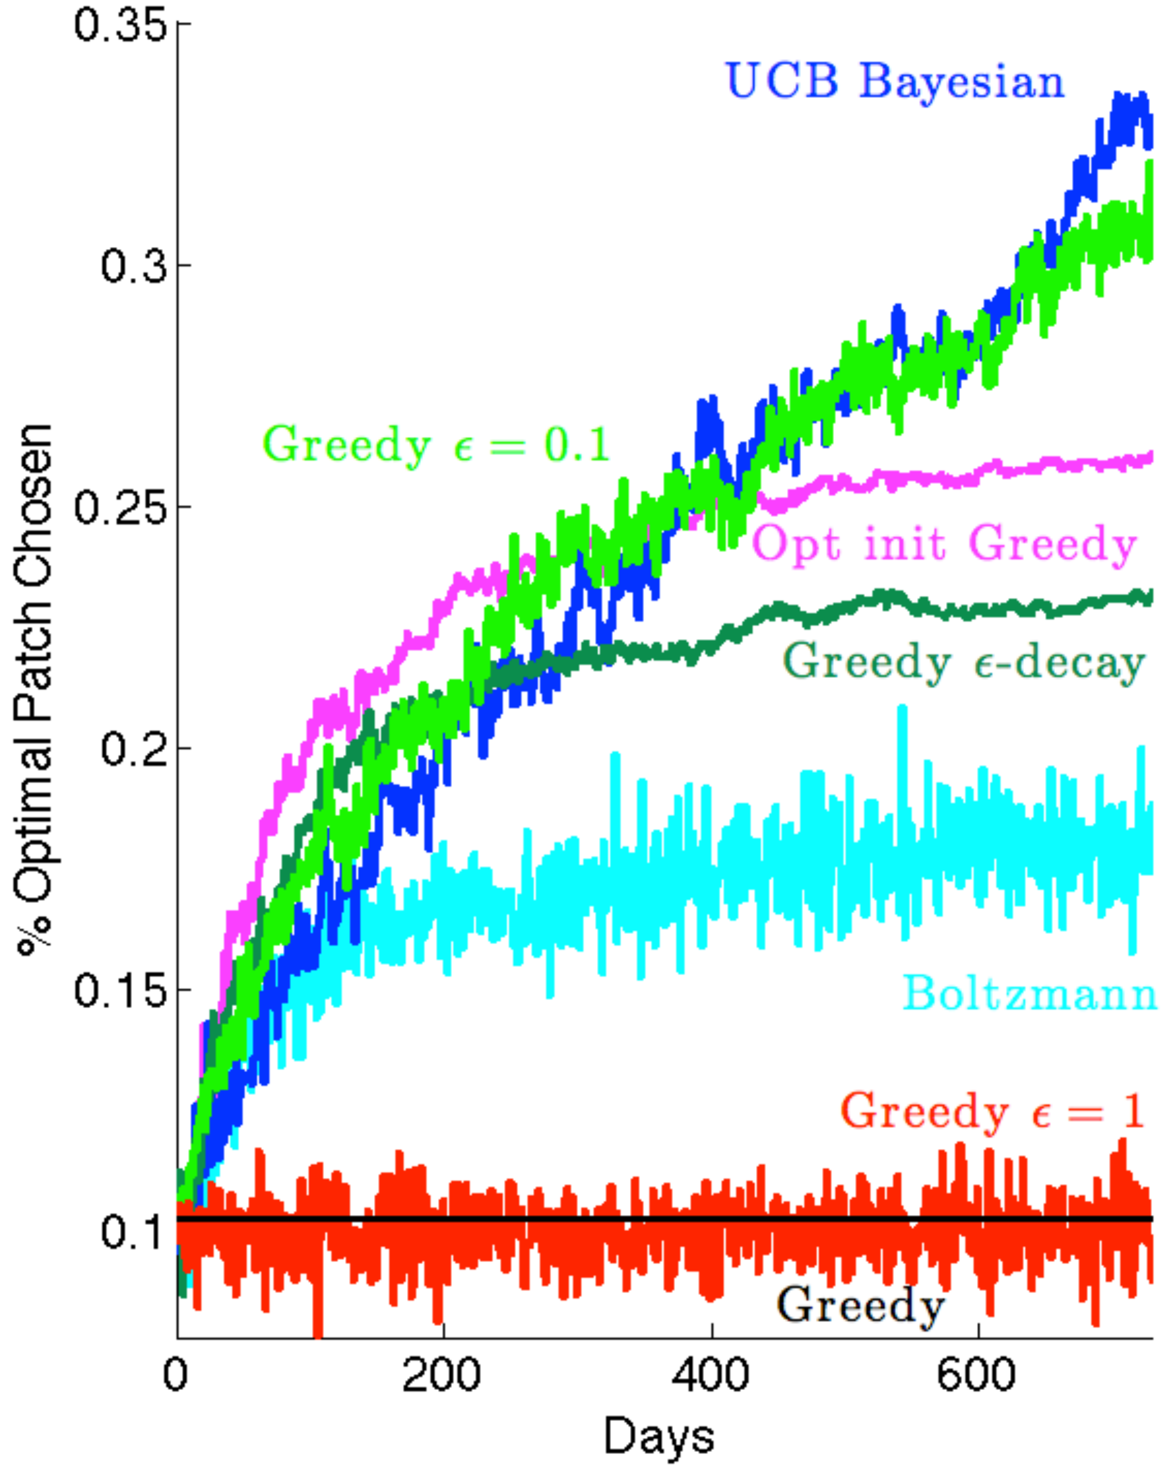

Results: Upper Bound on Tensor-to-Scalar Ratio

Results: Upper Bound on Tensor-to-Scalar Ratio

Pessimistic

Optimistic

Results: Upper Bound on Tensor-to-Scalar Ratio

Pessimistic

Optimistic

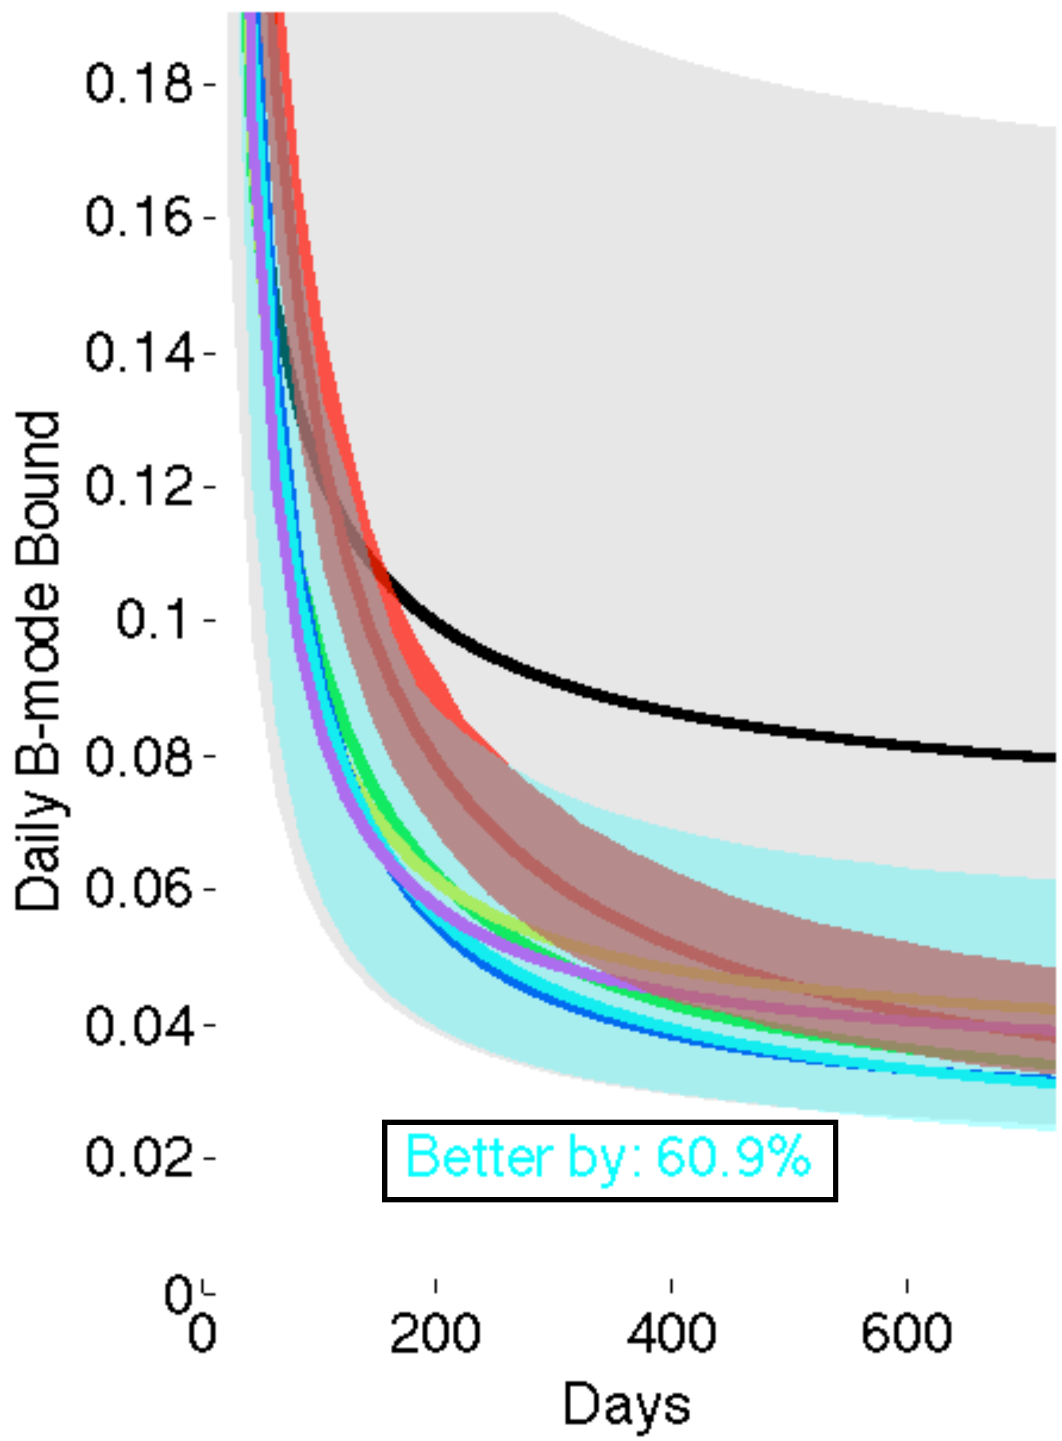

Results: Upper Bound on Tensor-to-Scalar Ratio

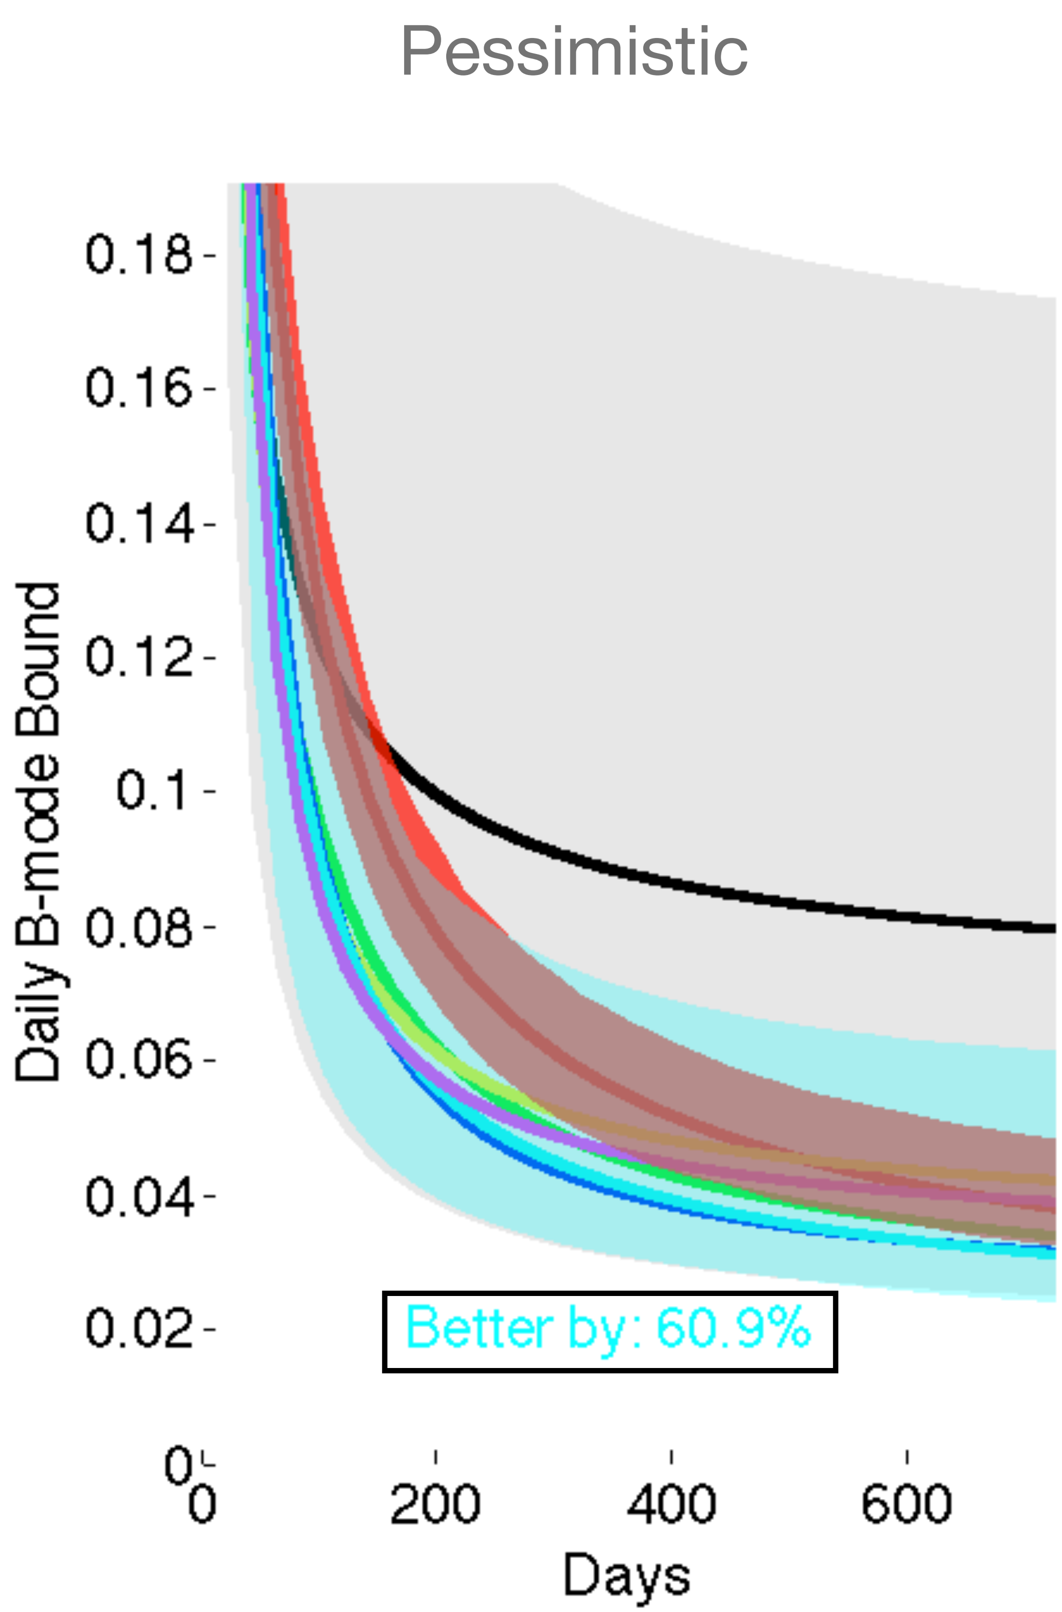

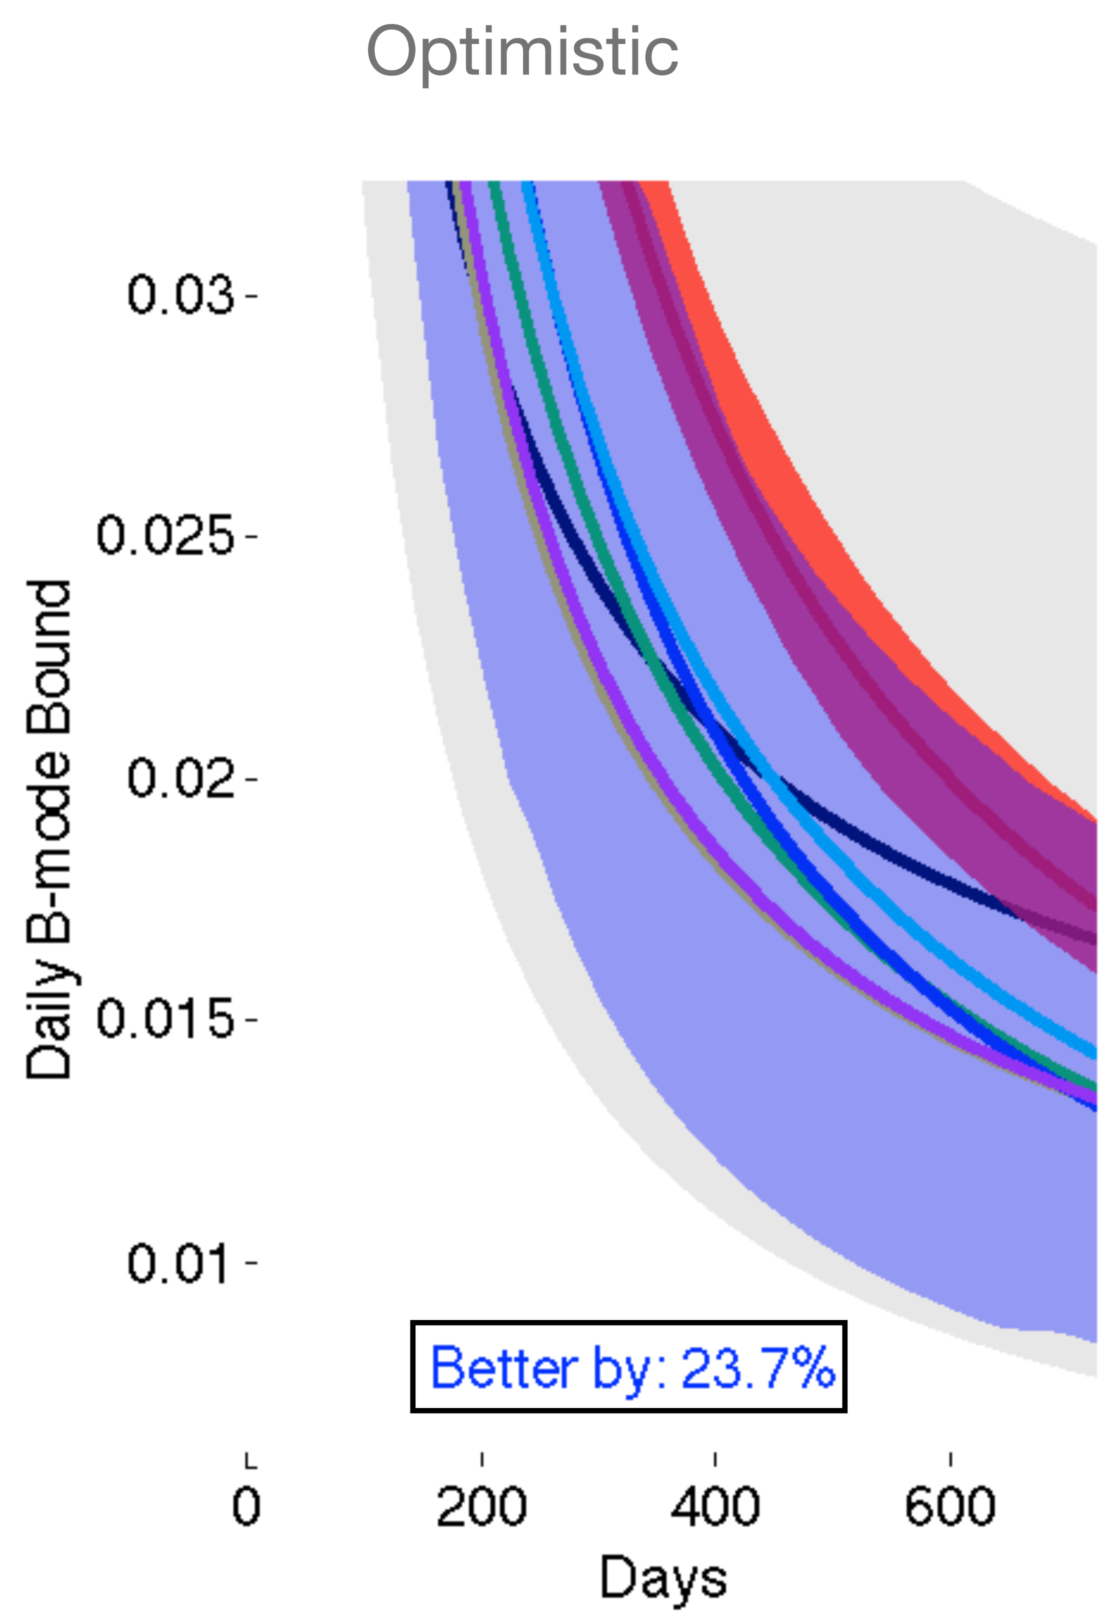

Outline

Outline

- The Multi-Armed-Bandit Problem
- Heuristic Solution Algorithms
- Mitigating CMB B-mode Foregrounds
- MAB Strategies Elsewhere

MAB Strategies Elsewhere

MAB Strategies Elsewhere

- 21-cm stochastic fluctuations.

MAB Strategies Elsewhere

- 21-cm stochastic fluctuations.

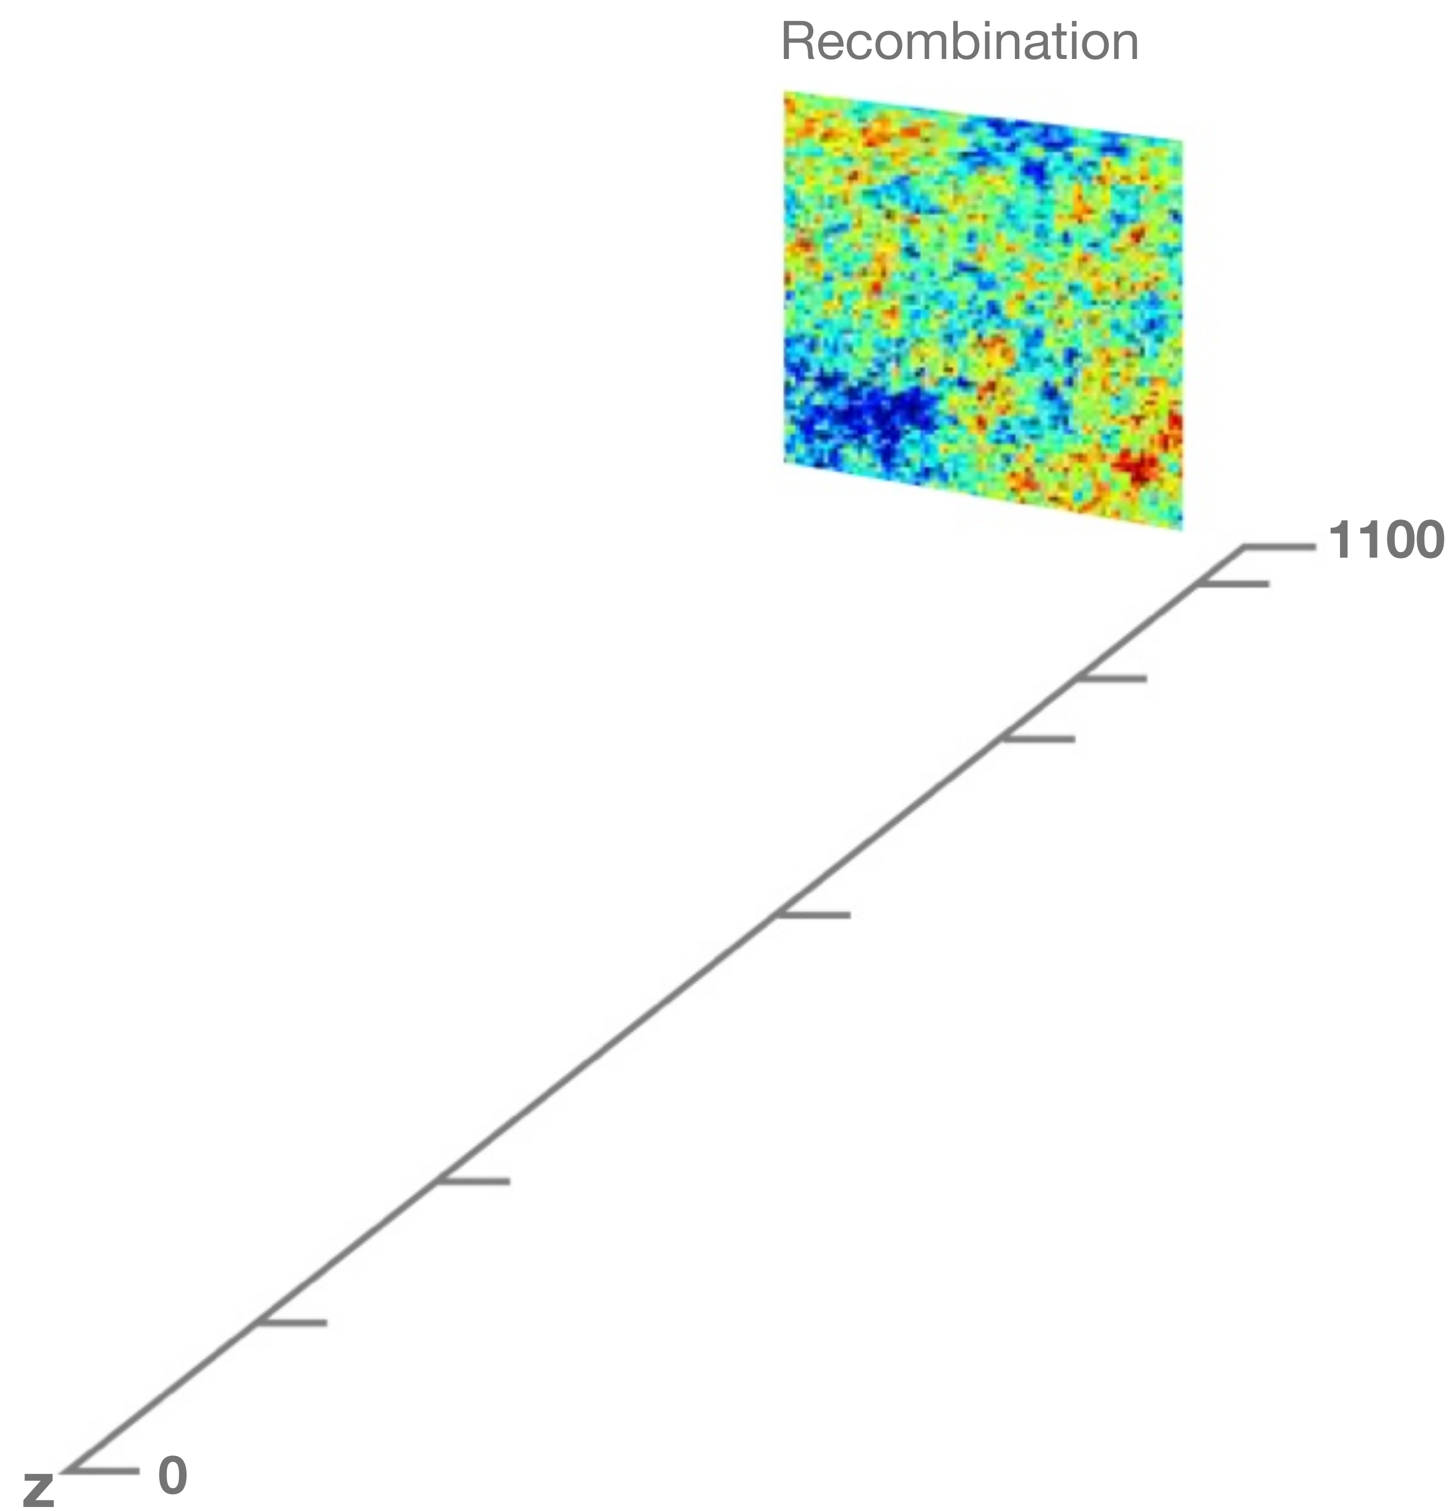

MAB Strategies Elsewhere

- 21-cm stochastic fluctuations.

  A 3D-bandit problem.

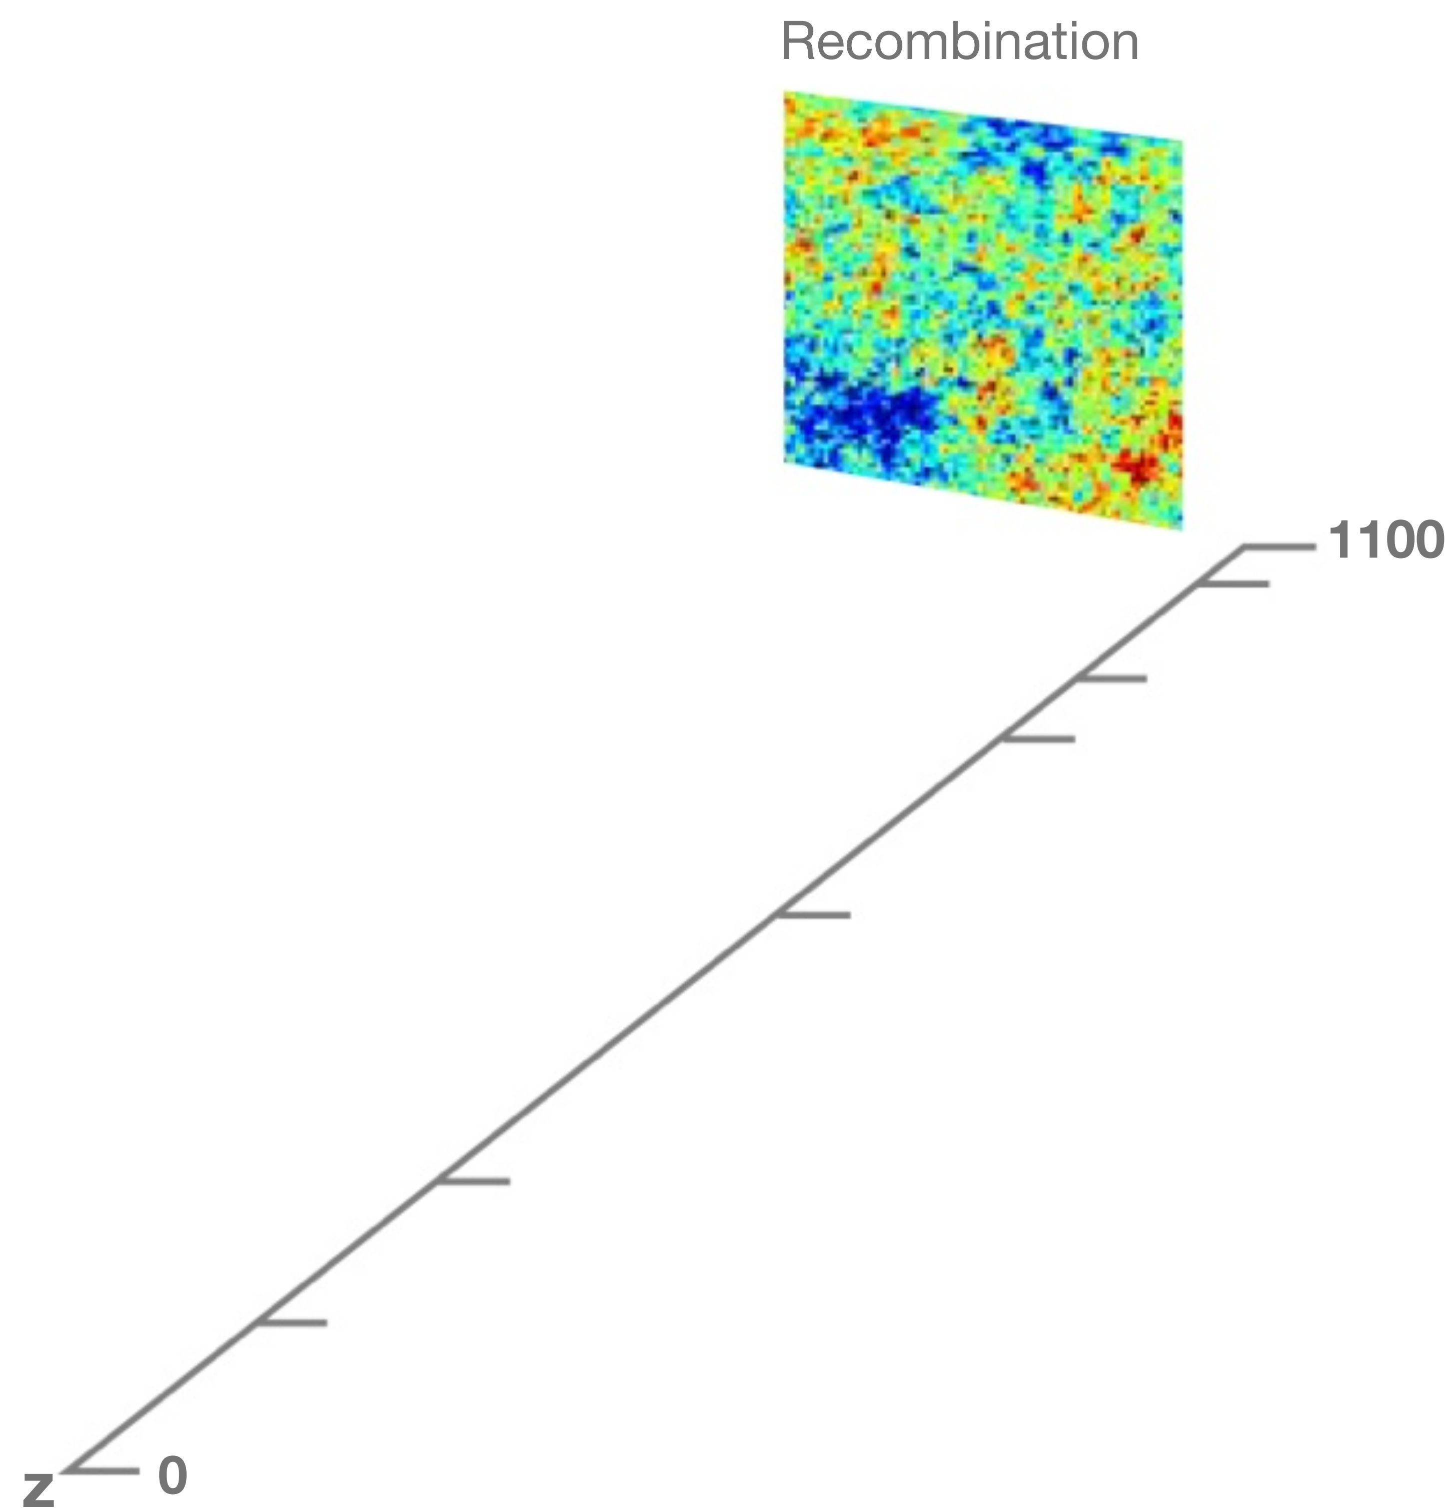

MAB Strategies Elsewhere

- 21-cm stochastic fluctuations.

  A 3D-bandit problem.

Radio Interferometer

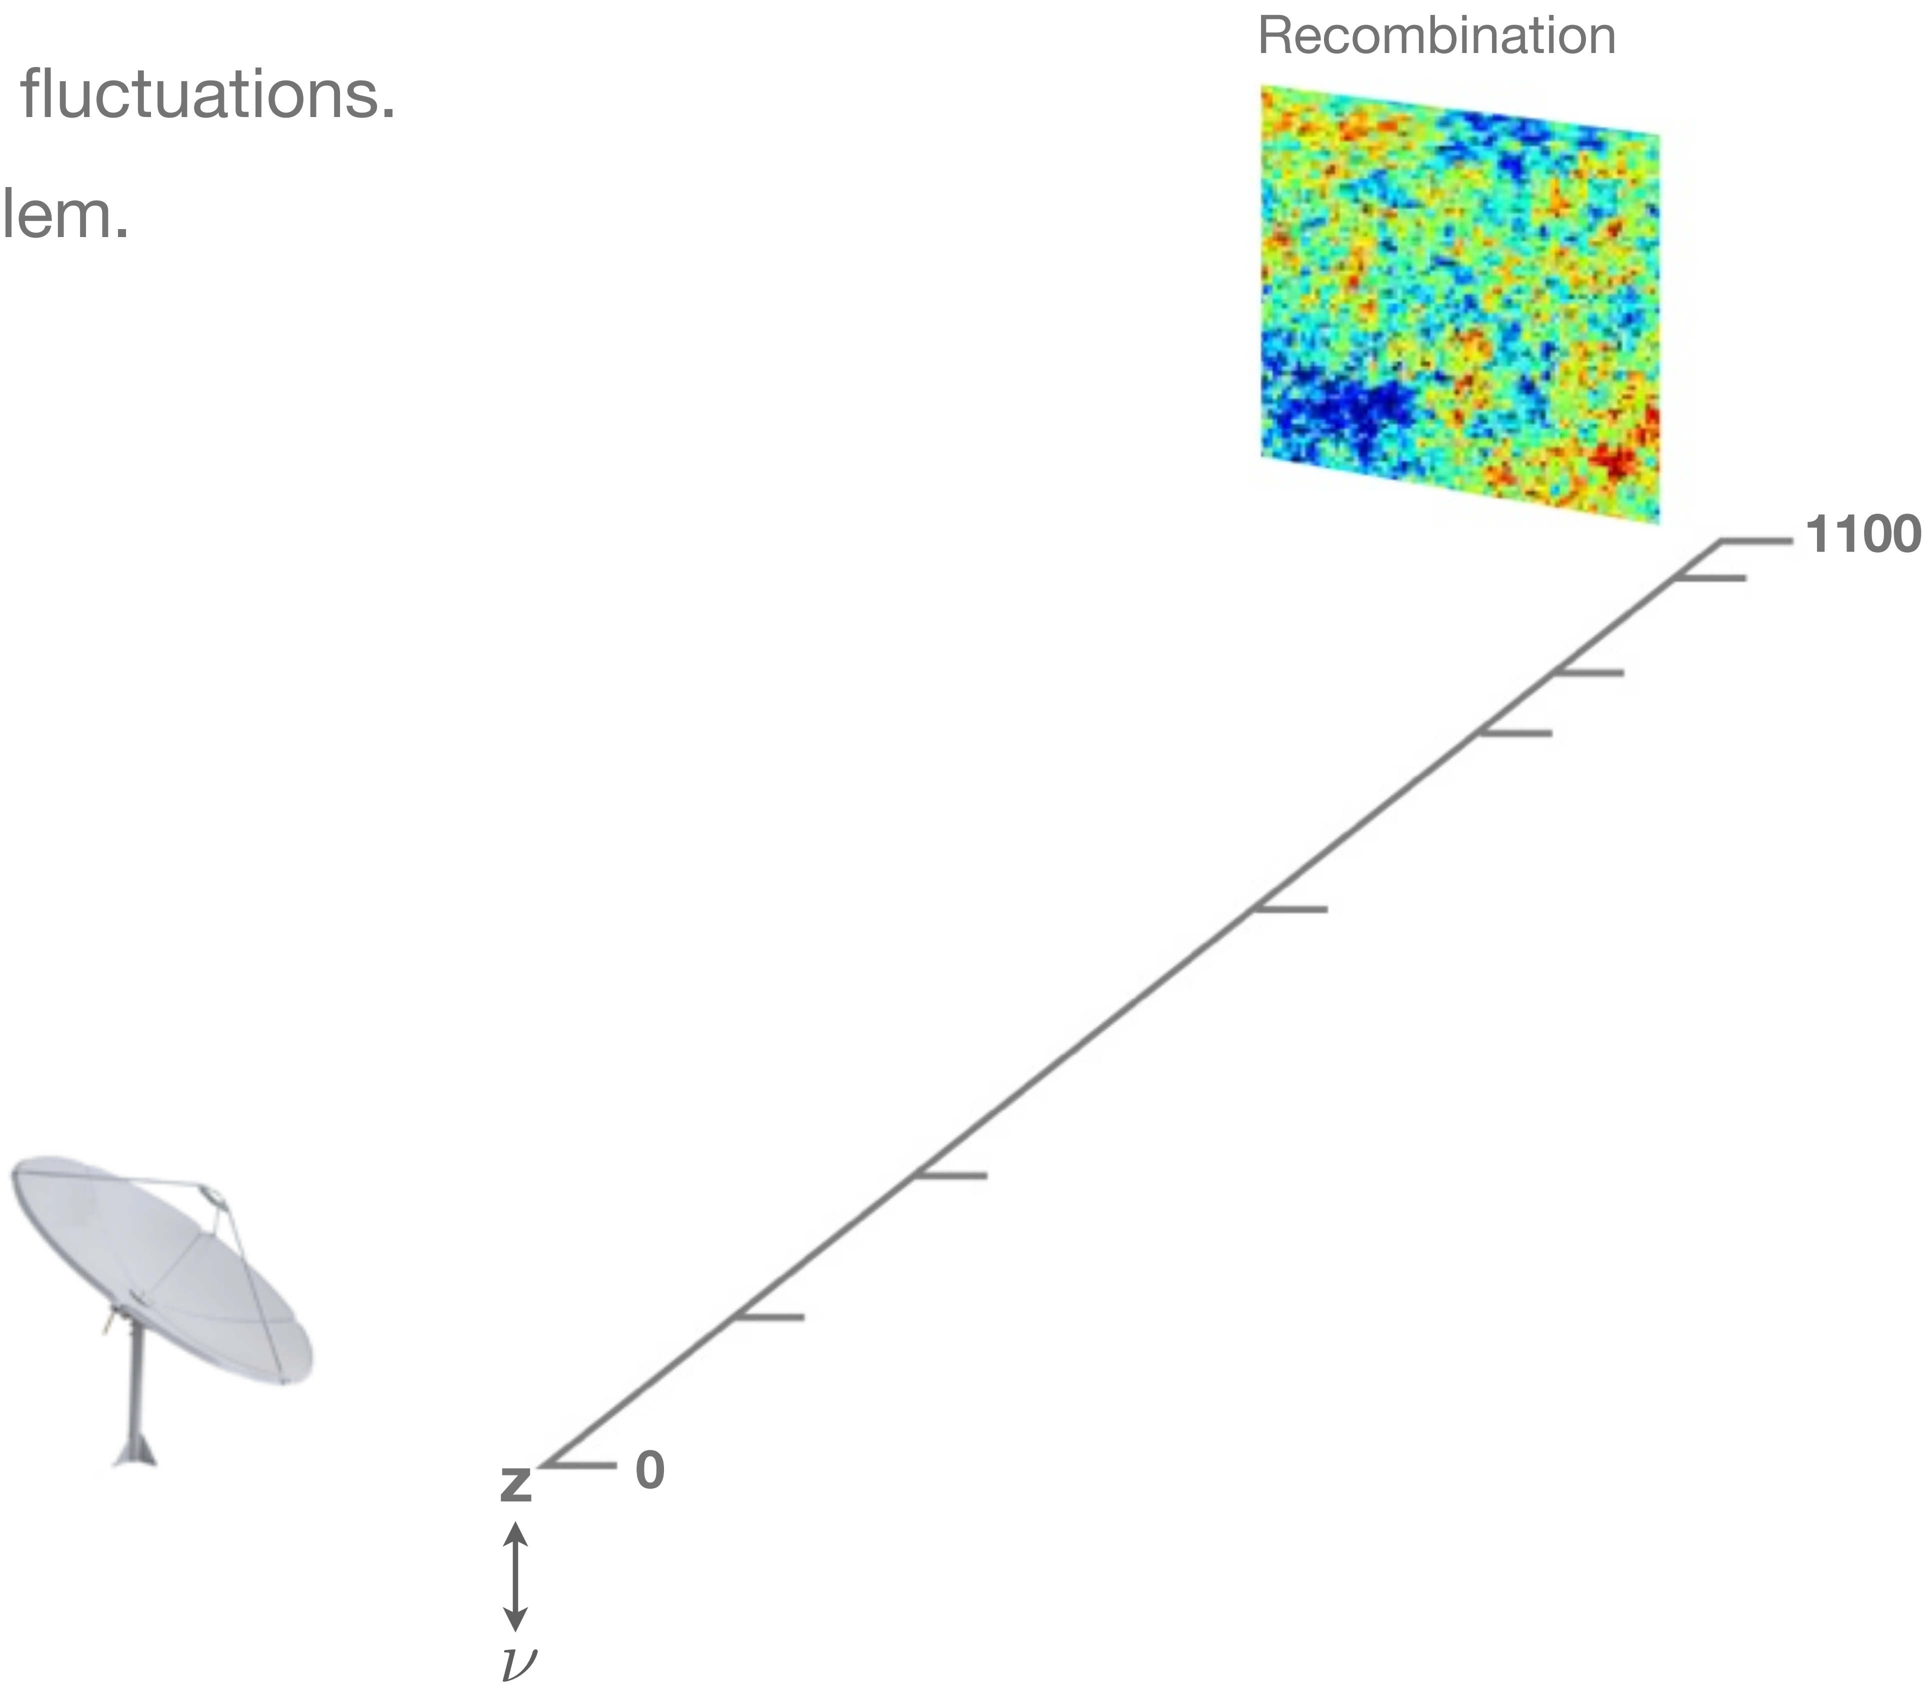

MAB Strategies Elsewhere

- 21-cm stochastic fluctuations.

  A 3D-bandit problem.

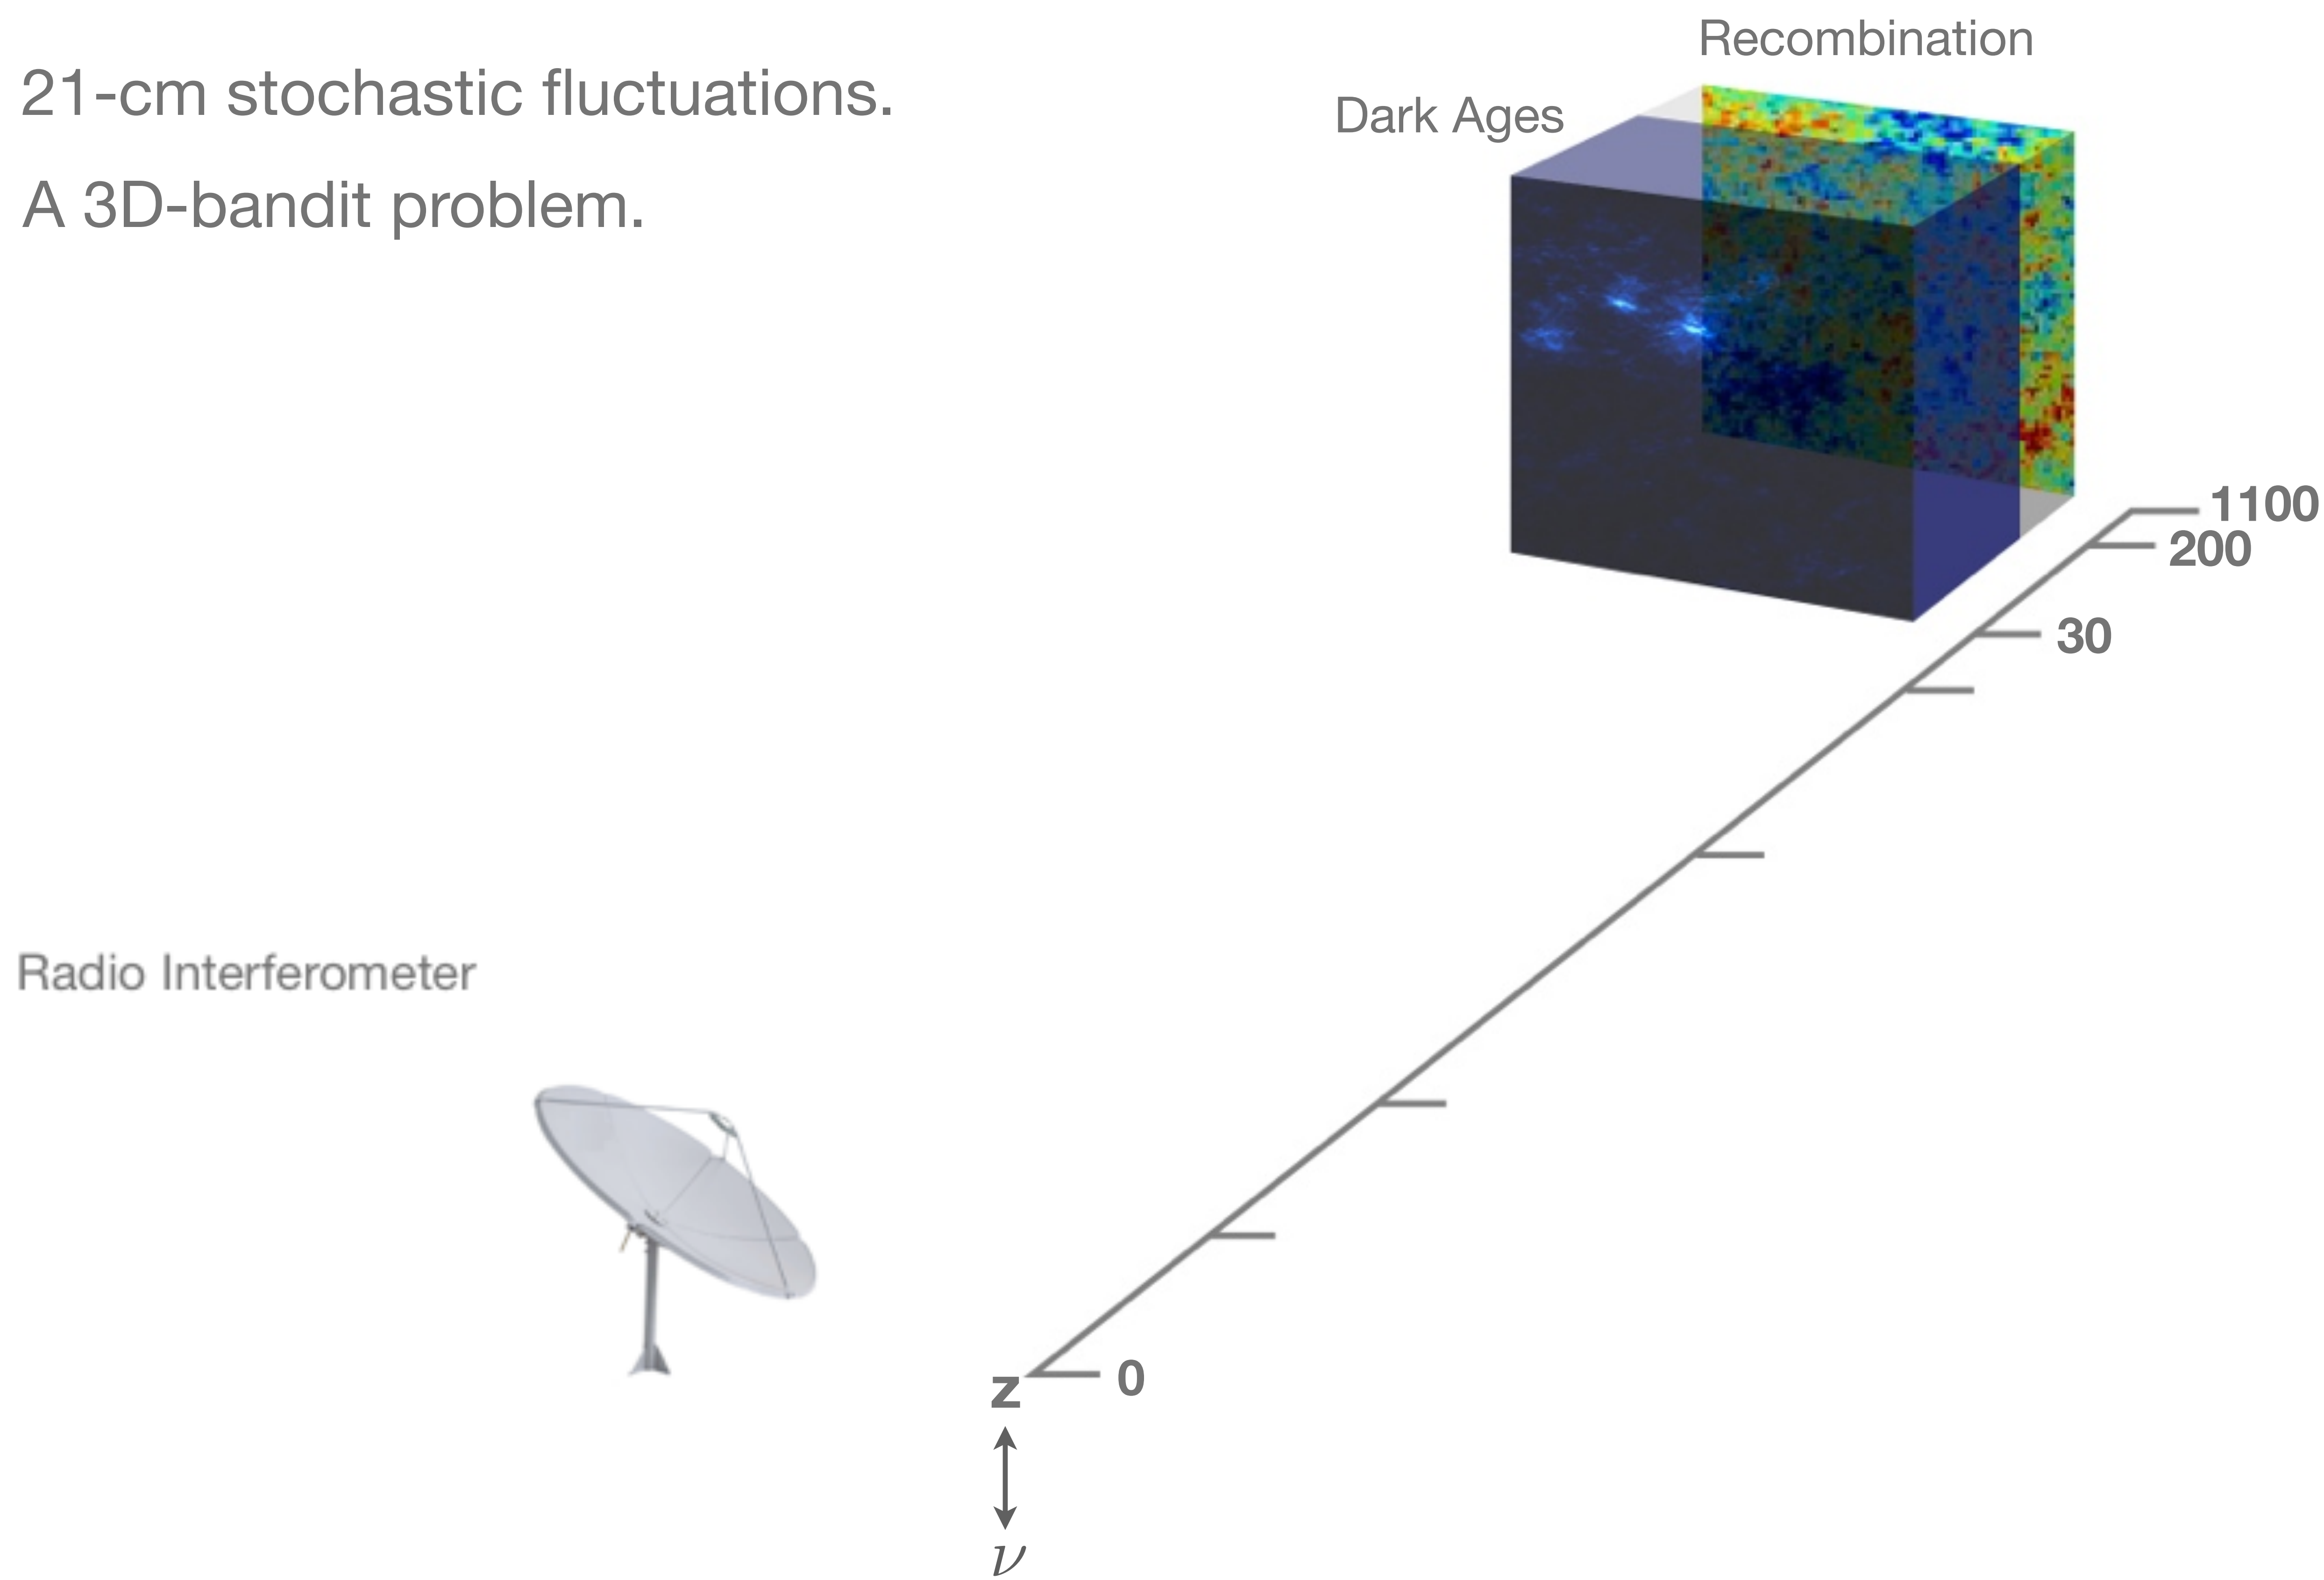

MAB Strategies Elsewhere

- 21-cm stochastic fluctuations.

  A 3D-bandit problem.

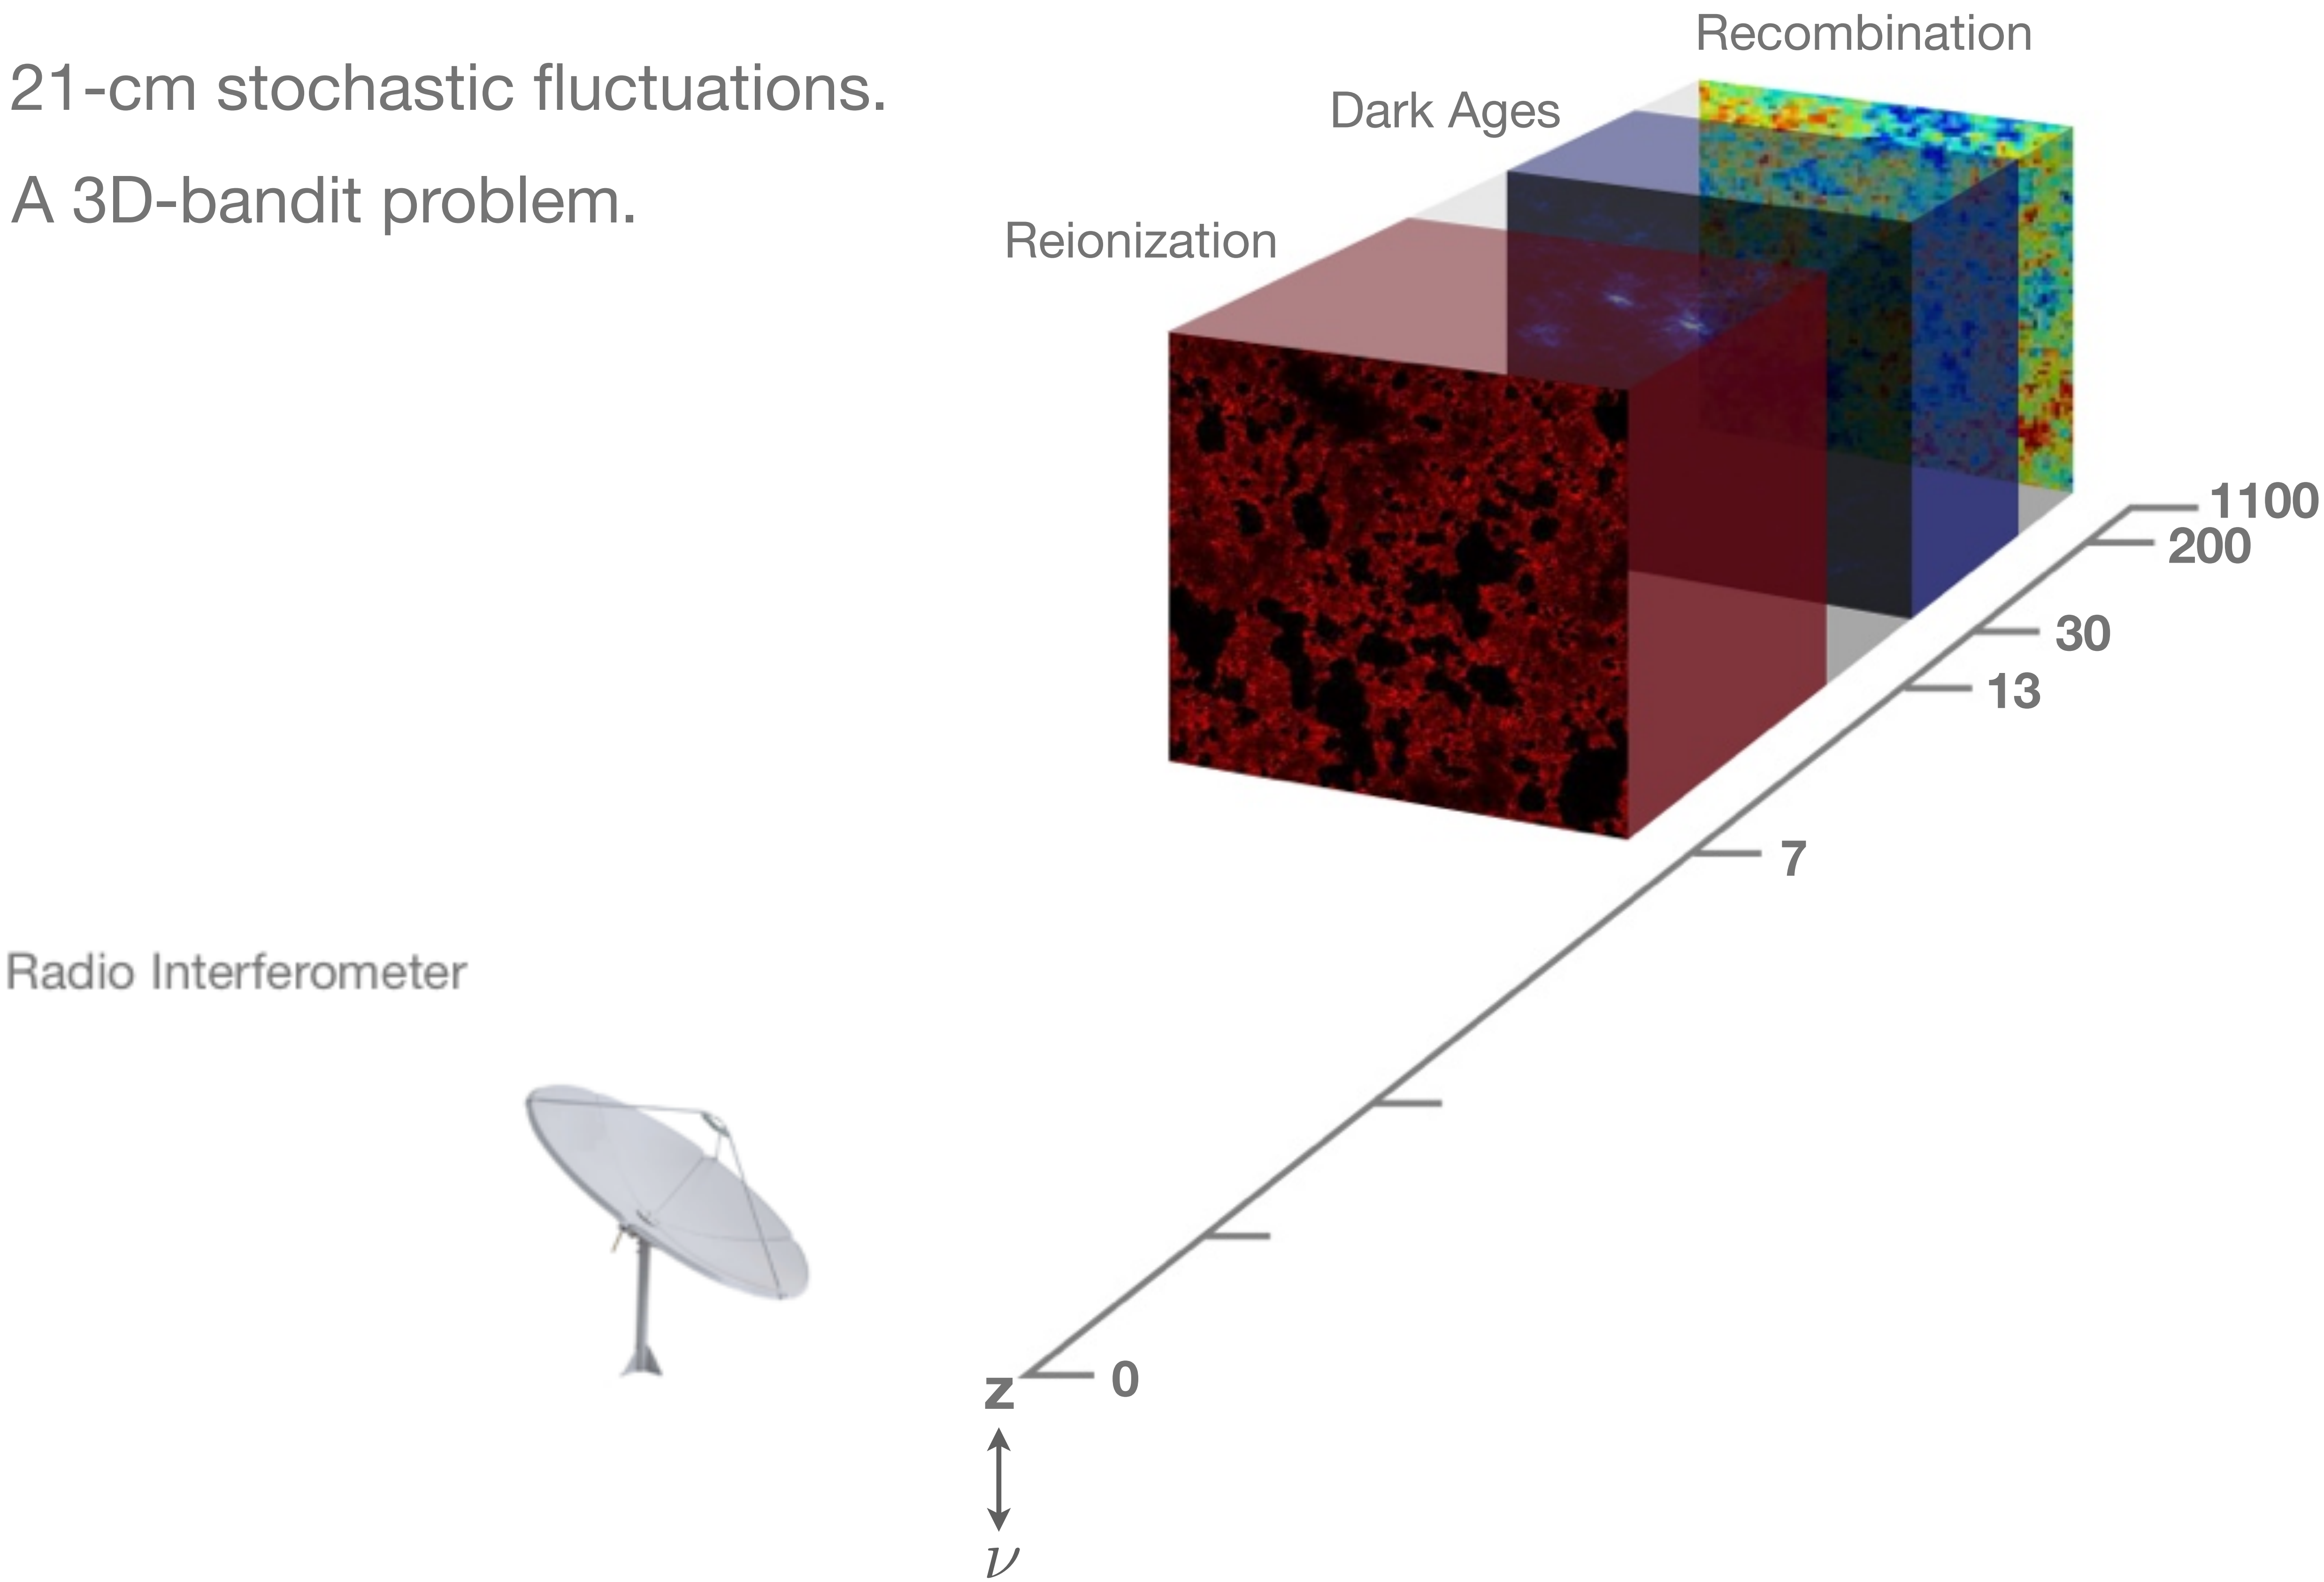

MAB Strategies Elsewhere

- 21-cm stochastic fluctuations.

  A 3D-bandit problem.

  Adds another dimension.

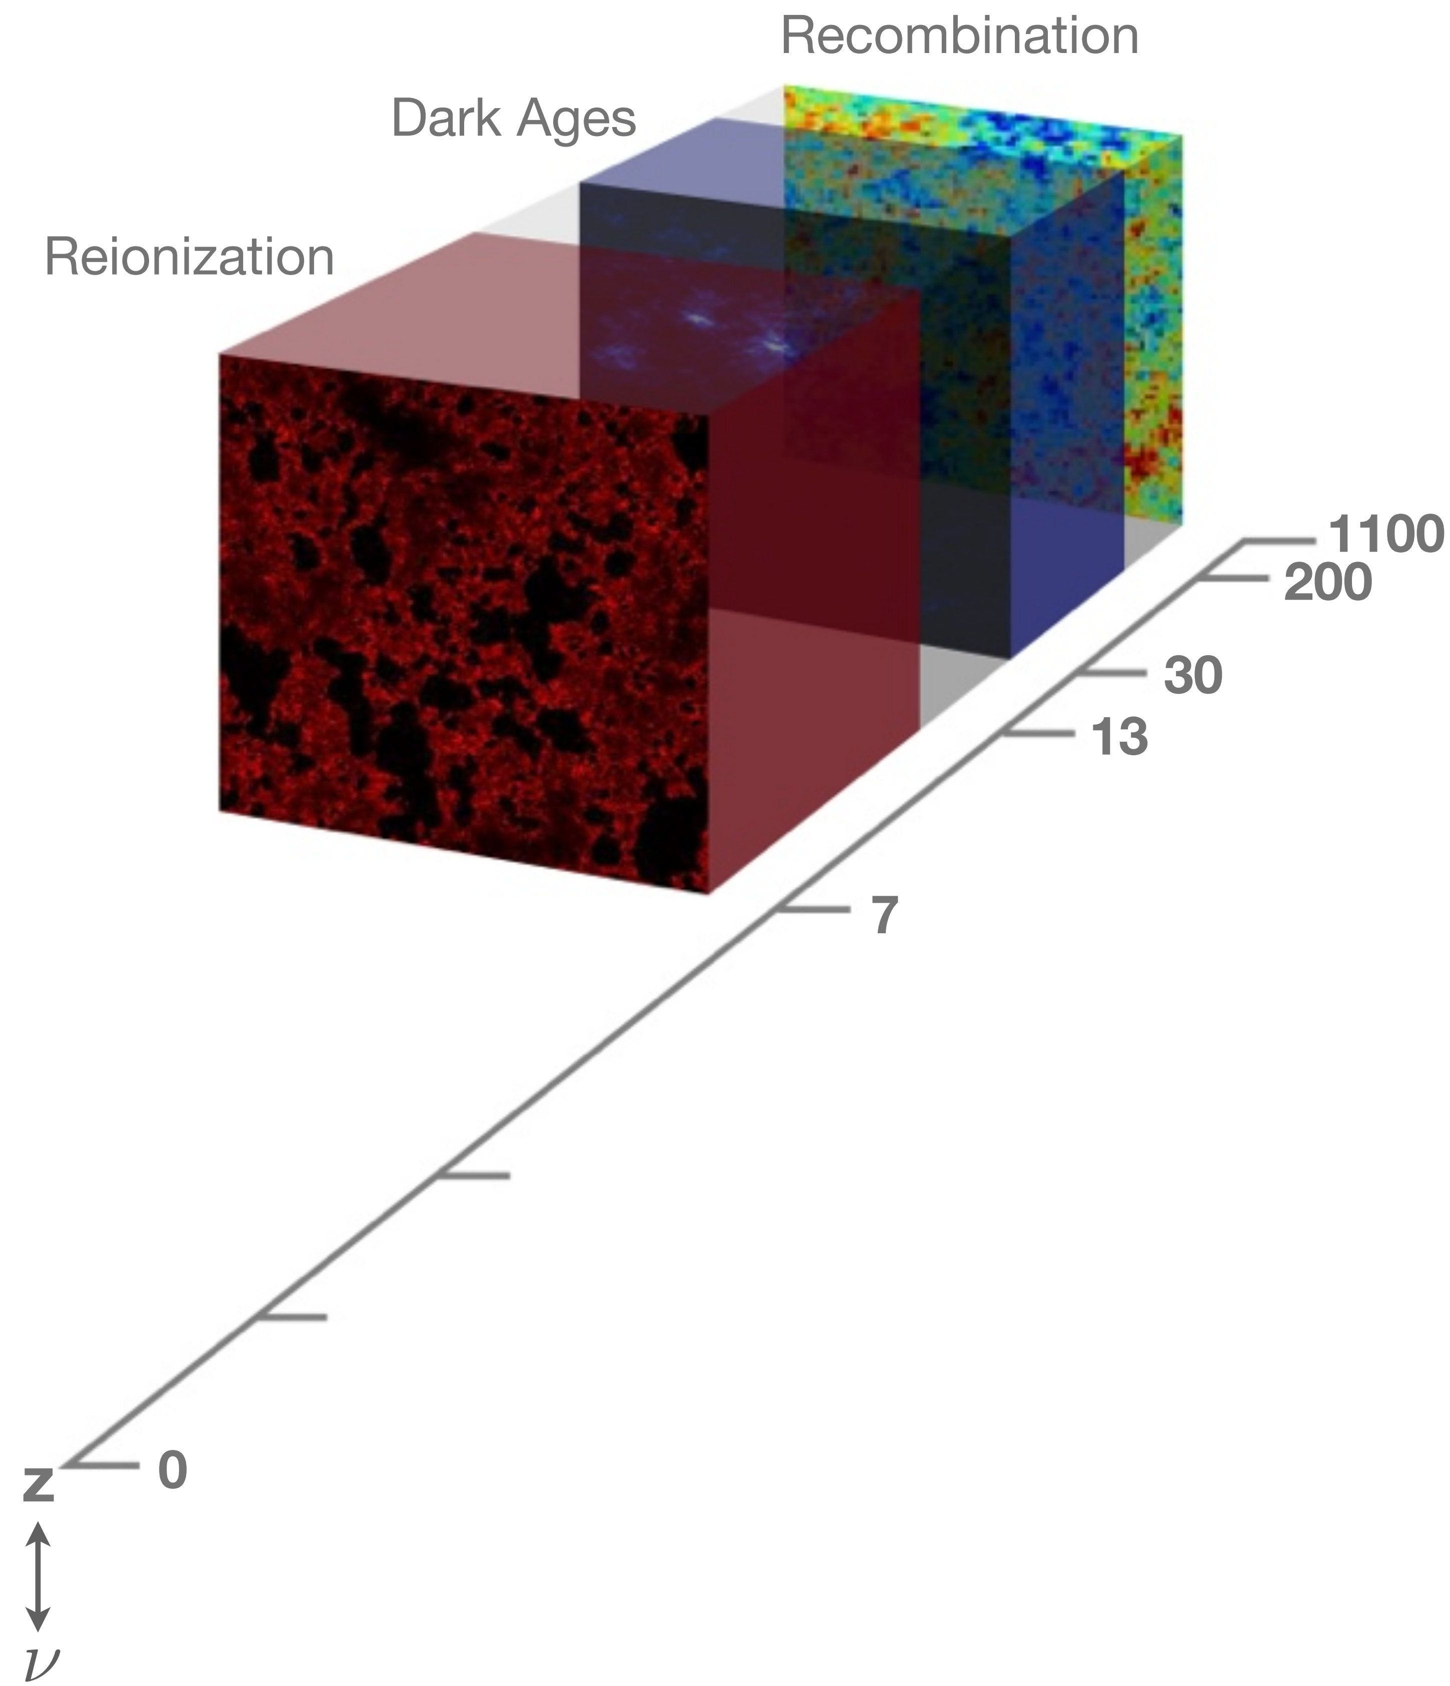

Radio Interferometer

MAB Strategies Elsewhere

- 21-cm stochastic fluctuations.

  A 3D-bandit problem.

  Adds another dimension.

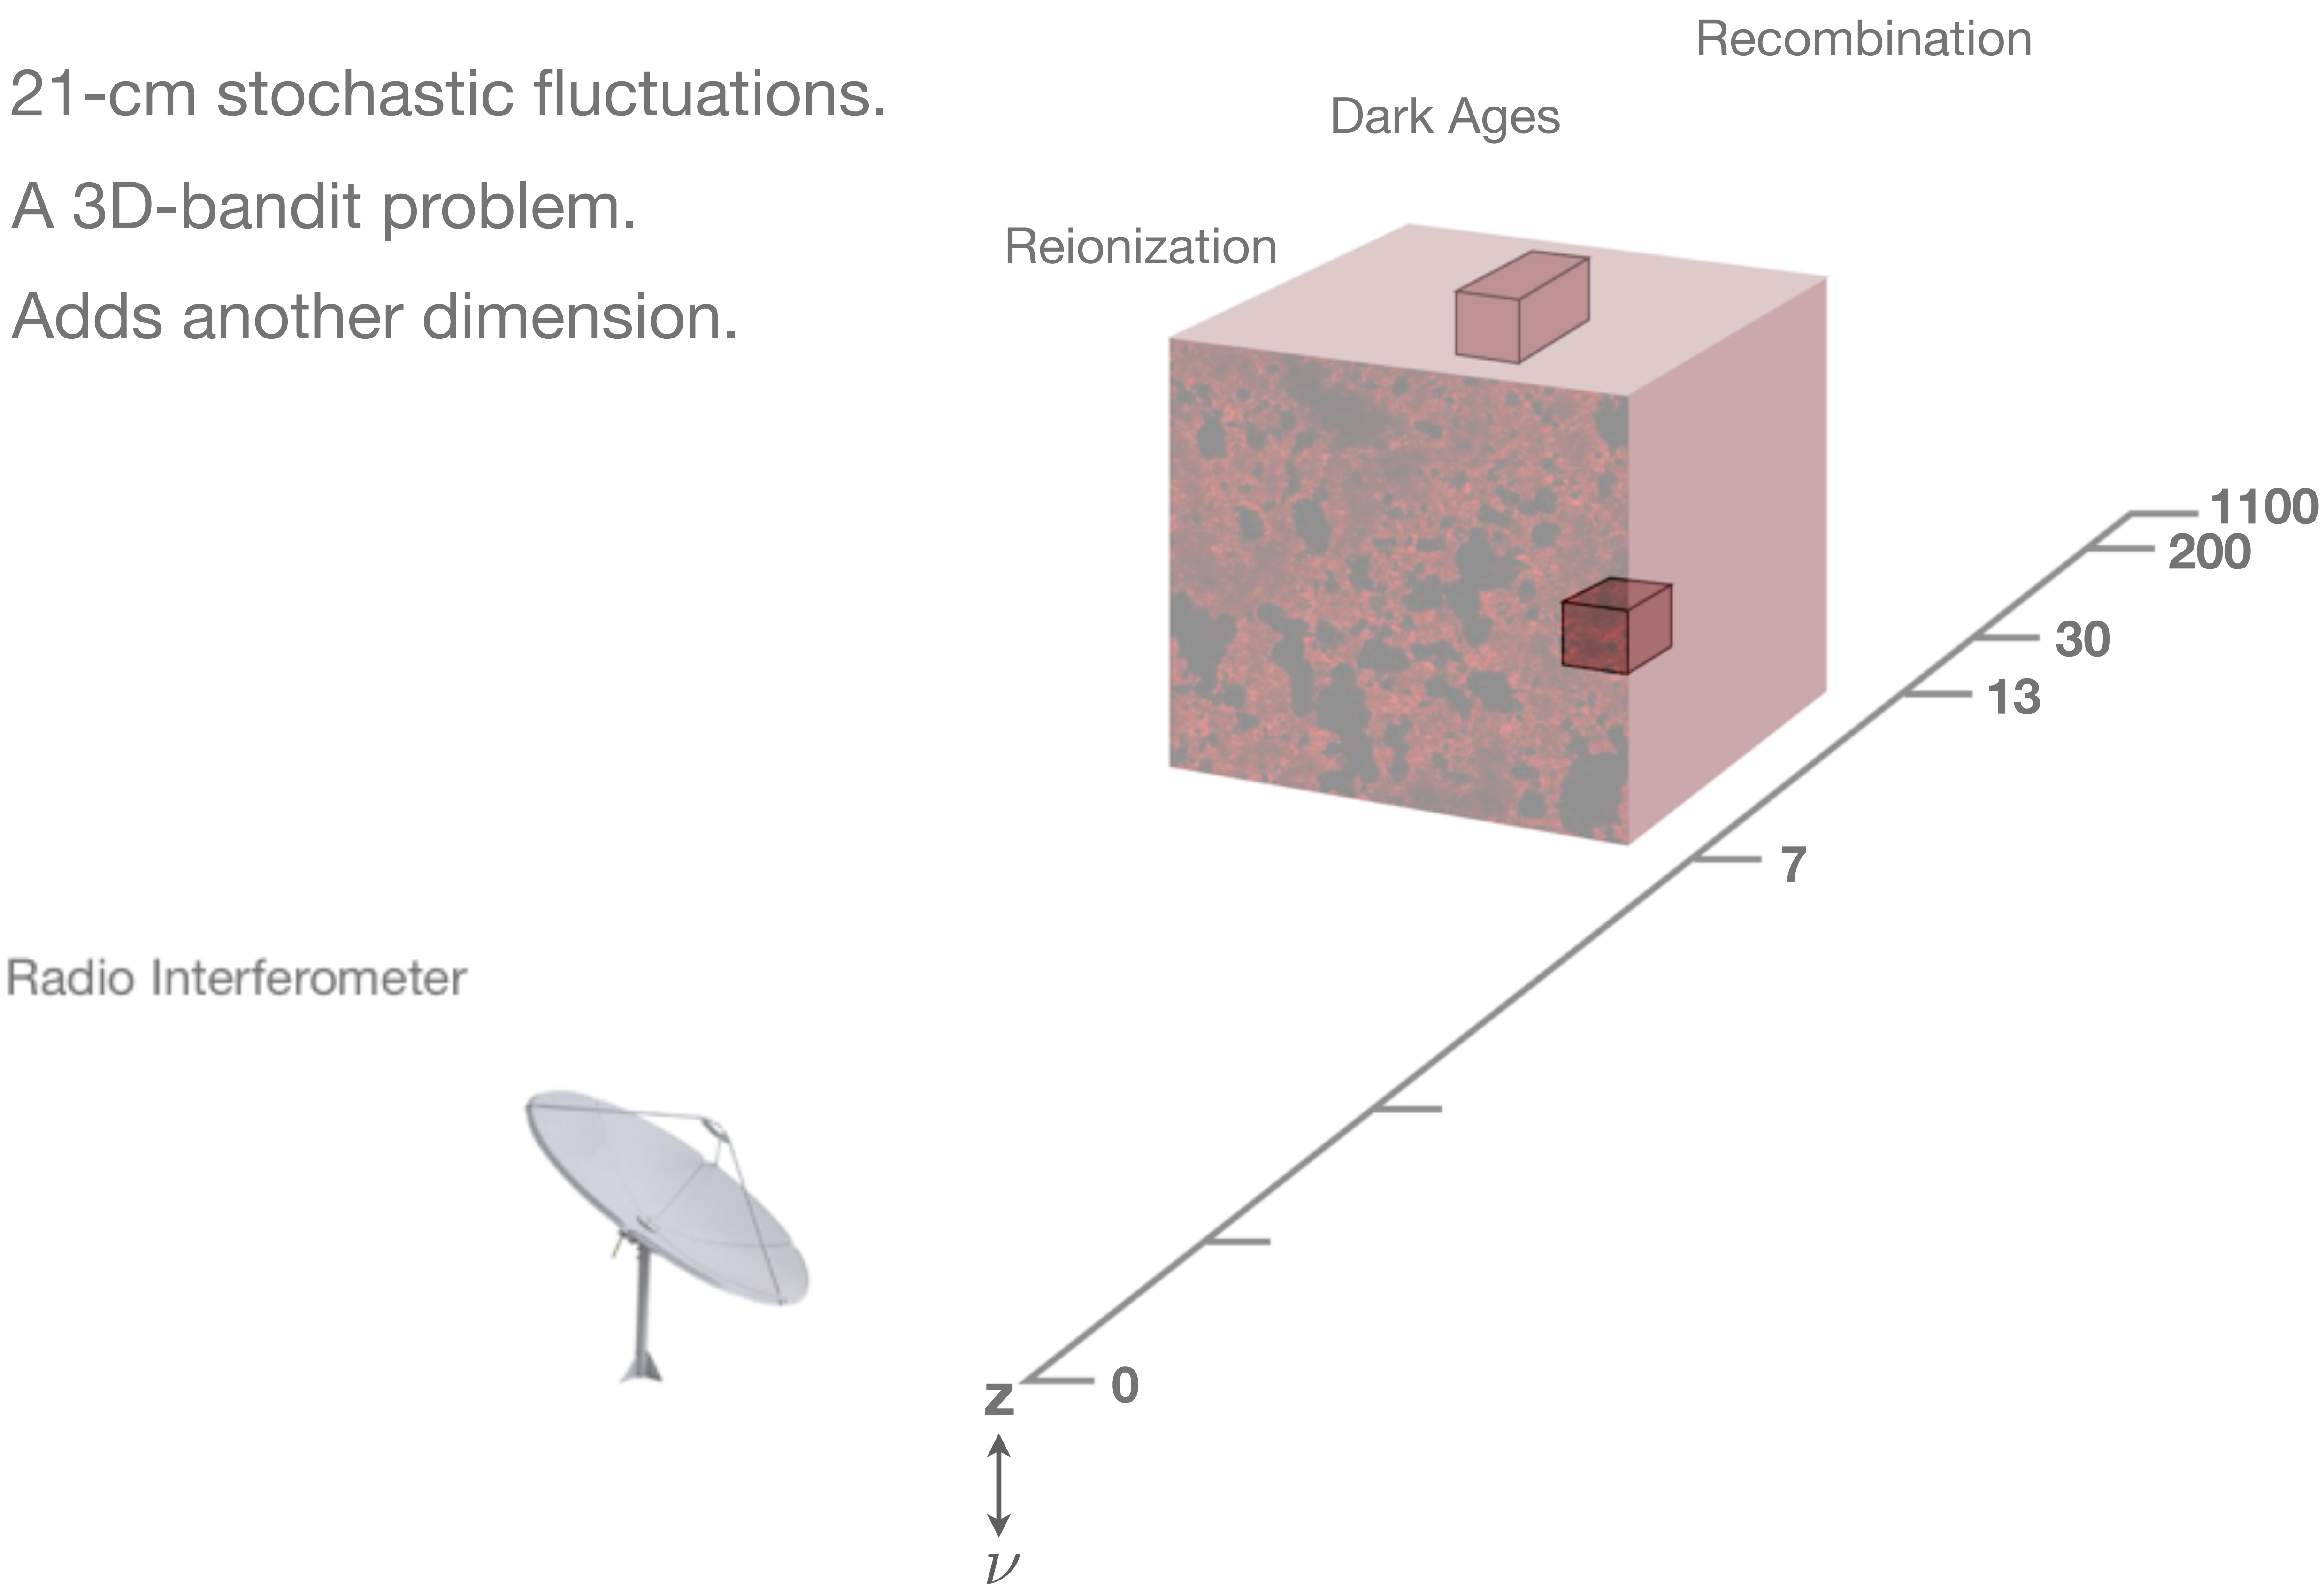

MAB Strategies Elsewhere

- 21-cm stochastic fluctuations.

  A 3D-bandit problem.

  Adds another dimension.

  MWA: Greedy $\epsilon = 0$ vs. 1

  (Thyagarajan et al. arXiv:1308.0565)

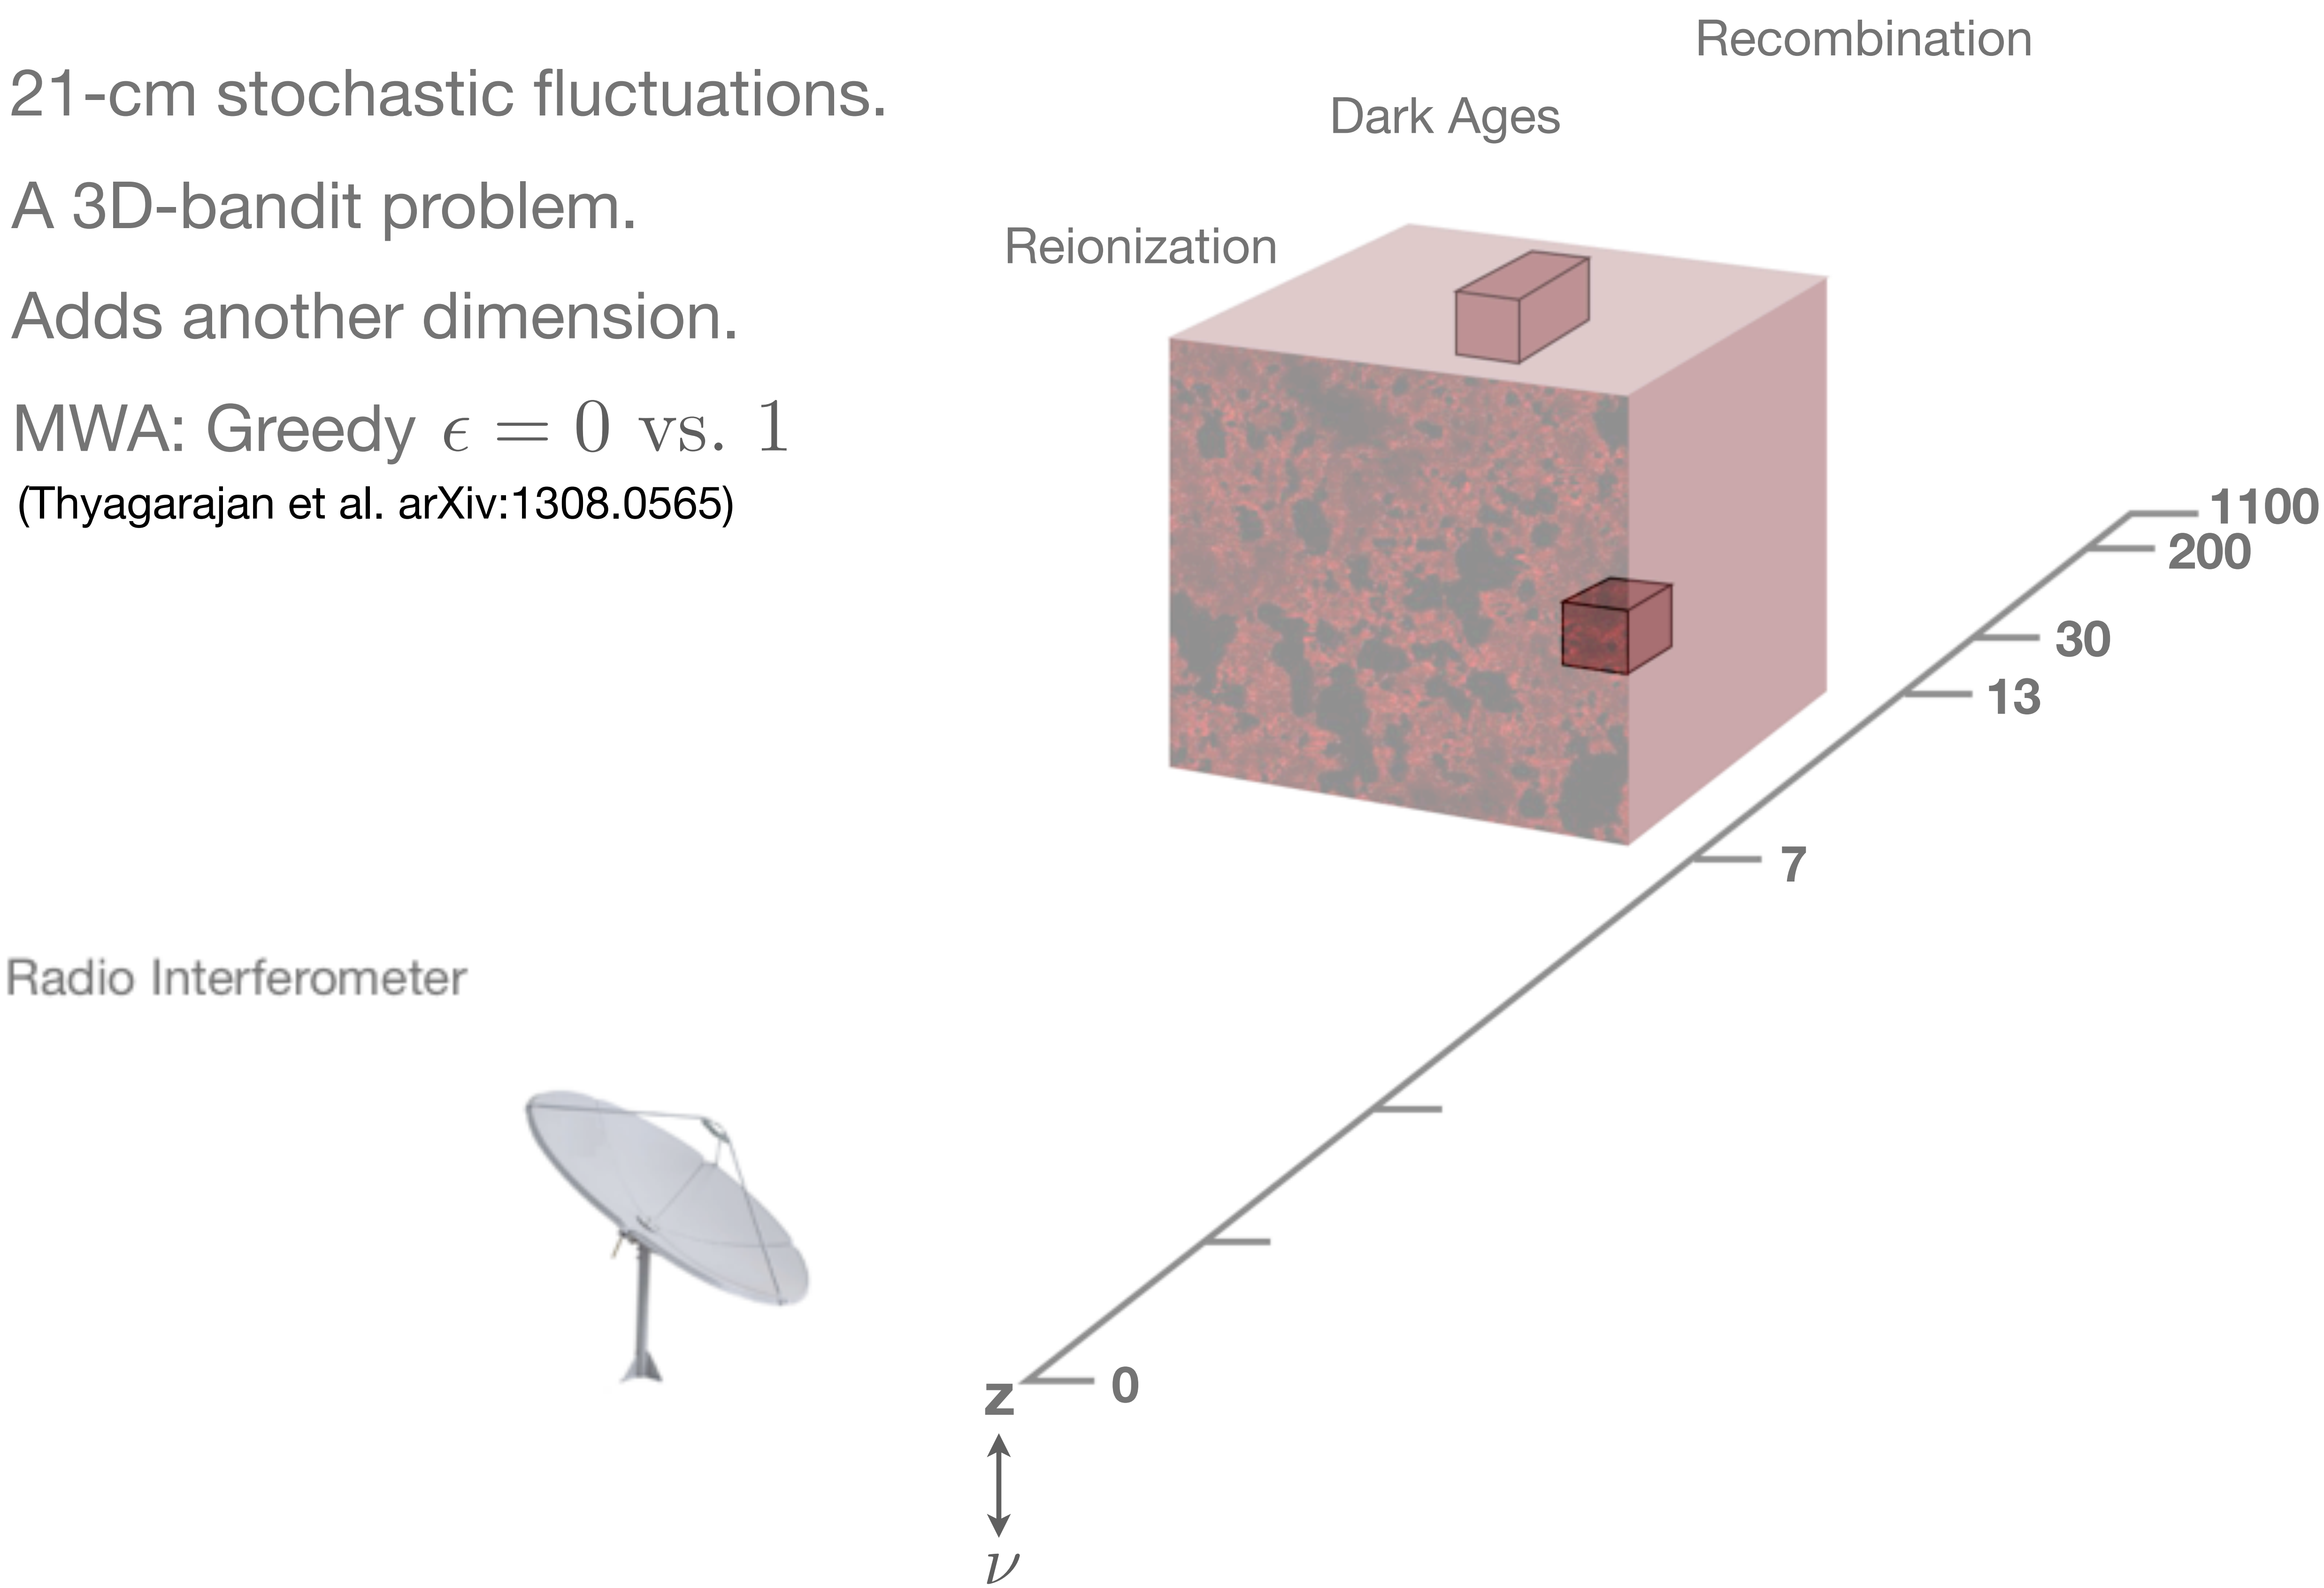

MAB Elsewhere

MAB Elsewhere

- Deep-field imaging:

  From HST to JWST?

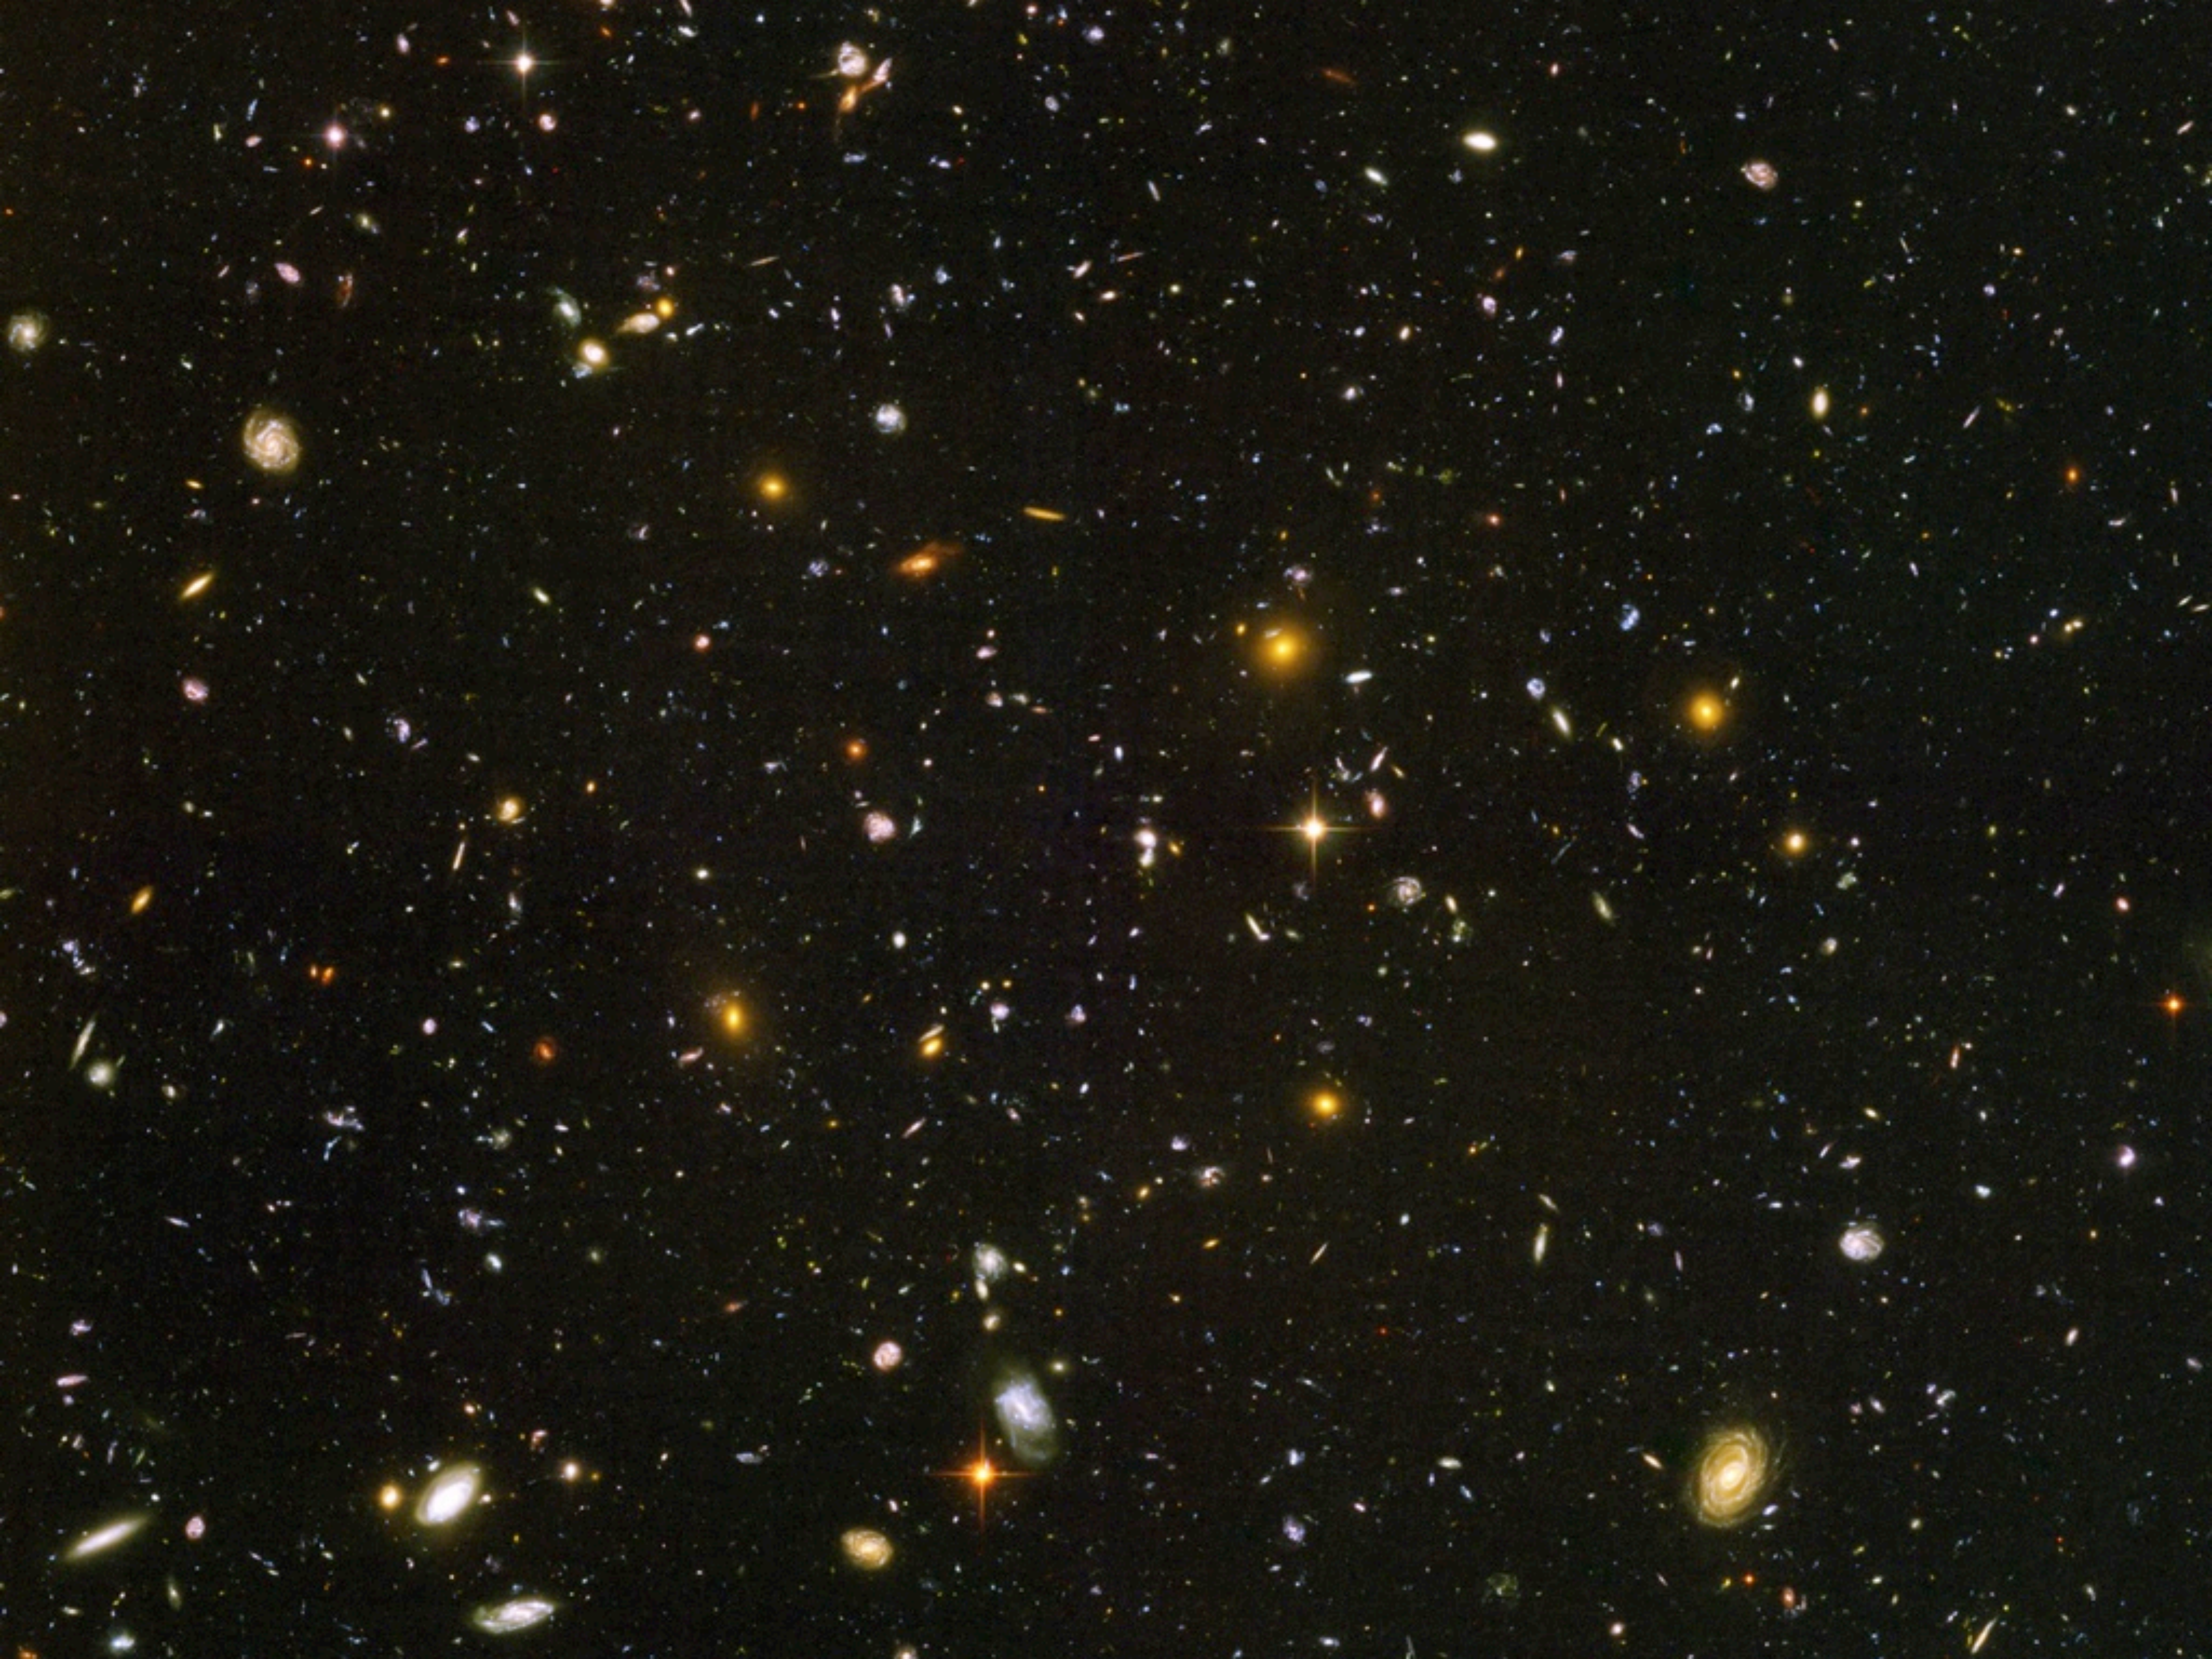

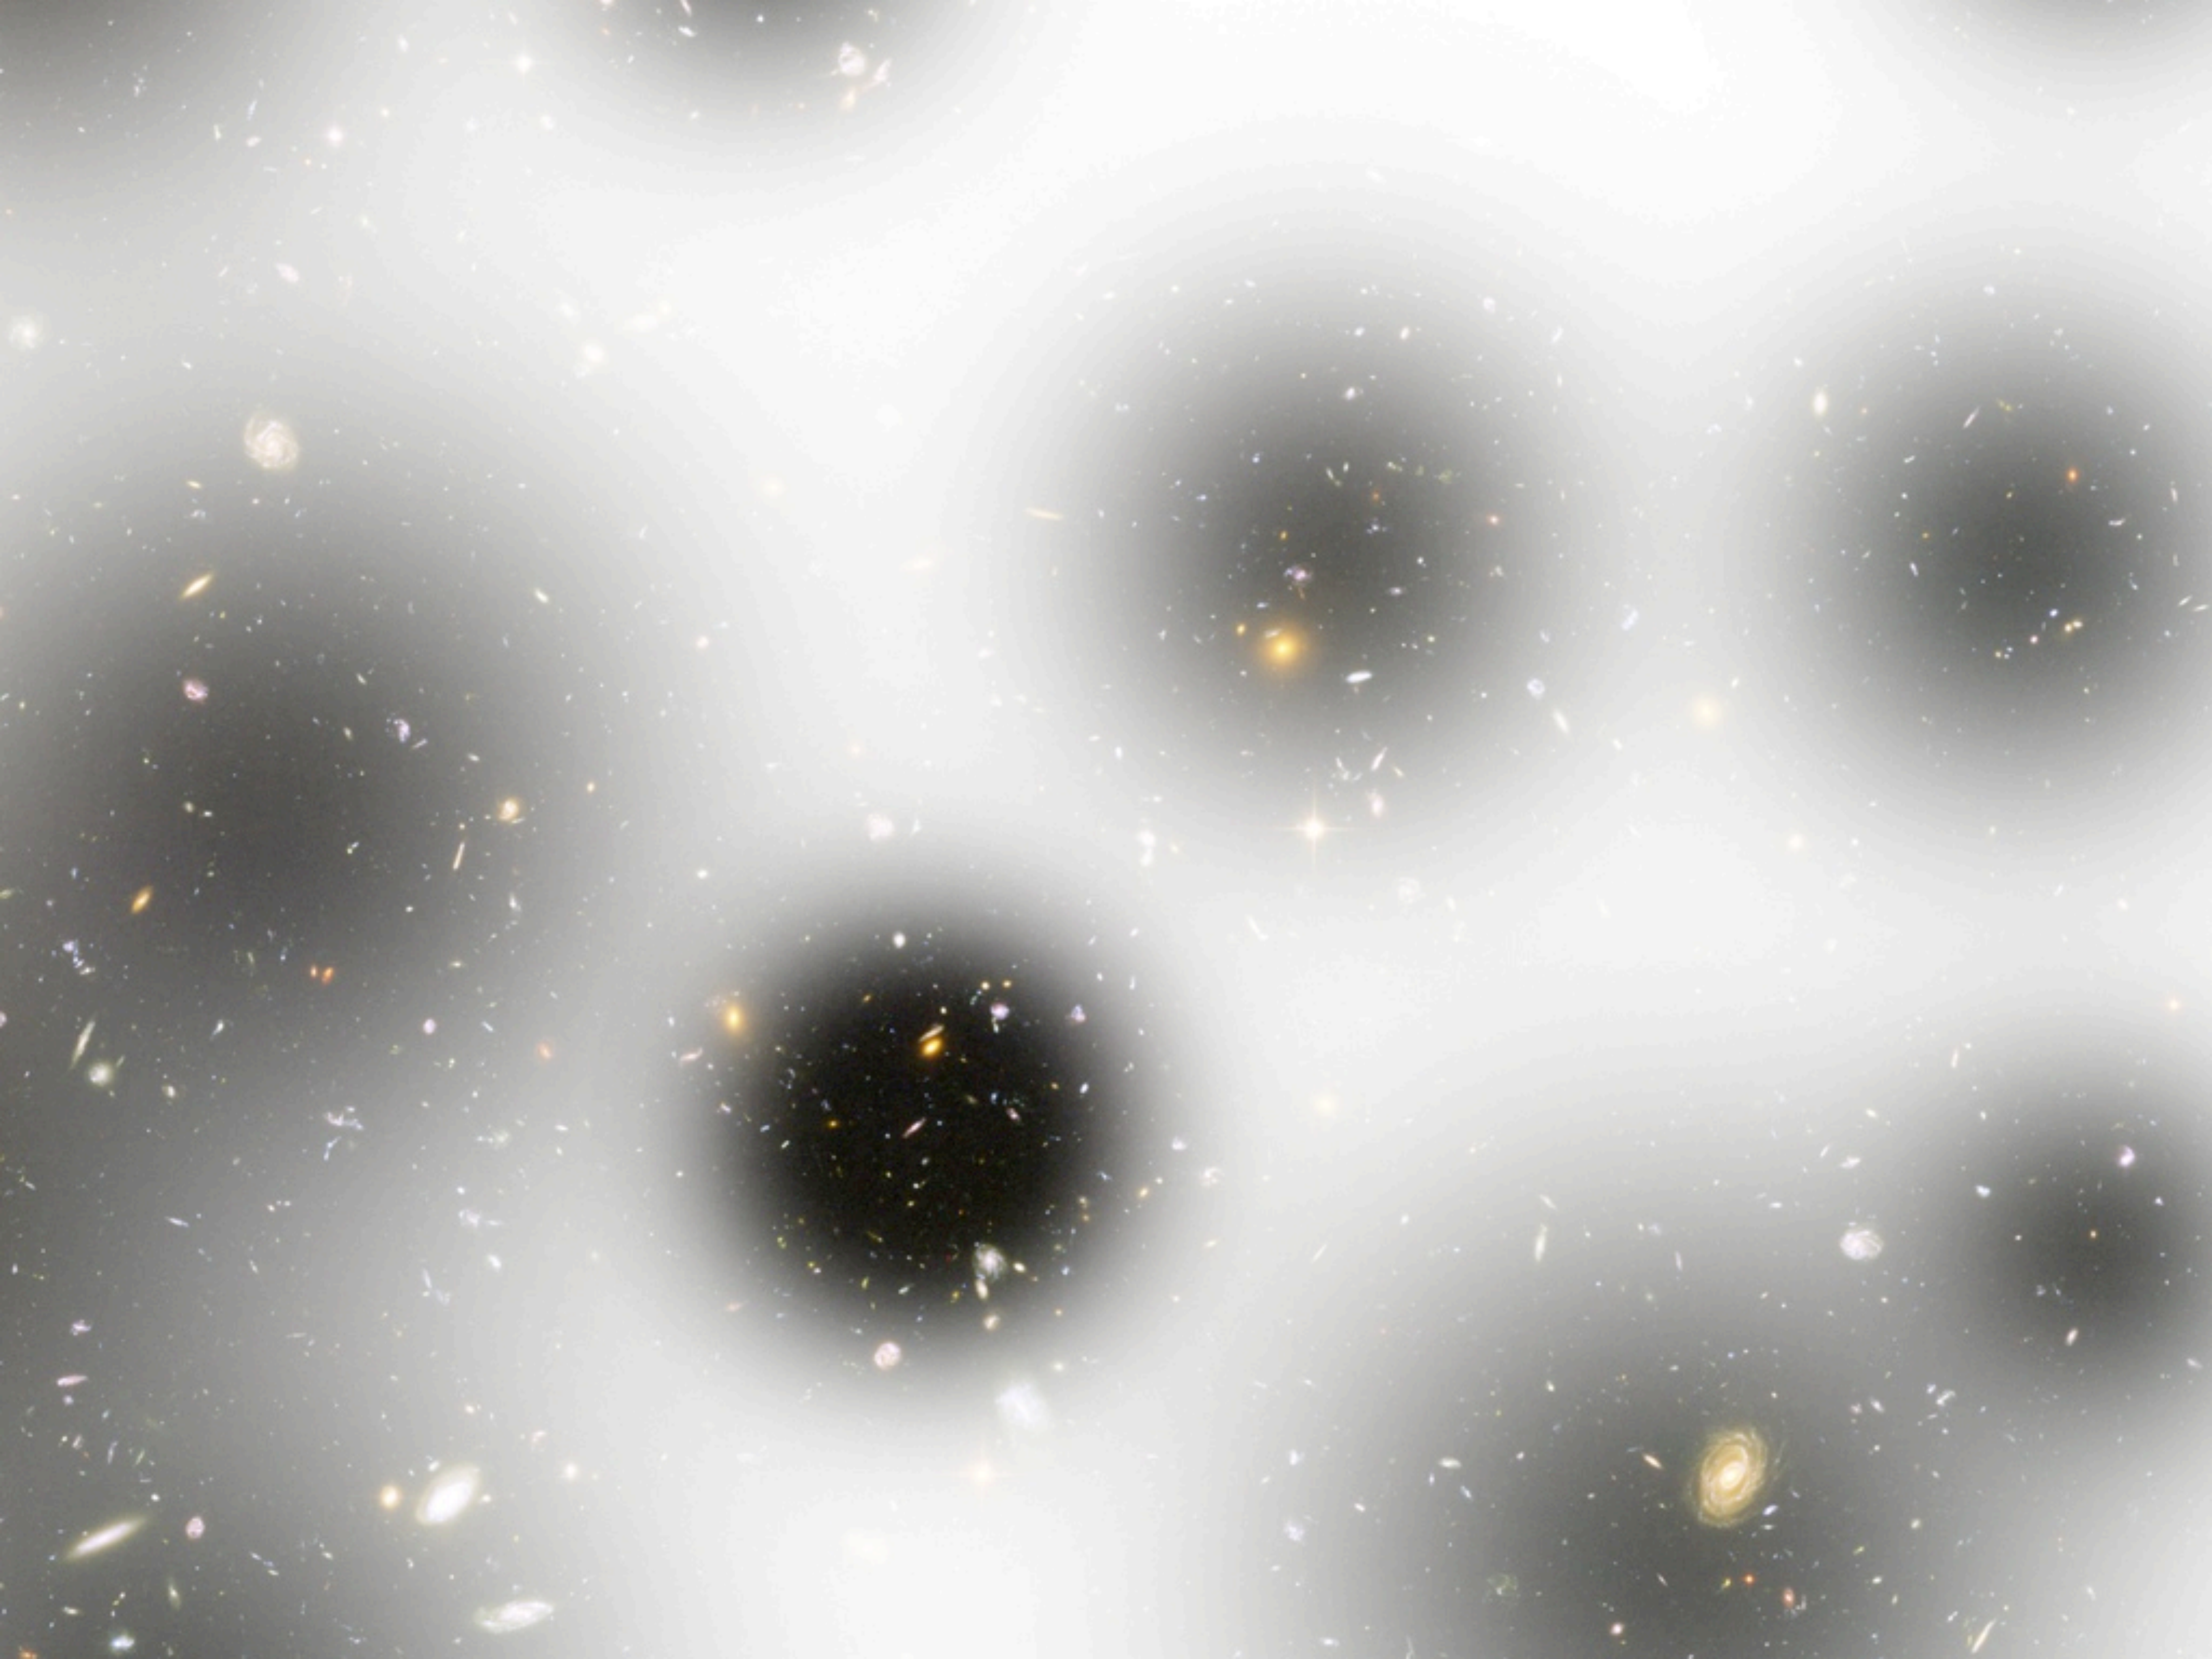

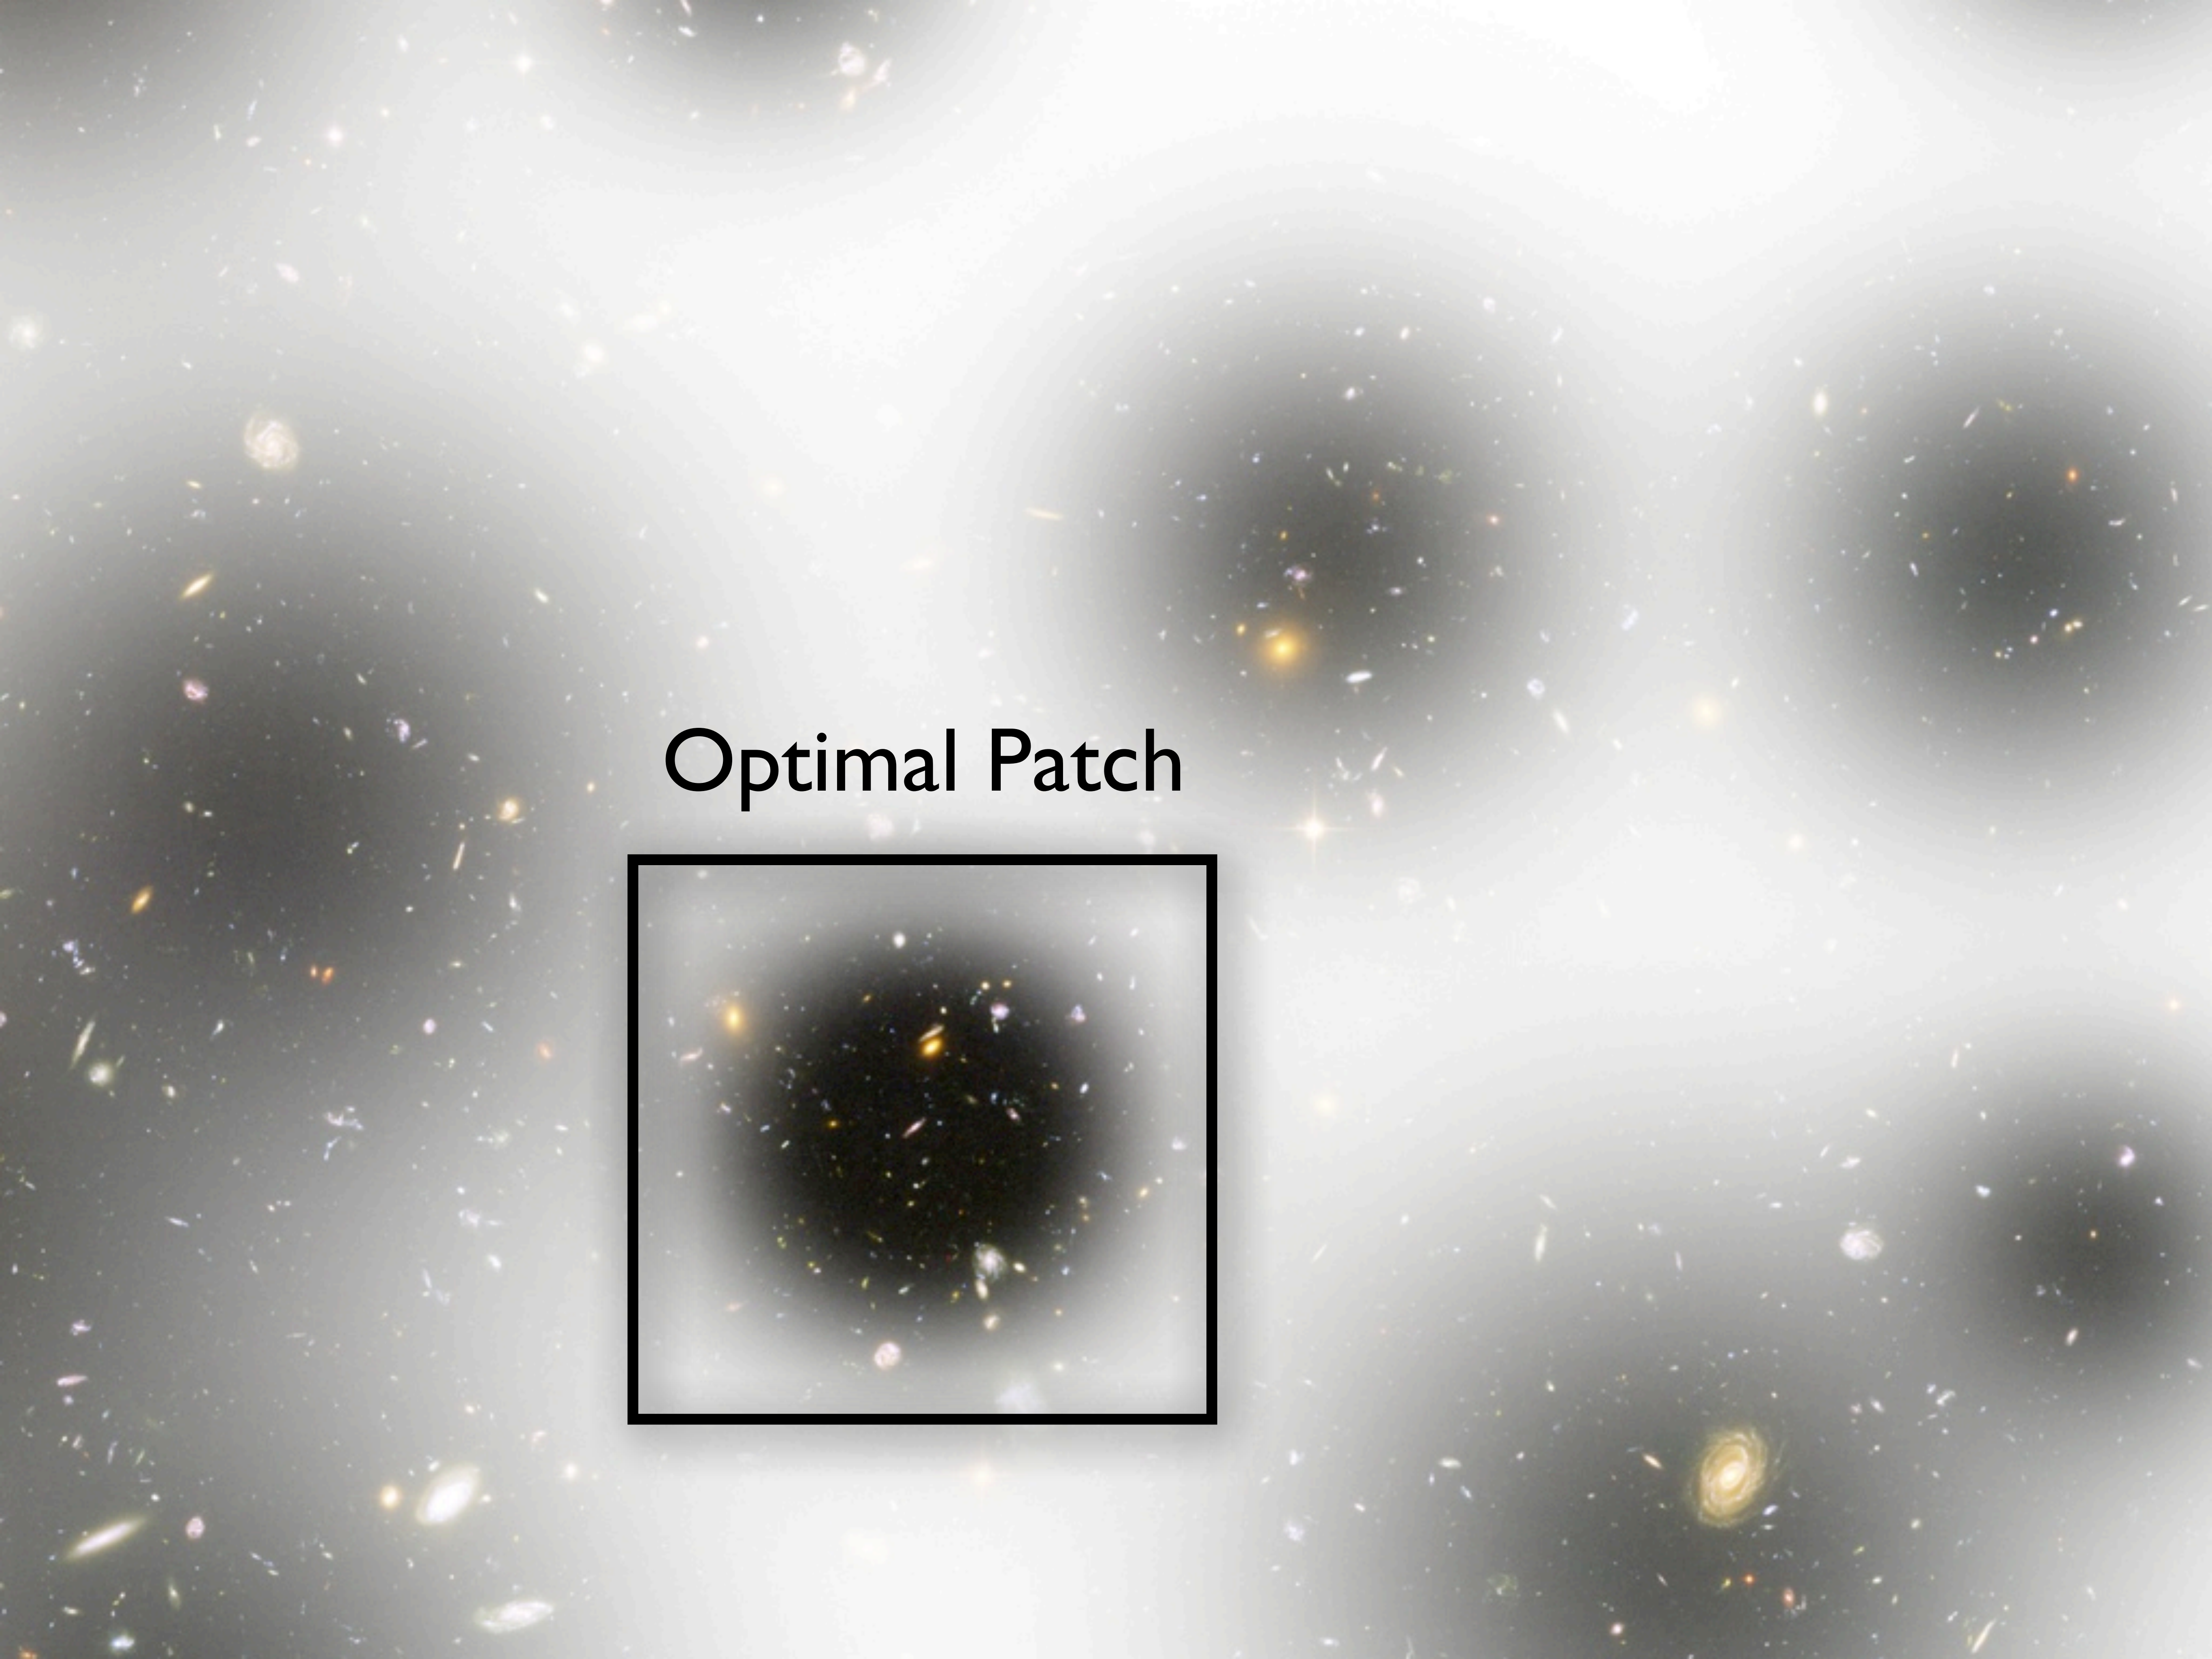
Optimal Patch

Optimal Patch

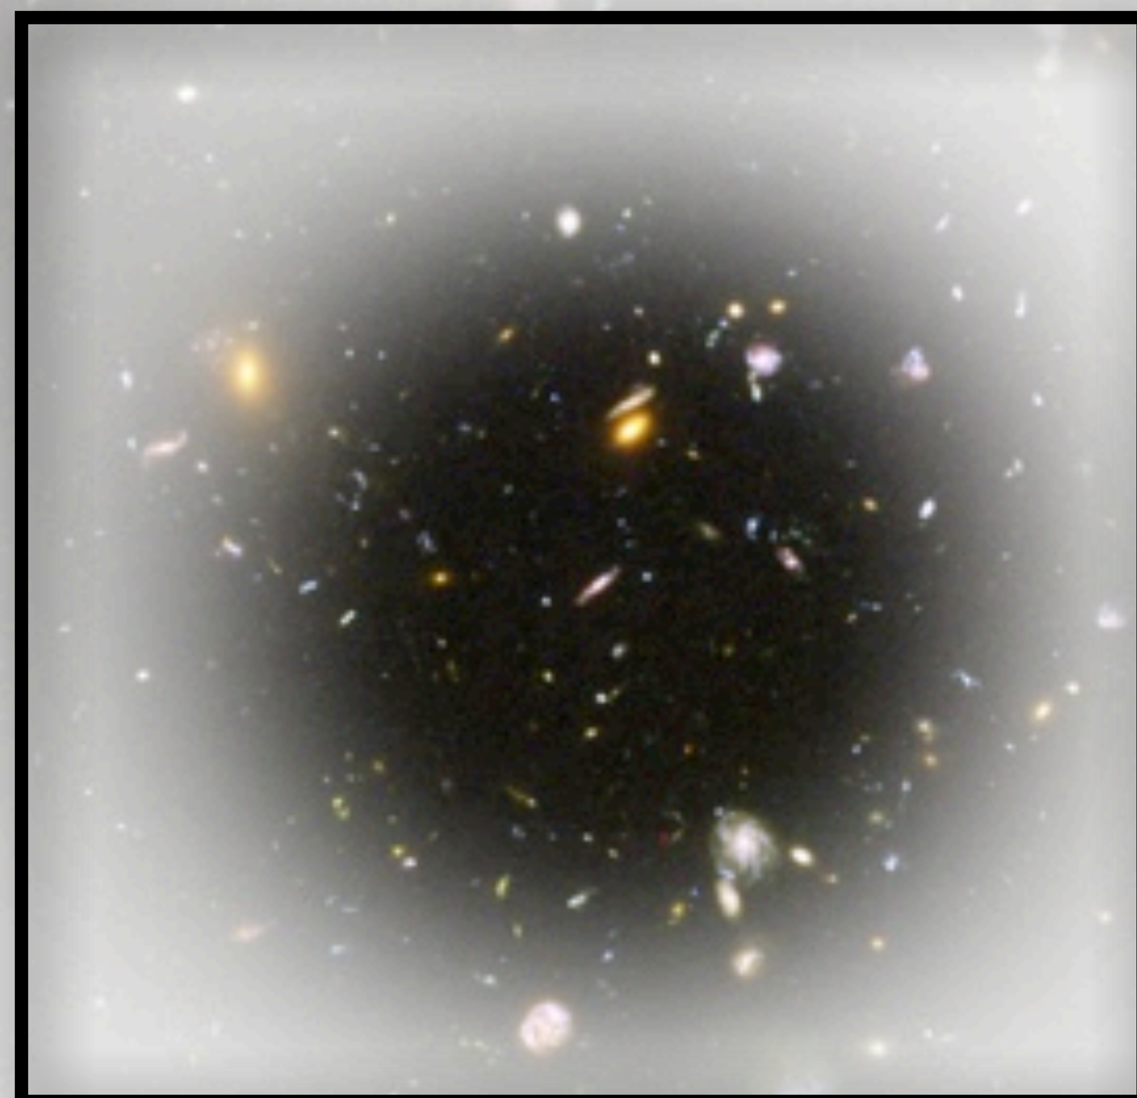

MAB Strategies Elsewhere

MAB Strategies Elsewhere

Measuring the Luminosity function and calibrating simulations. (with P. Behroozi @CANDELS)

MAB Strategies Elsewhere

Measuring the Luminosity function and calibrating simulations. (with P. Behroozi @CANDELS)

- Cosmic variance uncertainty plays a huge role.

MAB Strategies Elsewhere

Measuring the Luminosity function and calibrating simulations. (with P. Behroozi @CANDELS)

- Cosmic variance uncertainty plays a huge role.

(Oesch et al. arXiv:1201.0755)

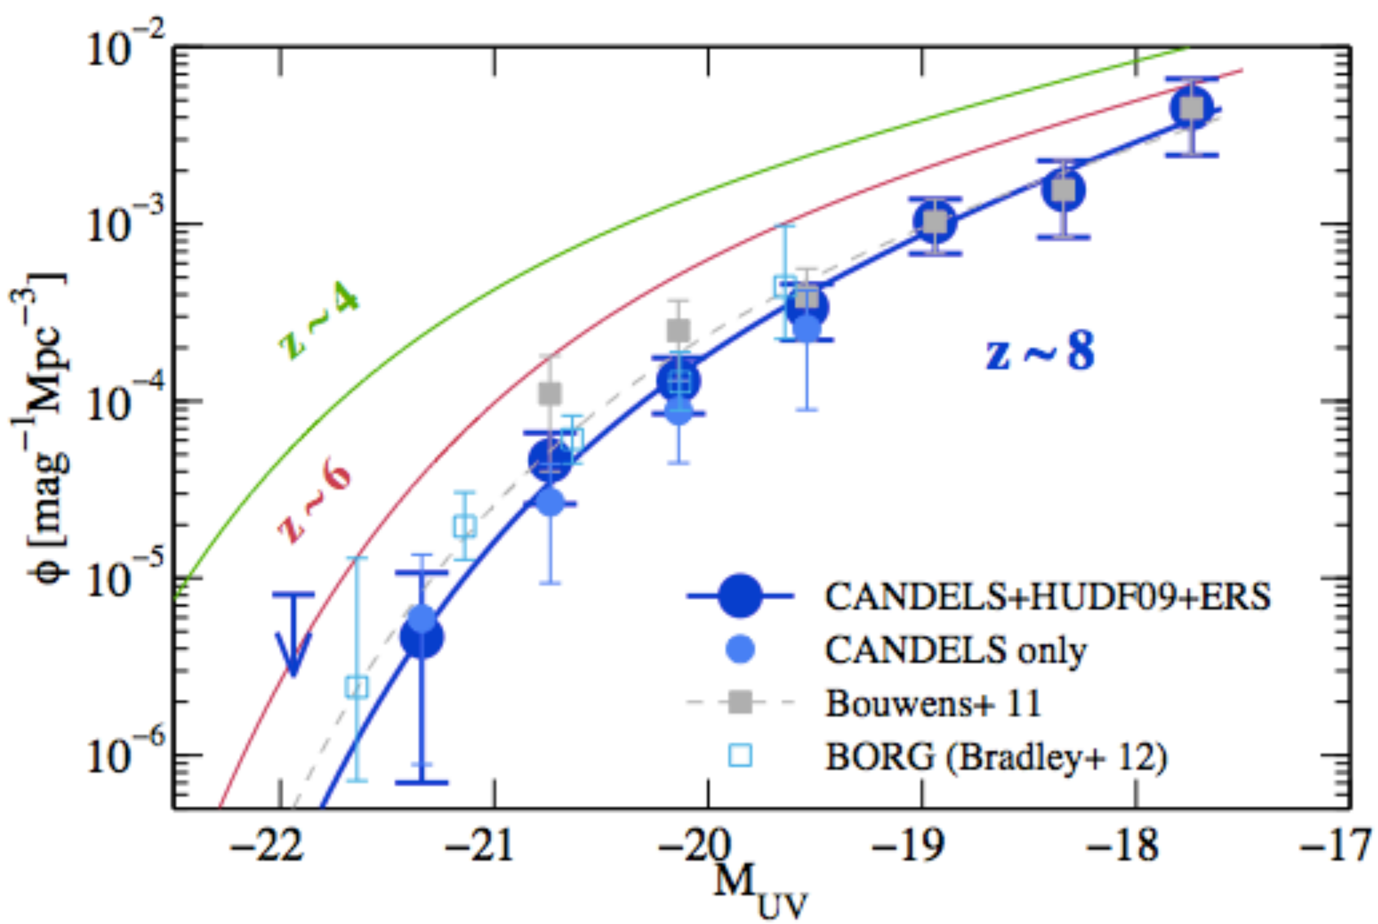

(Grogin et al. arXiv:1105.3753)

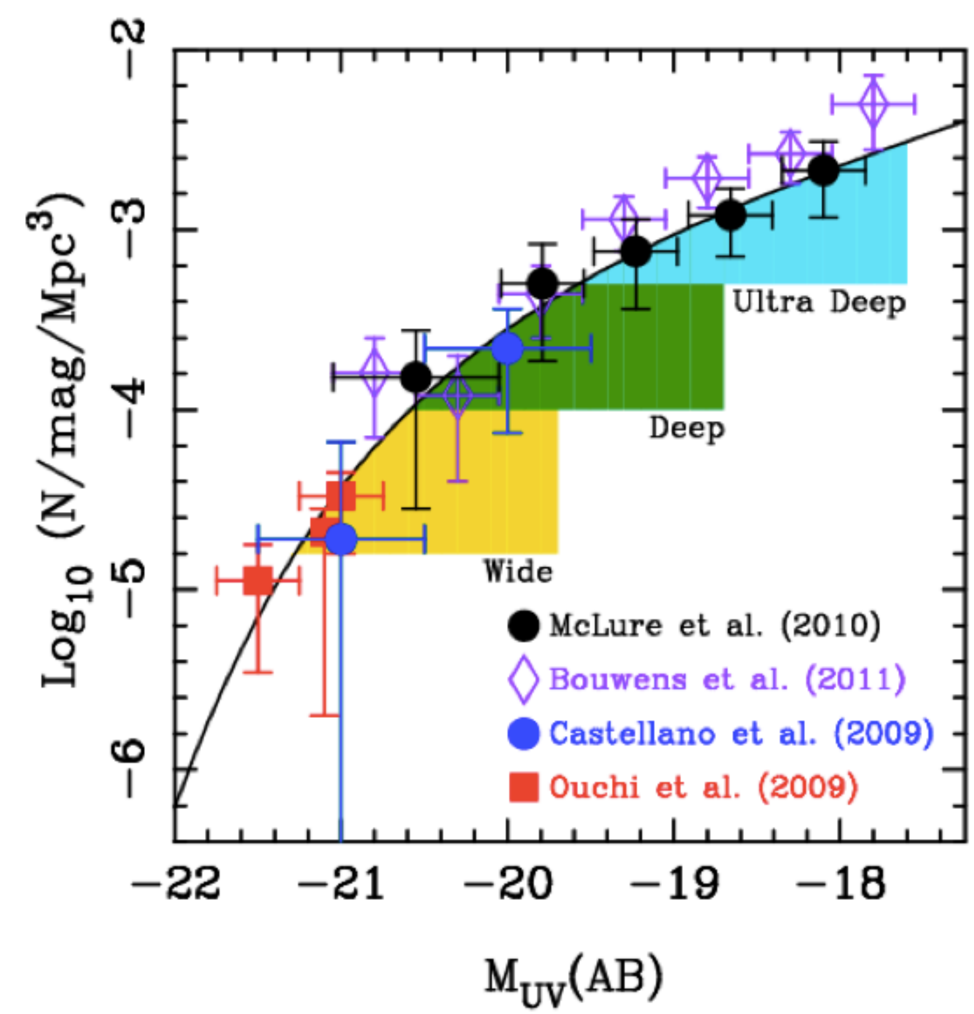

MAB Strategies Elsewhere

Measuring the Luminosity function and calibrating simulations. (with P. Behroozi @CANDELS)

- Cosmic variance uncertainty plays a huge role.

(Oesch et al. arXiv:1201.0755)

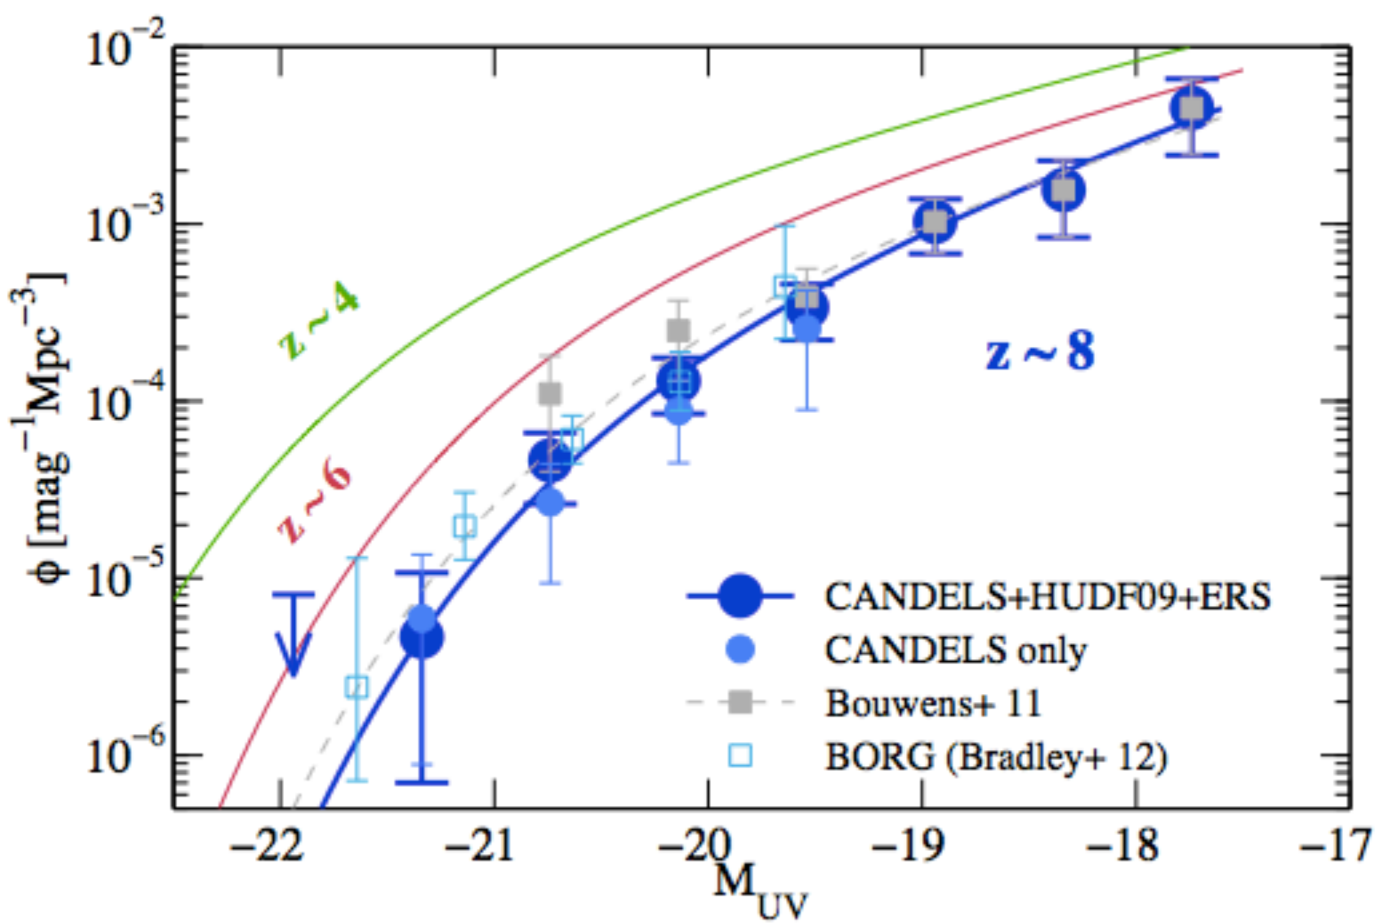

(Grogin et al. arXiv:1105.3753)

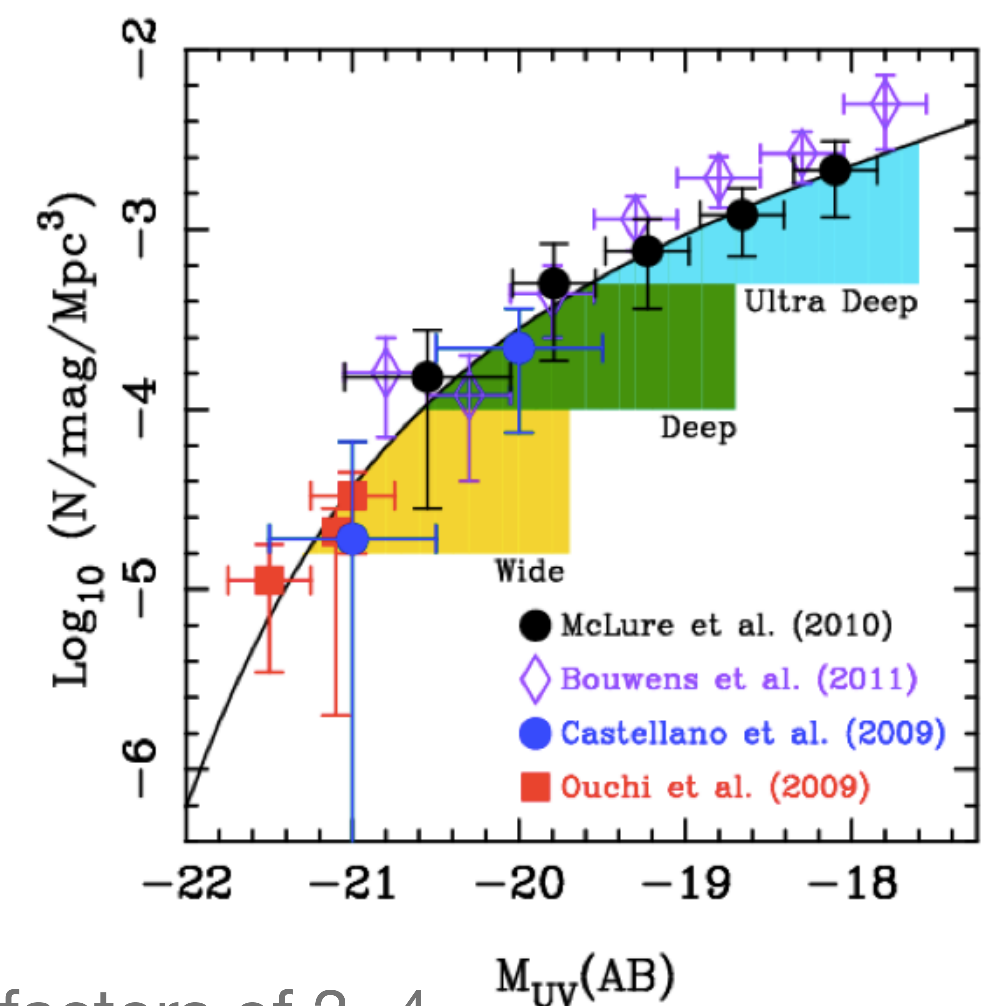

CANDELS fields and HUDF differ at z~7,8 by factors of 3–4.

MAB Strategies Elsewhere

Measuring the Luminosity function and calibrating simulations. (with P. Behroozi @CANDELS)

- Cosmic variance uncertainty plays a huge role.

(Oesch et al. arXiv:1201.0755)

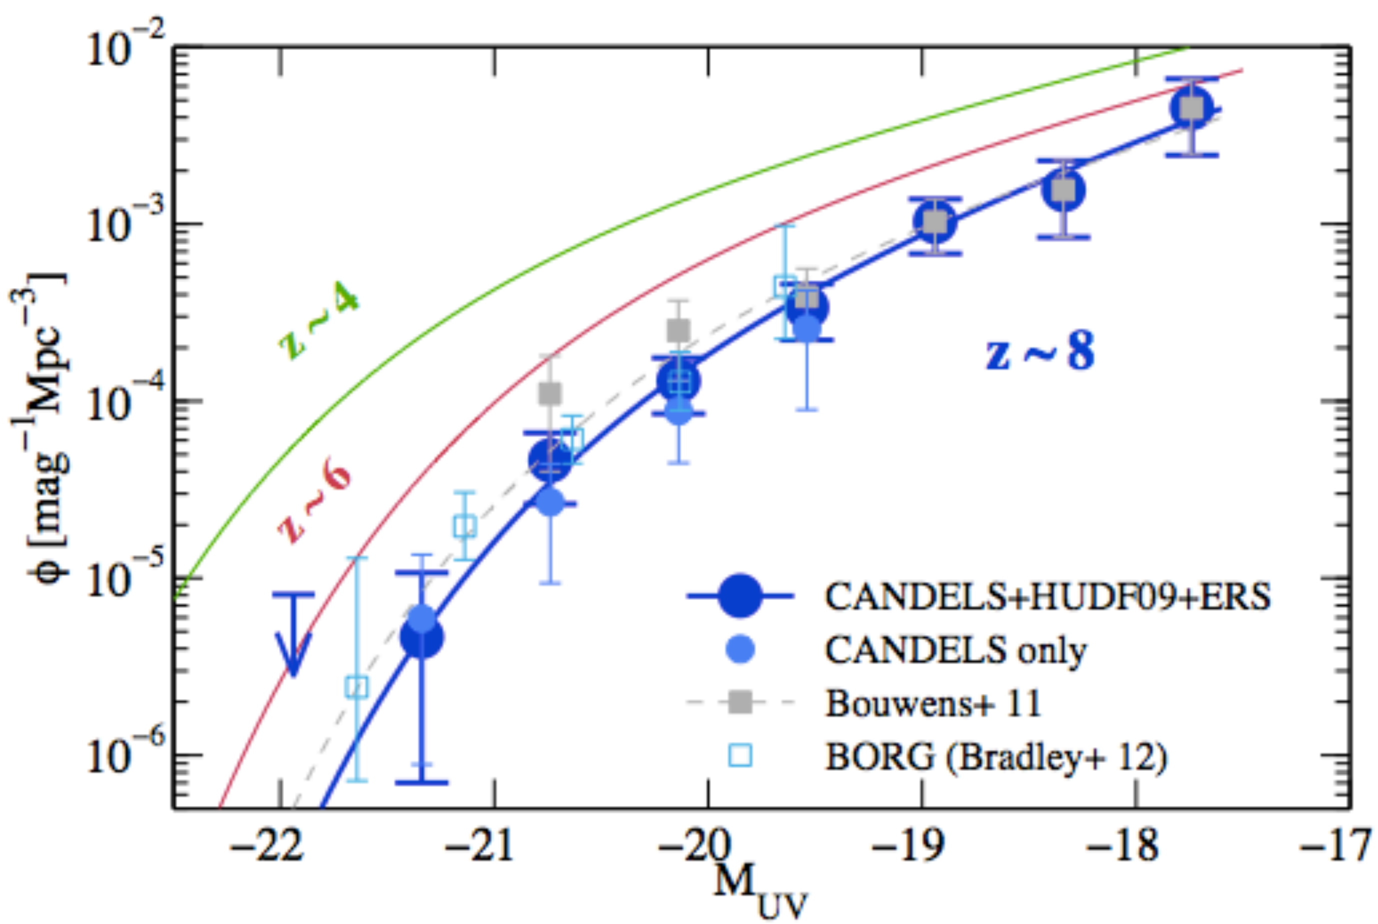

(Grogin et al. arXiv:1105.3753)

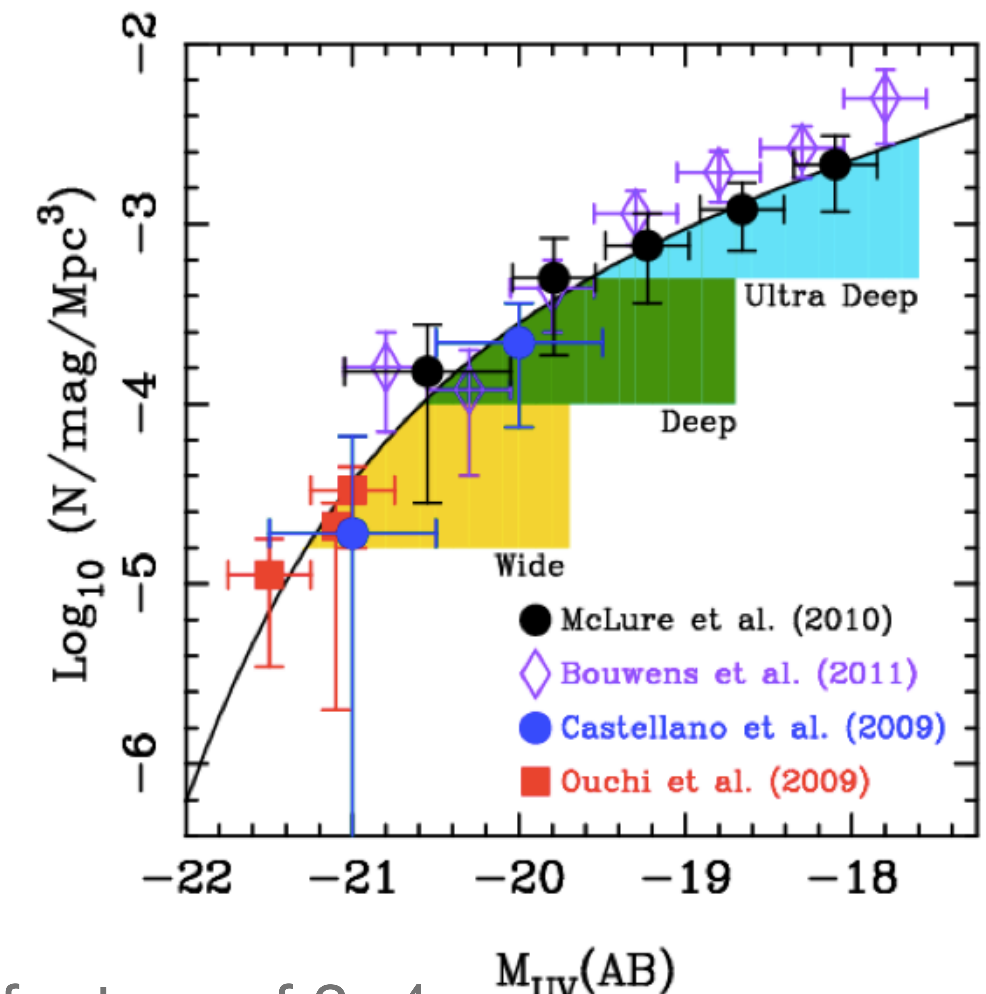

CANDELS fields and HUDF differ at z~7,8 by factors of 3–4.

- An efficient adaptive strategy would converge onto the *cosmic mean*.

  (Not useful in a Casino, but may save considerable telescope time)

Thank you!

Ely D. Kovetz
University of Texas at Austin
ITC Seminar, Dec. 10th, 2013

Thank you!

Ely D. Kovetz
University of Texas at Austin
ITC Seminar, Dec. 10th, 2013

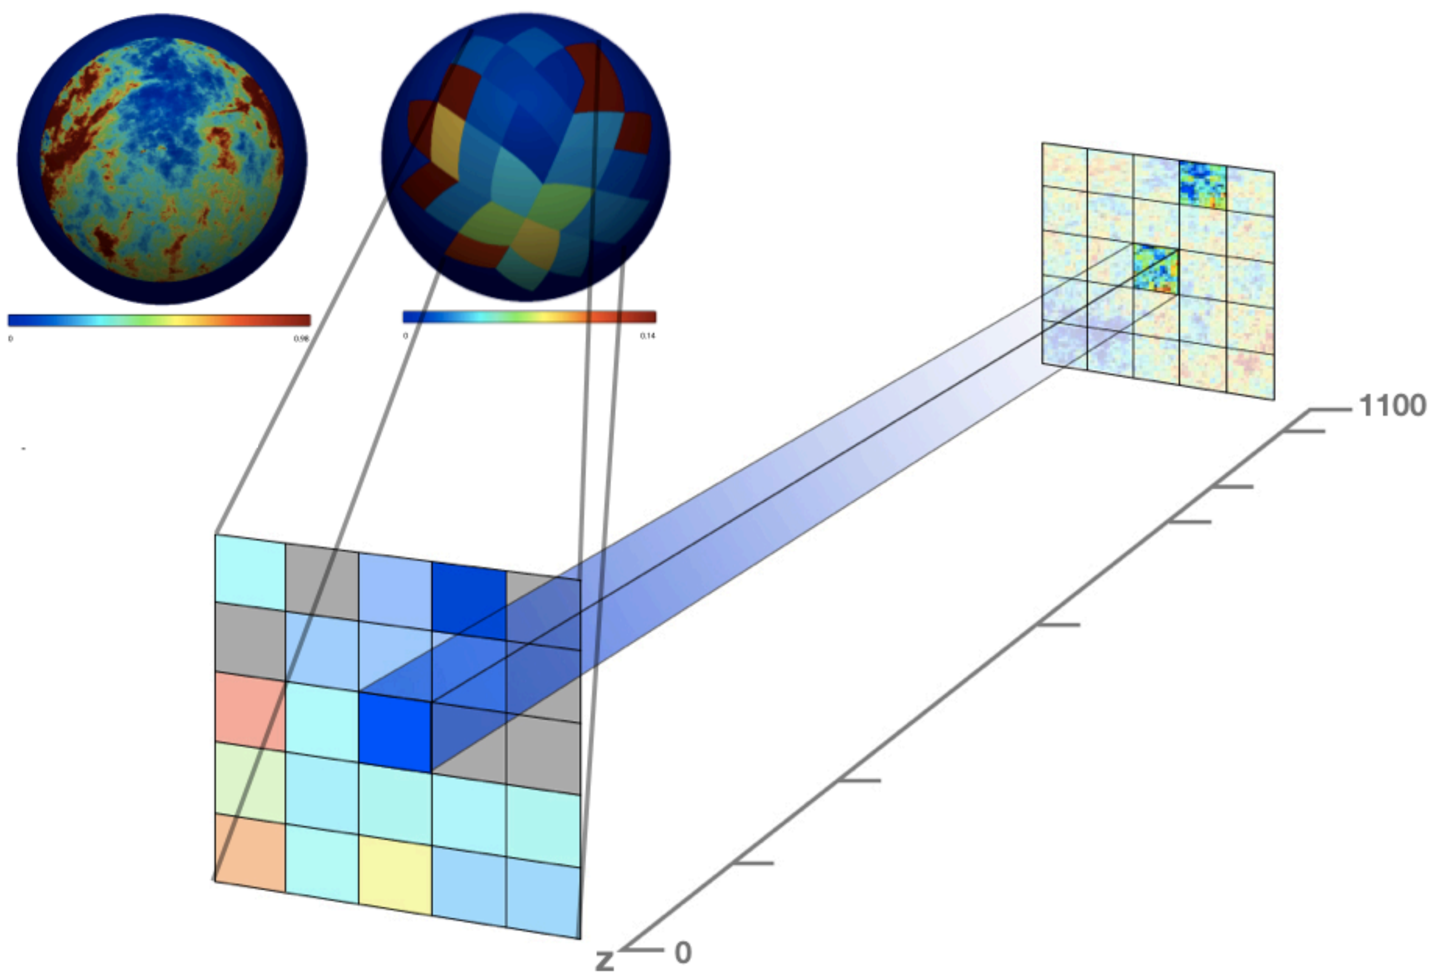

Based on: arXiv:1308.1404
and current work
with M. Kamionkowski (JHU)

Thank you!

Ely D. Kovetz
University of Texas at Austin
ITC Seminar, Dec. 10th, 2013

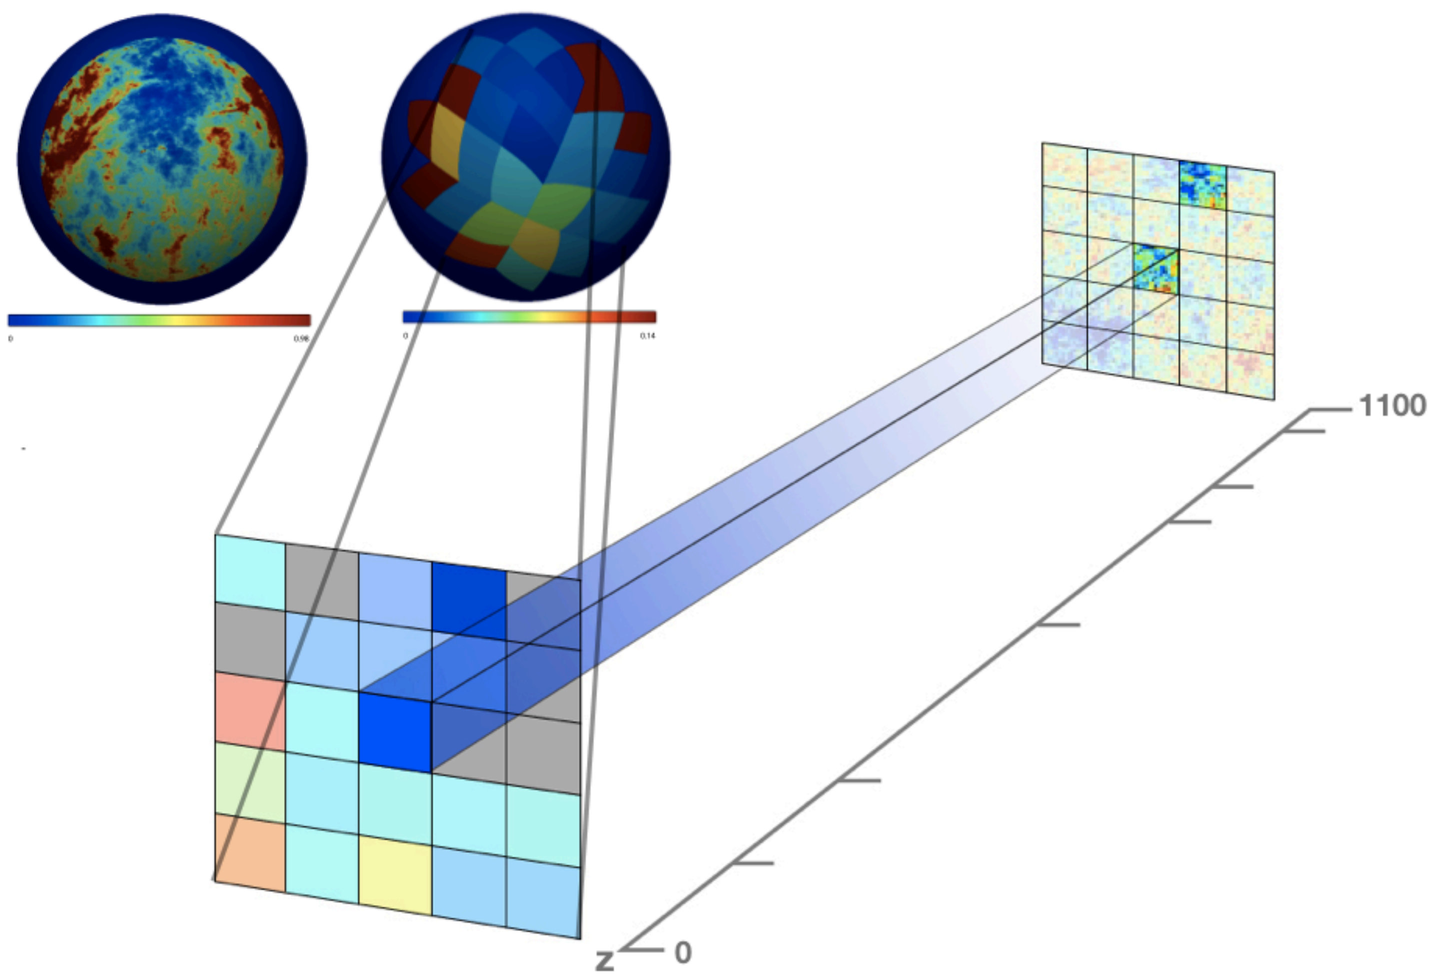

Based on: arXiv:1308.1404
and current work
with M. Kamionkowski (JHU)

Discussion: B-mode detection

Discussion: B-mode detection

Recall some assumptions:

Discussion: B-mode detection

Recall some assumptions:

- Single frequency. In practice, remain with (non-zero) foreground residuals.

Discussion: B-mode detection

Recall some assumptions:

- Single frequency. In practice, remain with (non-zero) foreground residuals.
- Specific template (FGPol). Real foregrounds may differ, probably not by much.

Discussion: B-mode detection

Recall some assumptions:

- Single frequency. In practice, remain with (non-zero) foreground residuals.
- Specific template (FGPol). Real foregrounds may differ, probably not by much.
- Power-law power spectrum. Exact shape can be calculated from template.

Discussion: B-mode detection

Recall some assumptions:

- Single frequency. In practice, remain with (non-zero) foreground residuals.
- Specific template (FGPol). Real foregrounds may differ, probably not by much.
- Power-law power spectrum. Exact shape can be calculated from template.
- Gaussian likelihood curves. Full simulation would draw from χ^2 at each ℓ.

Discussion: B-mode detection

Recall some assumptions:

- Single frequency. In practice, remain with (non-zero) foreground residuals.
- Specific template (FGPol). Real foregrounds may differ, probably not by much.
- Power-law power spectrum. Exact shape can be calculated from template.
- Gaussian likelihood curves. Full simulation would draw from χ^2 at each ℓ.
- Limited scenarios. More depth between pessimistic and optimistic.

Discussion: B-mode detection

Recall some assumptions:

- Single frequency. In practice, remain with (non-zero) foreground residuals.
- Specific template (FGPol). Real foregrounds may differ, probably not by much.
- Power-law power spectrum. Exact shape can be calculated from template.
- Gaussian likelihood curves. Full simulation would draw from χ^2 at each ℓ.
- Limited scenarios. More depth between pessimistic and optimistic.
- Cost of moving telescope target. This *should* be taken into account.

Discussion: B-mode detection

Recall some assumptions:

- Single frequency. In practice, remain with (non-zero) foreground residuals.
- Specific template (FGPol). Real foregrounds may differ, probably not by much.
- Power-law power spectrum. Exact shape can be calculated from template.
- Gaussian likelihood curves. Full simulation would draw from χ^2 at each ℓ.
- Limited scenarios. More depth between pessimistic and optimistic.
- Cost of moving telescope target. This *should* be taken into account.

Actual performance may deteriorate. However:

Discussion: B-mode detection

Recall some assumptions:

- Single frequency. In practice, remain with (non-zero) foreground residuals.
- Specific template (FGPol). Real foregrounds may differ, probably not by much.
- Power-law power spectrum. Exact shape can be calculated from template.
- Gaussian likelihood curves. Full simulation would draw from χ^2 at each ℓ.
- Limited scenarios. More depth between pessimistic and optimistic.
- Cost of moving telescope target. This *should* be taken into account.

Actual performance may deteriorate. However:

- Methods can be optimized further.

Discussion: B-mode detection

Recall some assumptions:

- Single frequency. In practice, remain with (non-zero) foreground residuals.
- Specific template (FGPol). Real foregrounds may differ, probably not by much.
- Power-law power spectrum. Exact shape can be calculated from template.
- Gaussian likelihood curves. Full simulation would draw from χ^2 at each ℓ.
- Limited scenarios. More depth between pessimistic and optimistic.
- Cost of moving telescope target. This *should* be taken into account.

Actual performance may deteriorate. However:

- Methods can be optimized further.
- Identification of optimal patches will reap future benefits.

Discussion

Discussion

Full B-mode simulation results:

Discussion

Full B-mode simulation results:

- 25%-75% improvement on average in different scenarios.

Discussion

Full B-mode simulation results:

- 25%-75% improvement on average in different scenarios.
- Similar improvements when comparing worst-case performances.

Discussion

Full B-mode simulation results:

- 25%-75% improvement on average in different scenarios.
- Similar improvements when comparing worst-case performances.
- Improvement in any experiment. Maximized with high resolution + sensitivity.

Discussion

Full B-mode simulation results:

- 25%-75% improvement on average in different scenarios.
- Similar improvements when comparing worst-case performances.
- Improvement in any experiment. Maximized with high resolution + sensitivity.

(Some) Caveats:

Discussion

Full B-mode simulation results:

- 25%-75% improvement on average in different scenarios.
- Similar improvements when comparing worst-case performances.
- Improvement in any experiment. Maximized with high resolution + sensitivity.

(Some) Caveats:

- Single frequency. In practice, remain with (non-zero) foreground residuals.

Discussion

Full B-mode simulation results:

- 25%-75% improvement on average in different scenarios.
- Similar improvements when comparing worst-case performances.
- Improvement in any experiment. Maximized with high resolution + sensitivity.

(Some) Caveats:

- Single frequency. In practice, remain with (non-zero) foreground residuals.
- Cost of moving telescope target. This *should* be taken into account.

Discussion

Full B-mode simulation results:

- 25%-75% improvement on average in different scenarios.
- Similar improvements when comparing worst-case performances.
- Improvement in any experiment. Maximized with high resolution + sensitivity.

(Some) Caveats:

- Single frequency. In practice, remain with (non-zero) foreground residuals.
- Cost of moving telescope target. This *should* be taken into account.

Actual performance may deteriorate. However:

Discussion

Full B-mode simulation results:

- 25%-75% improvement on average in different scenarios.
- Similar improvements when comparing worst-case performances.
- Improvement in any experiment. Maximized with high resolution + sensitivity.

(Some) Caveats:

- Single frequency. In practice, remain with (non-zero) foreground residuals.
- Cost of moving telescope target. This *should* be taken into account.

Actual performance may deteriorate. However:

- Methods can be optimized further.

Discussion

Full B-mode simulation results:

- 25%-75% improvement on average in different scenarios.
- Similar improvements when comparing worst-case performances.
- Improvement in any experiment. Maximized with high resolution + sensitivity.

(Some) Caveats:

- Single frequency. In practice, remain with (non-zero) foreground residuals.
- Cost of moving telescope target. This *should* be taken into account.

Actual performance may deteriorate. However:

- Methods can be optimized further.
- Identification of optimal patches will reap future benefits.