How Rocky Are They?

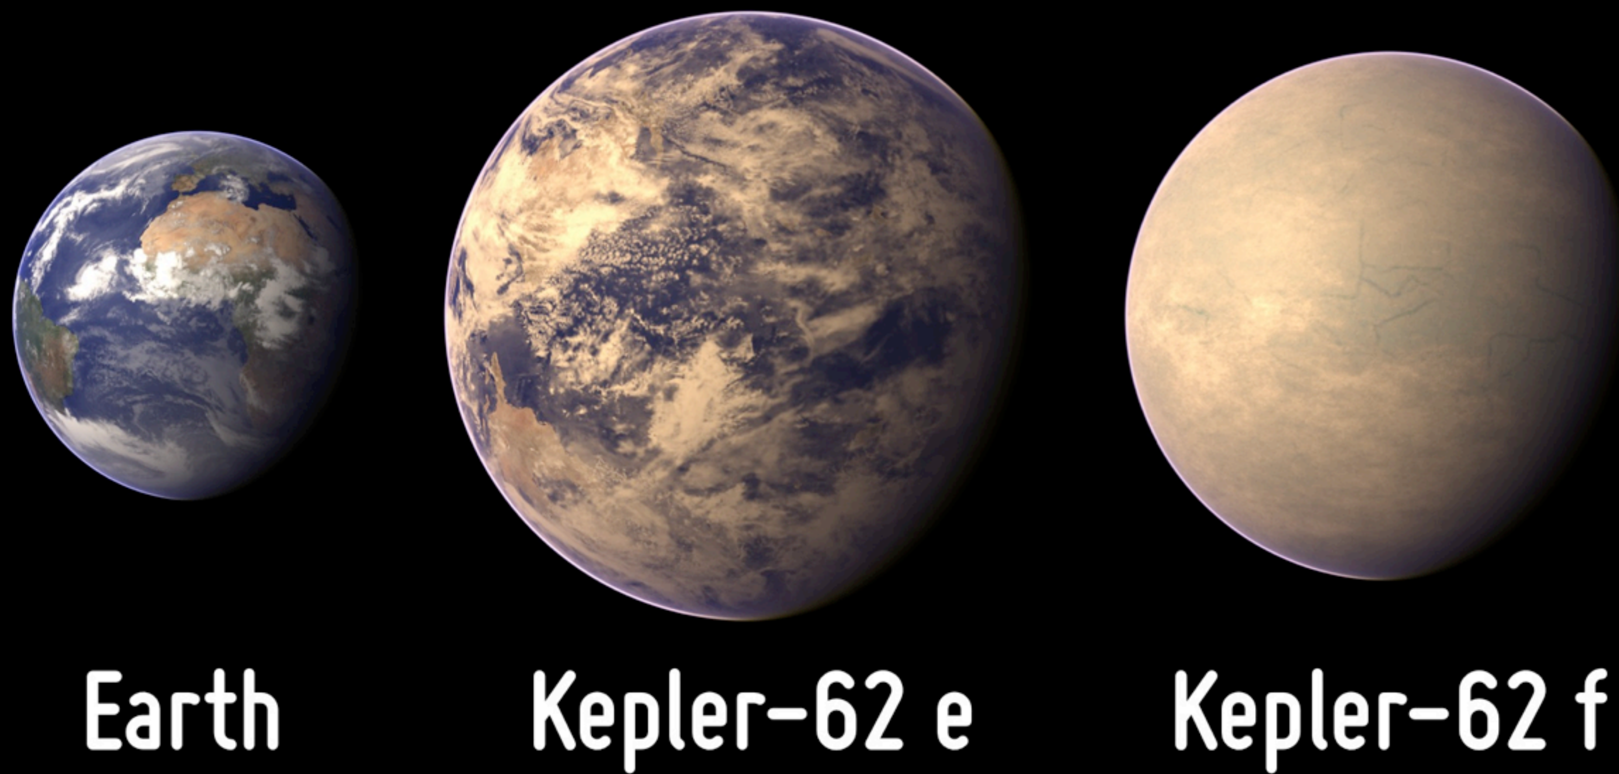

Credit: PHL @ UPR Arecibo

The Composition Distribution of *Kepler*'s Sub-Neptune Planet Candidates

(within 0.15 AU)

Angie Wolfgang, NSF Graduate Research Fellow

SAMSI Bayesian Exoplanet Populations Group, Eric Lopez, Gregory Laughlin

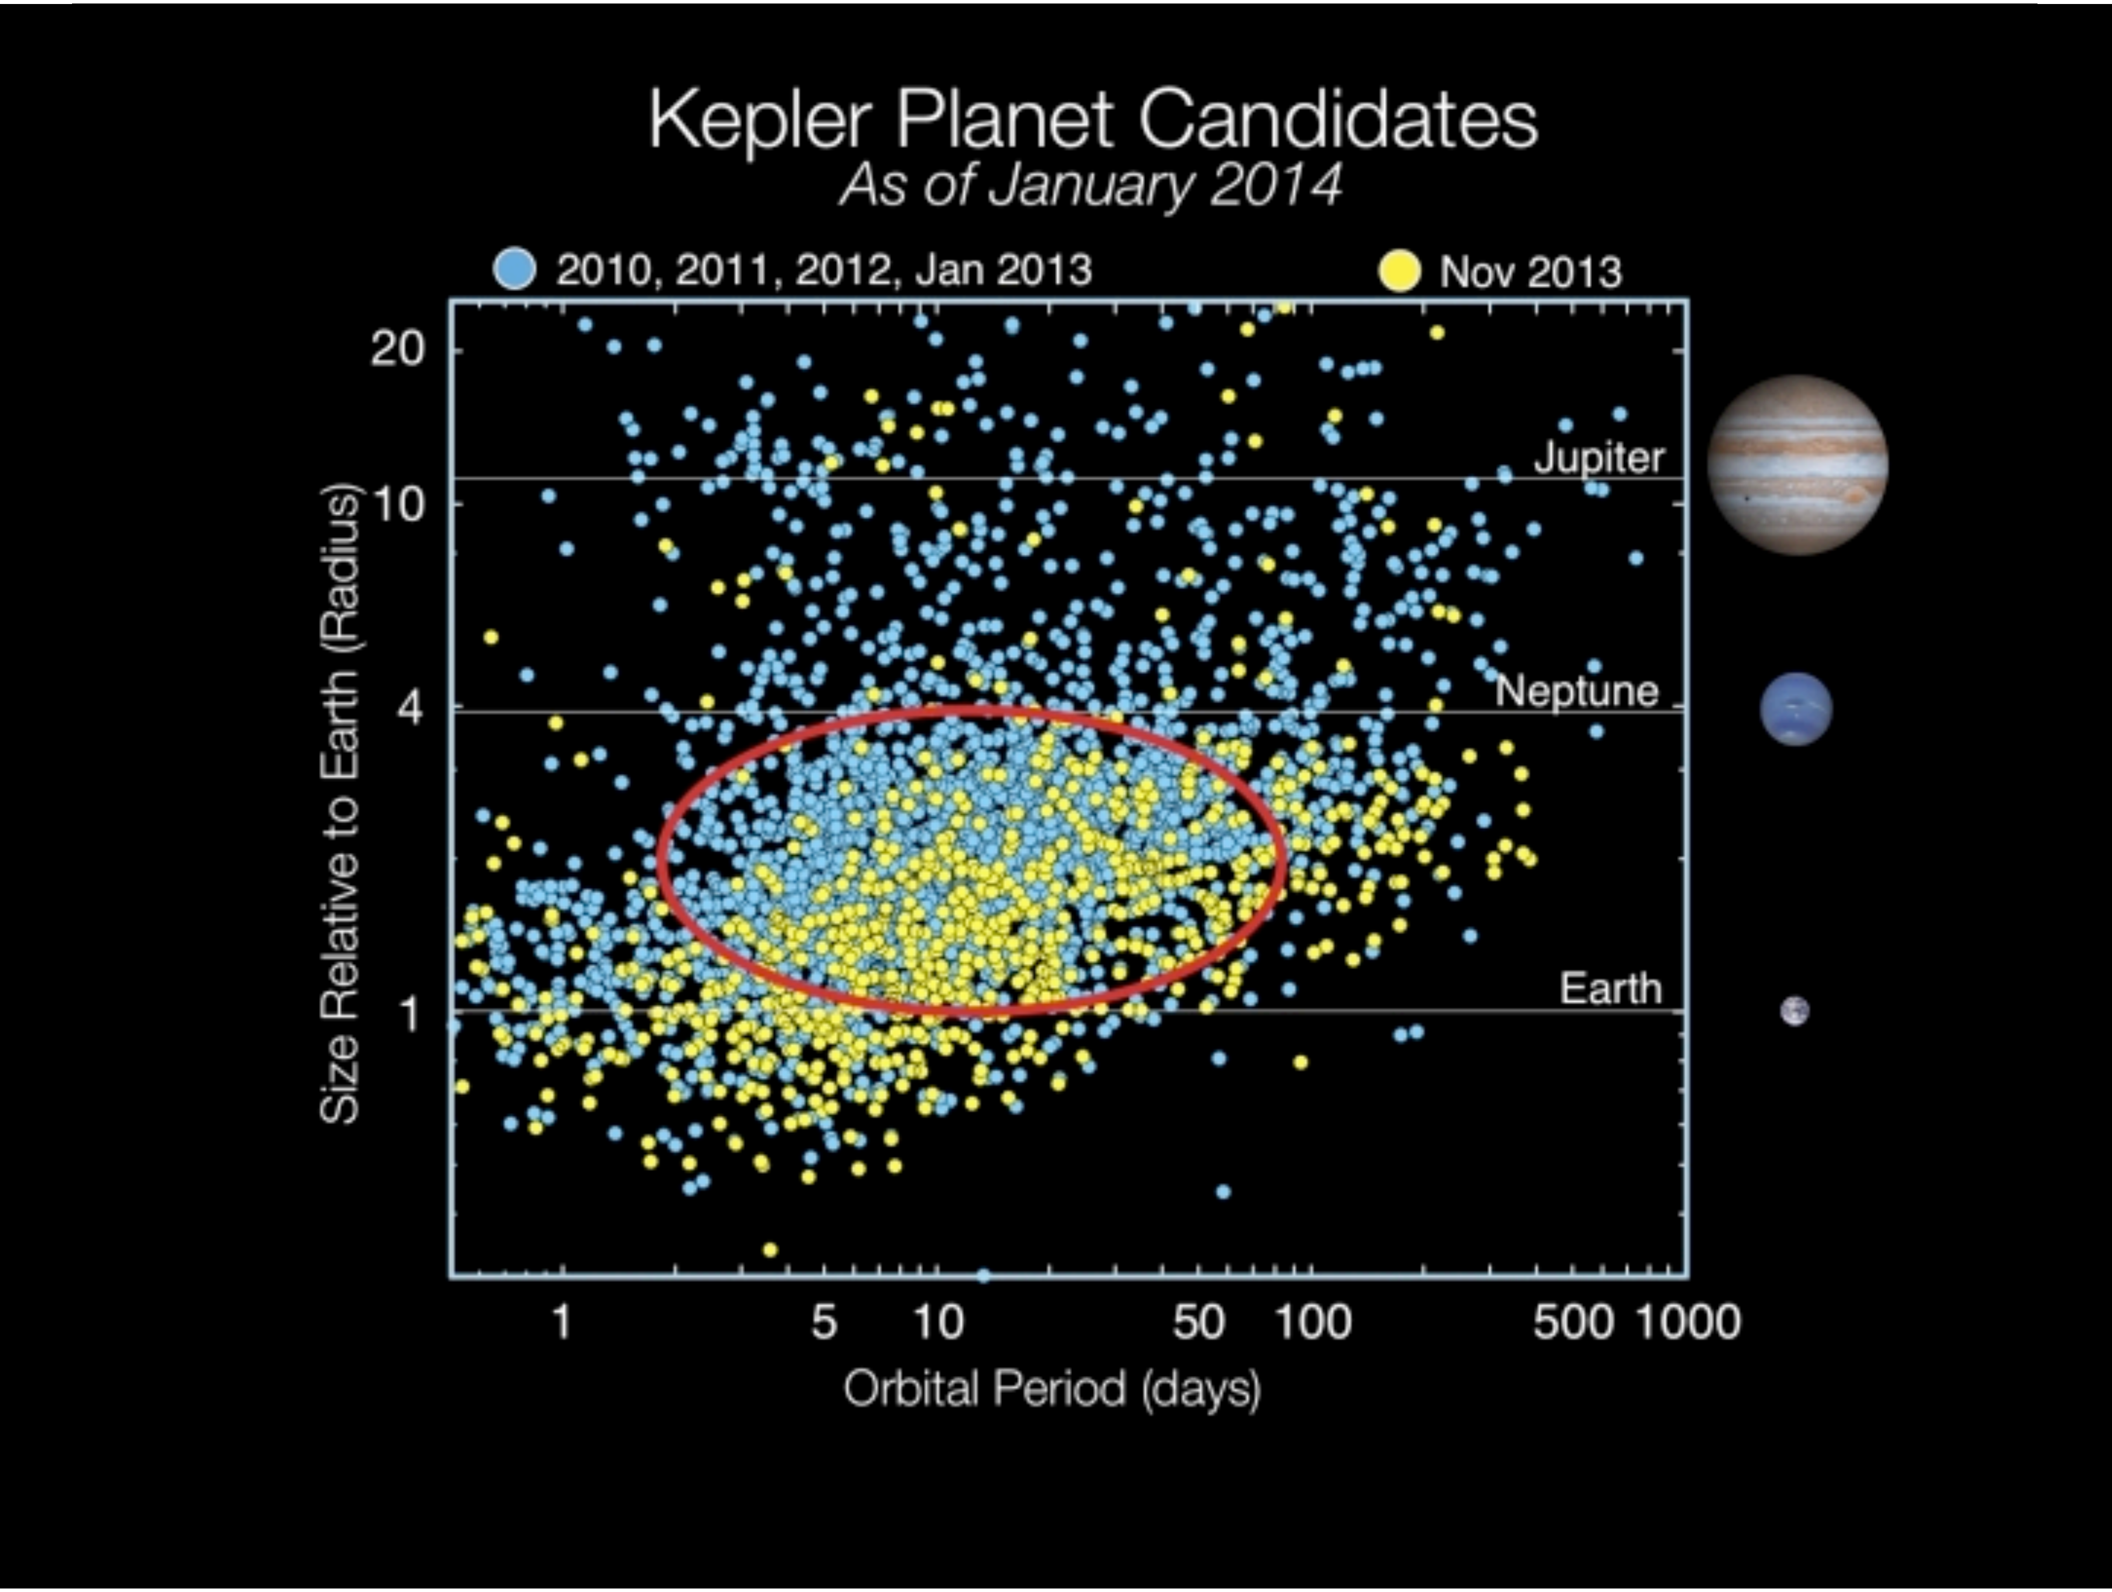
Kepler Planet Candidates
As of January 2014
2010, 2011, 2012, Jan 2013
Nov 2013
Size Relative to Earth (Radius)
20
10
4
1
Jupiter
Neptune
Earth
1
5
10
50
100
500
1000
Orbital Period (days)

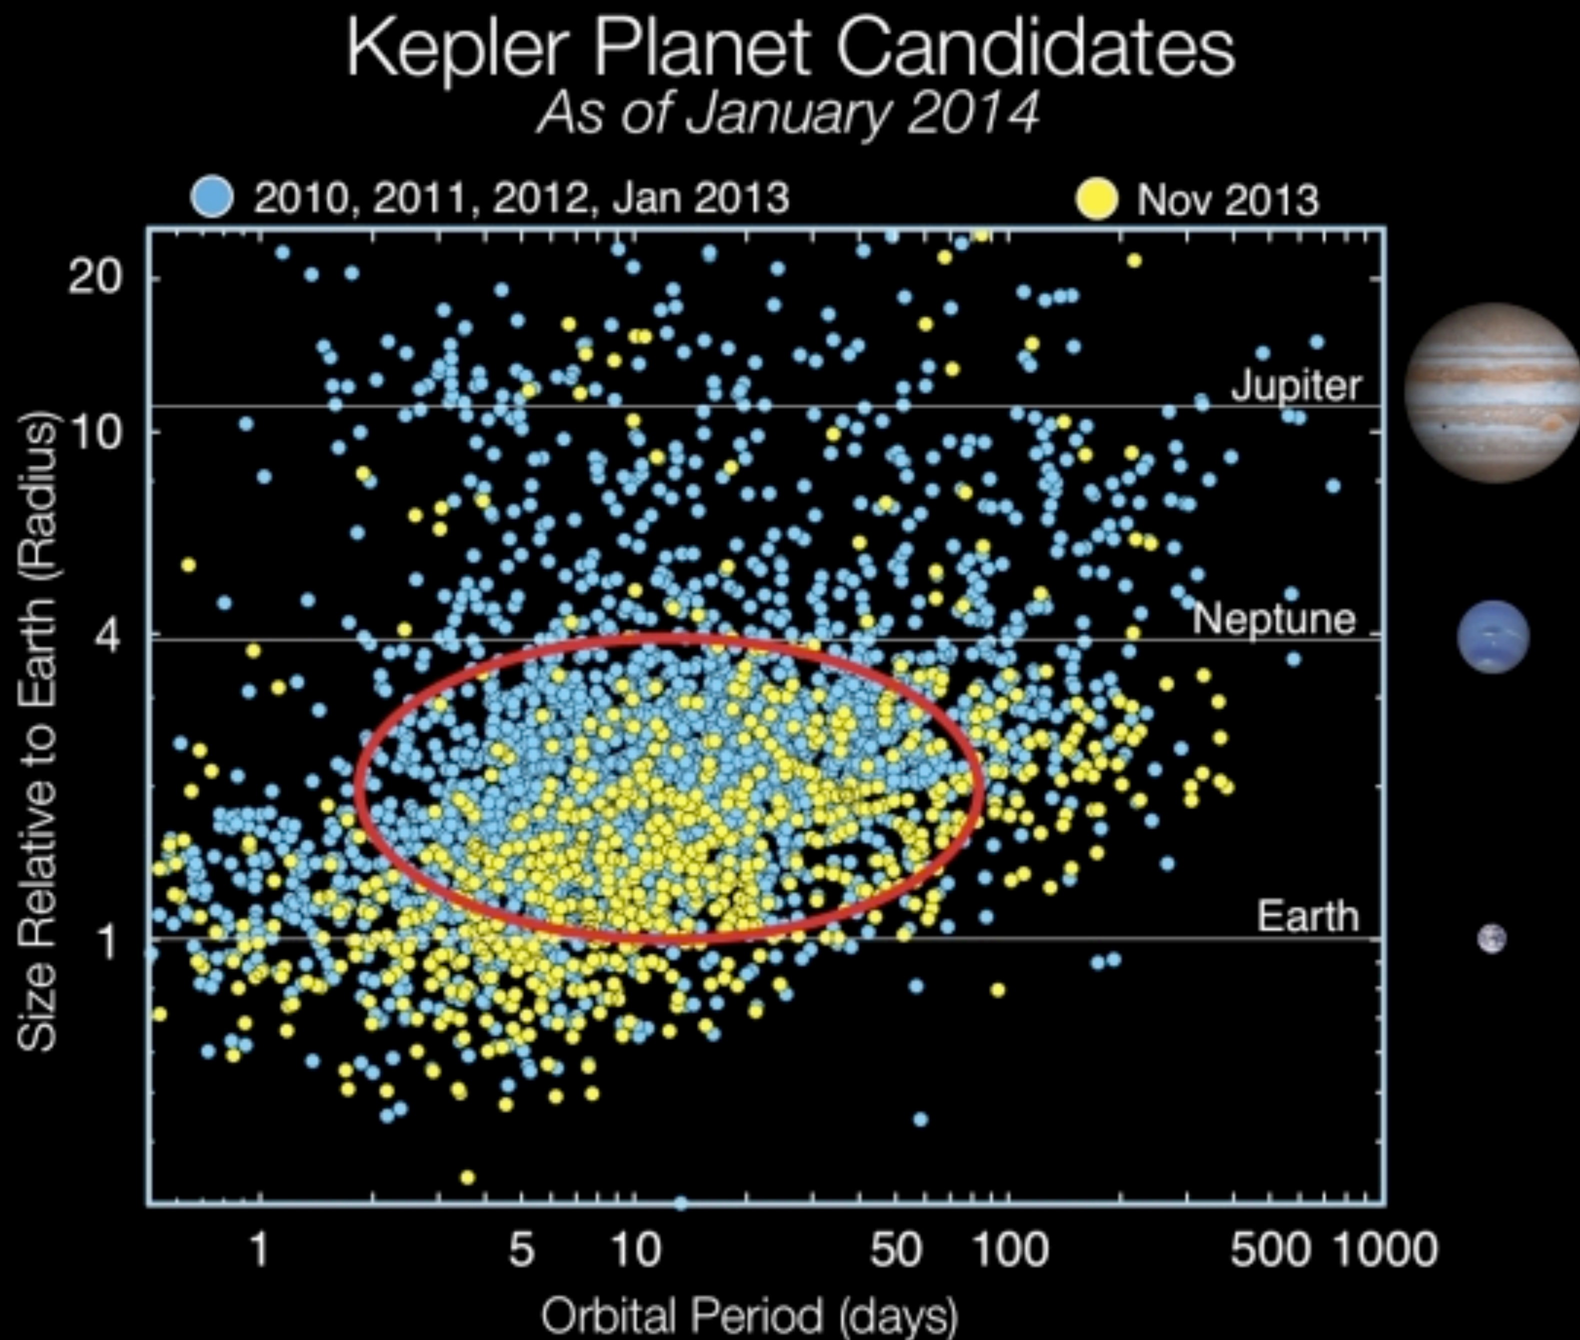

Planets with no Solar System analogs . . . what possible compositions?

From just Mass and Radius:

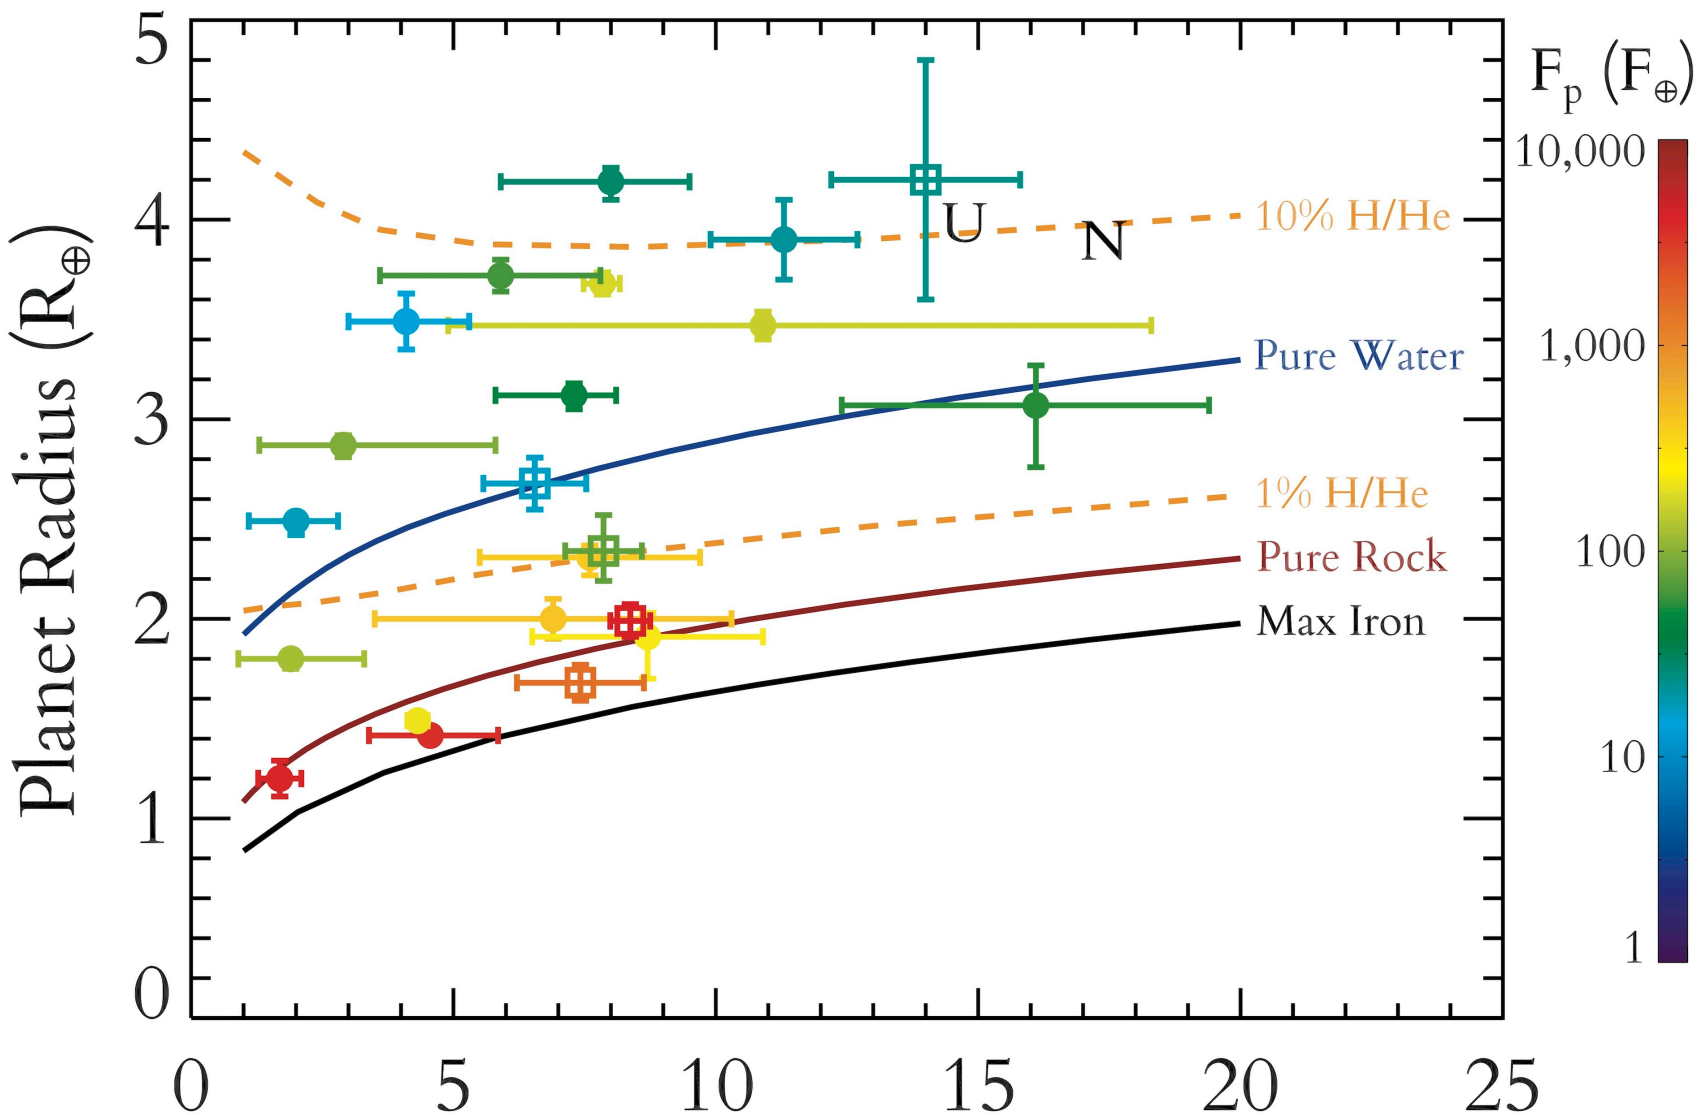

courtesy of Eric Lopez

From just Mass and Radius:

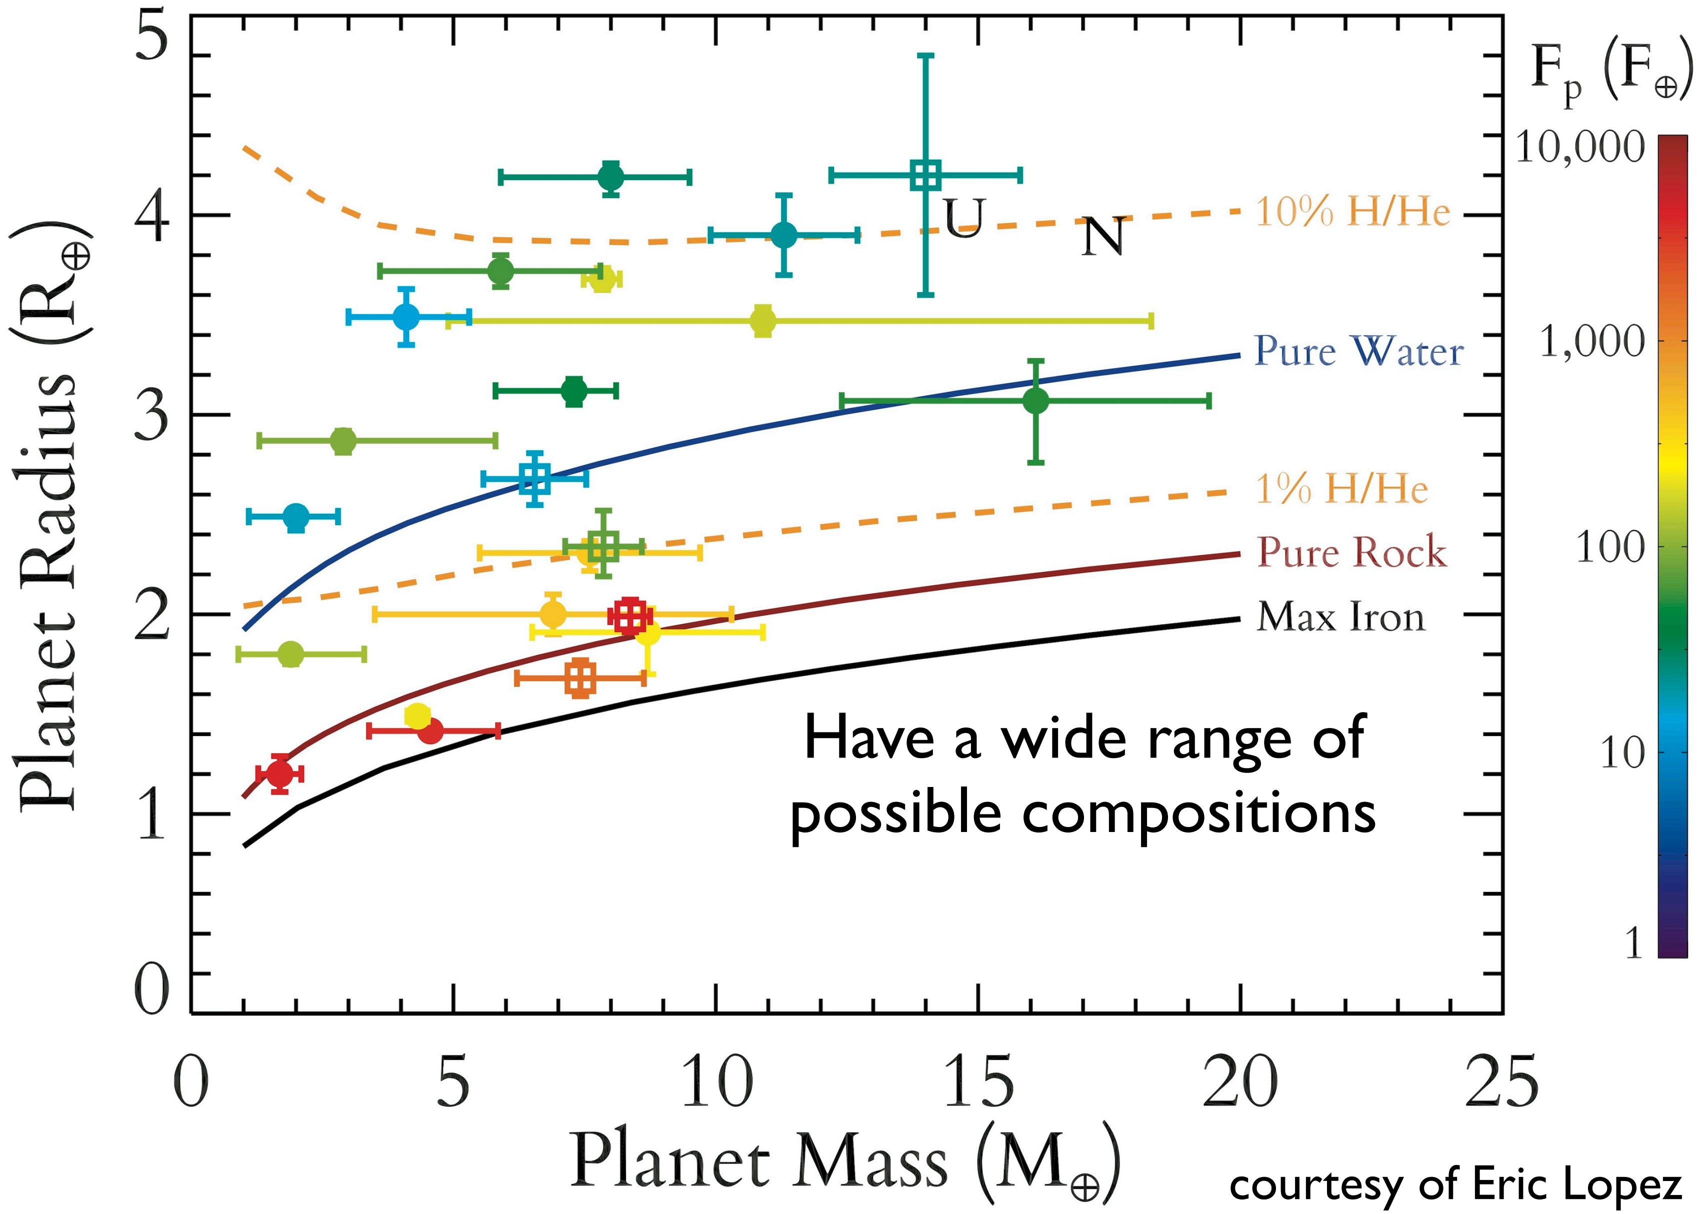

courtesy of Eric Lopez

From just Mass and Radius:

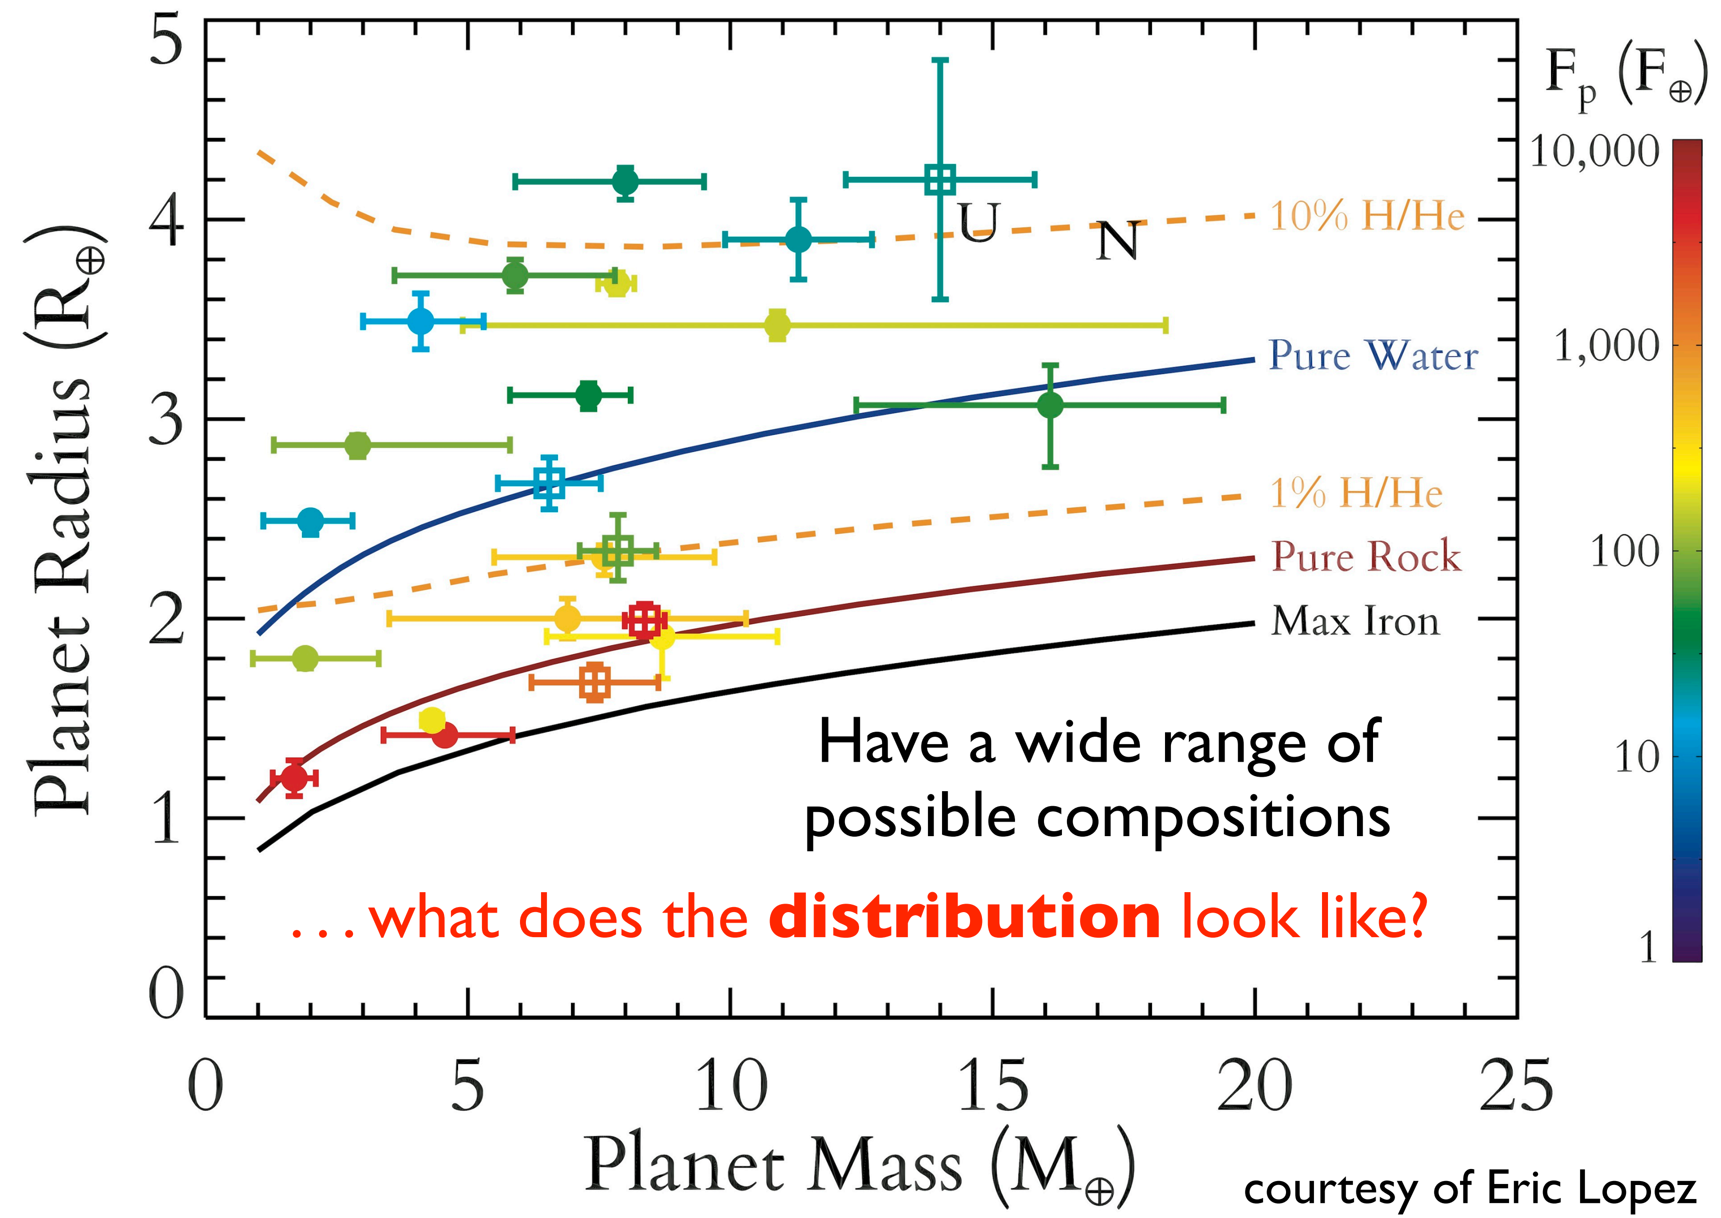

Have a wide range of possible compositions

…what does the **distribution** look like?

courtesy of Eric Lopez

From just Mass and Radius:

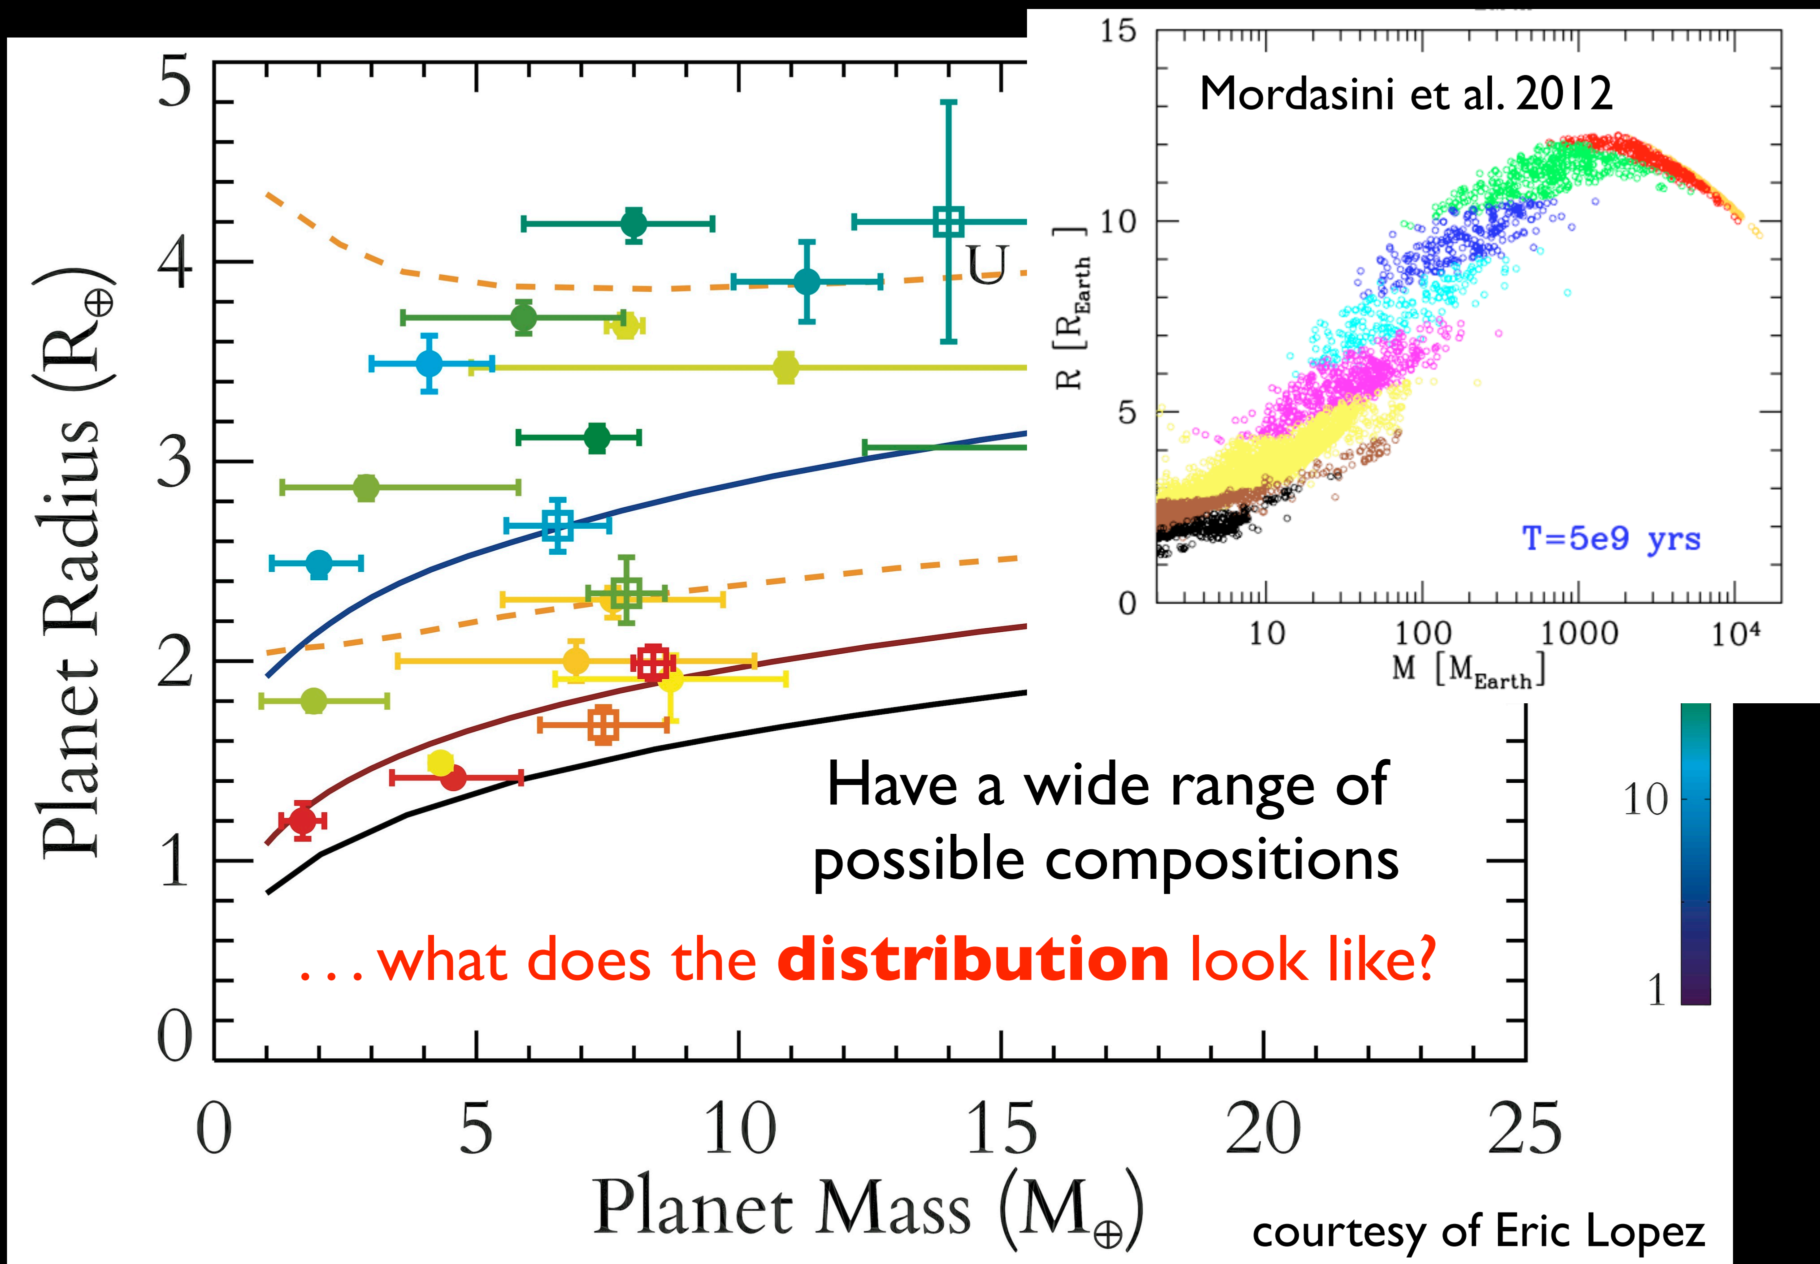

Have a wide range of possible compositions

. . . what does the **distribution** look like?

courtesy of Eric Lopez

From just Mass and Radius:

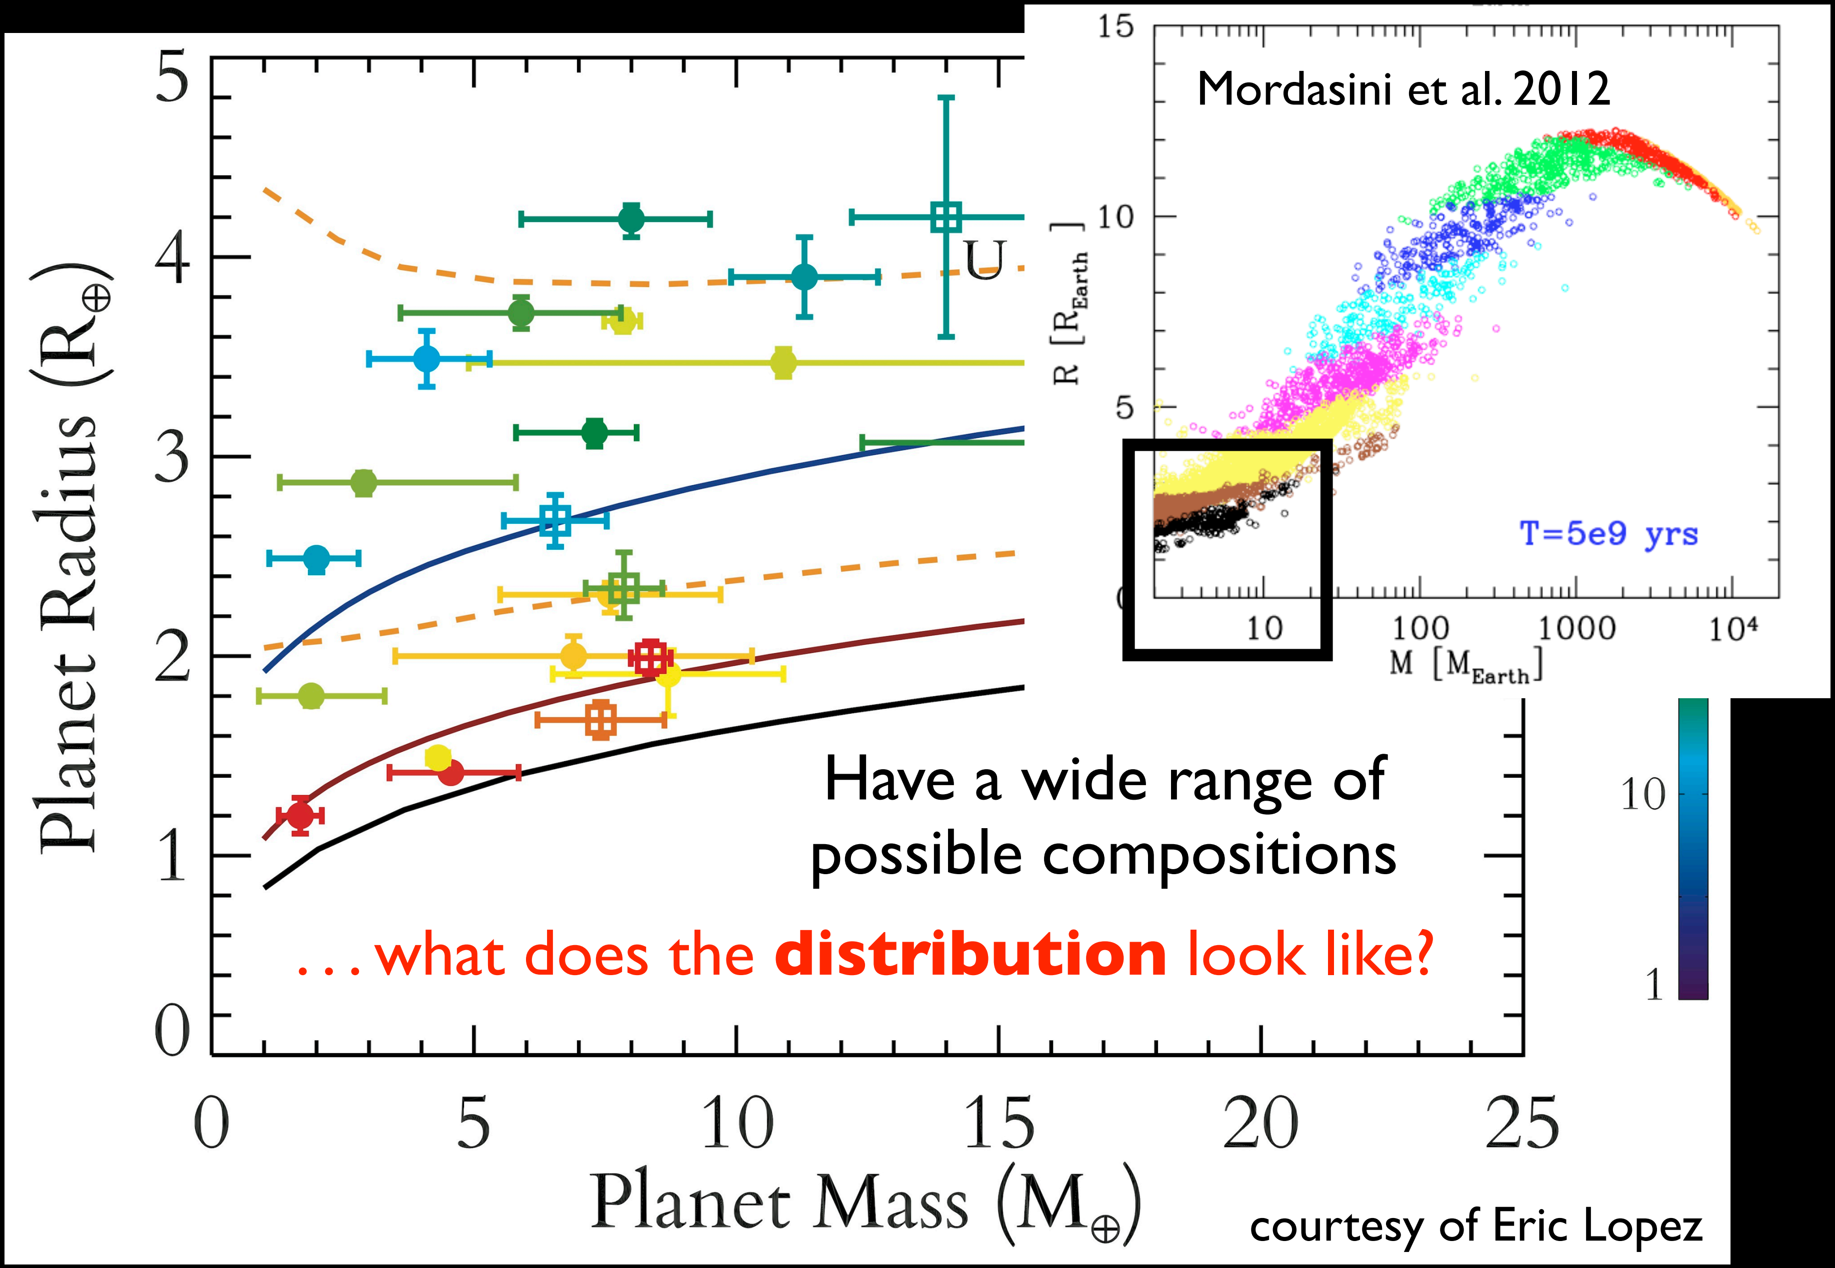

courtesy of Eric Lopez

Flat Mass-Radius Relations!

Earth-composition rocky core, H+He envelope

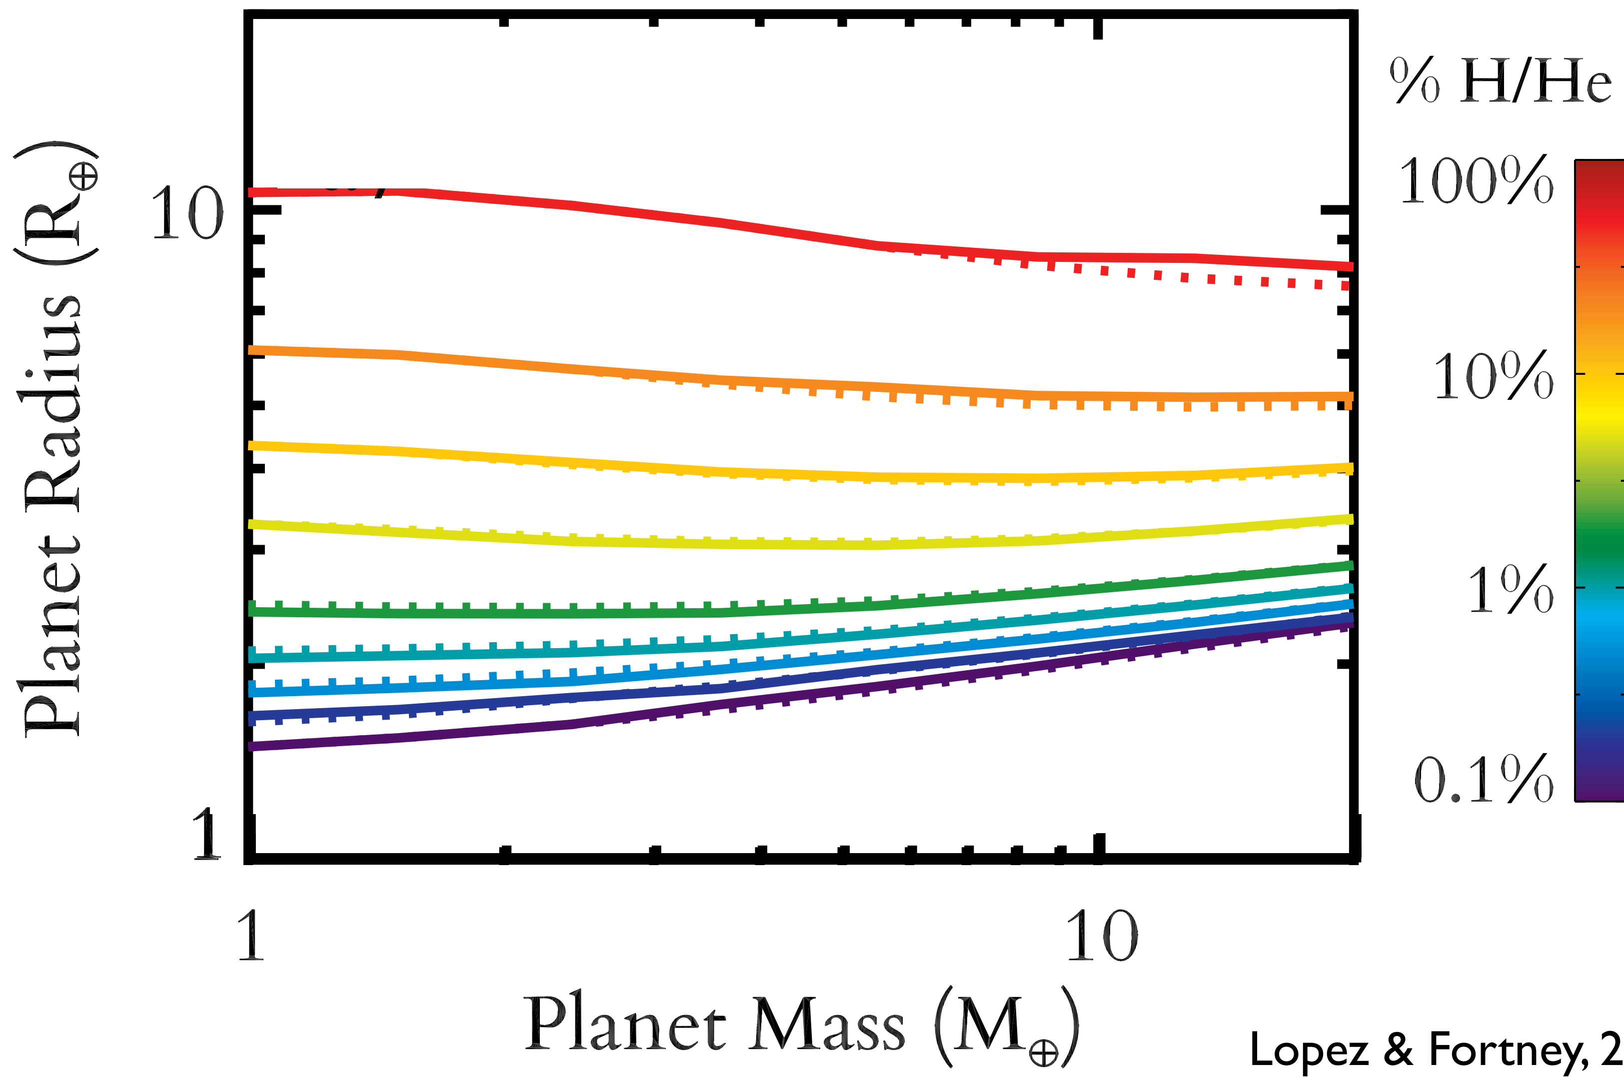

Lopez & Fortney, 2014

Flat Mass-Radius Relations!

Earth-composition rocky core, H+He envelope

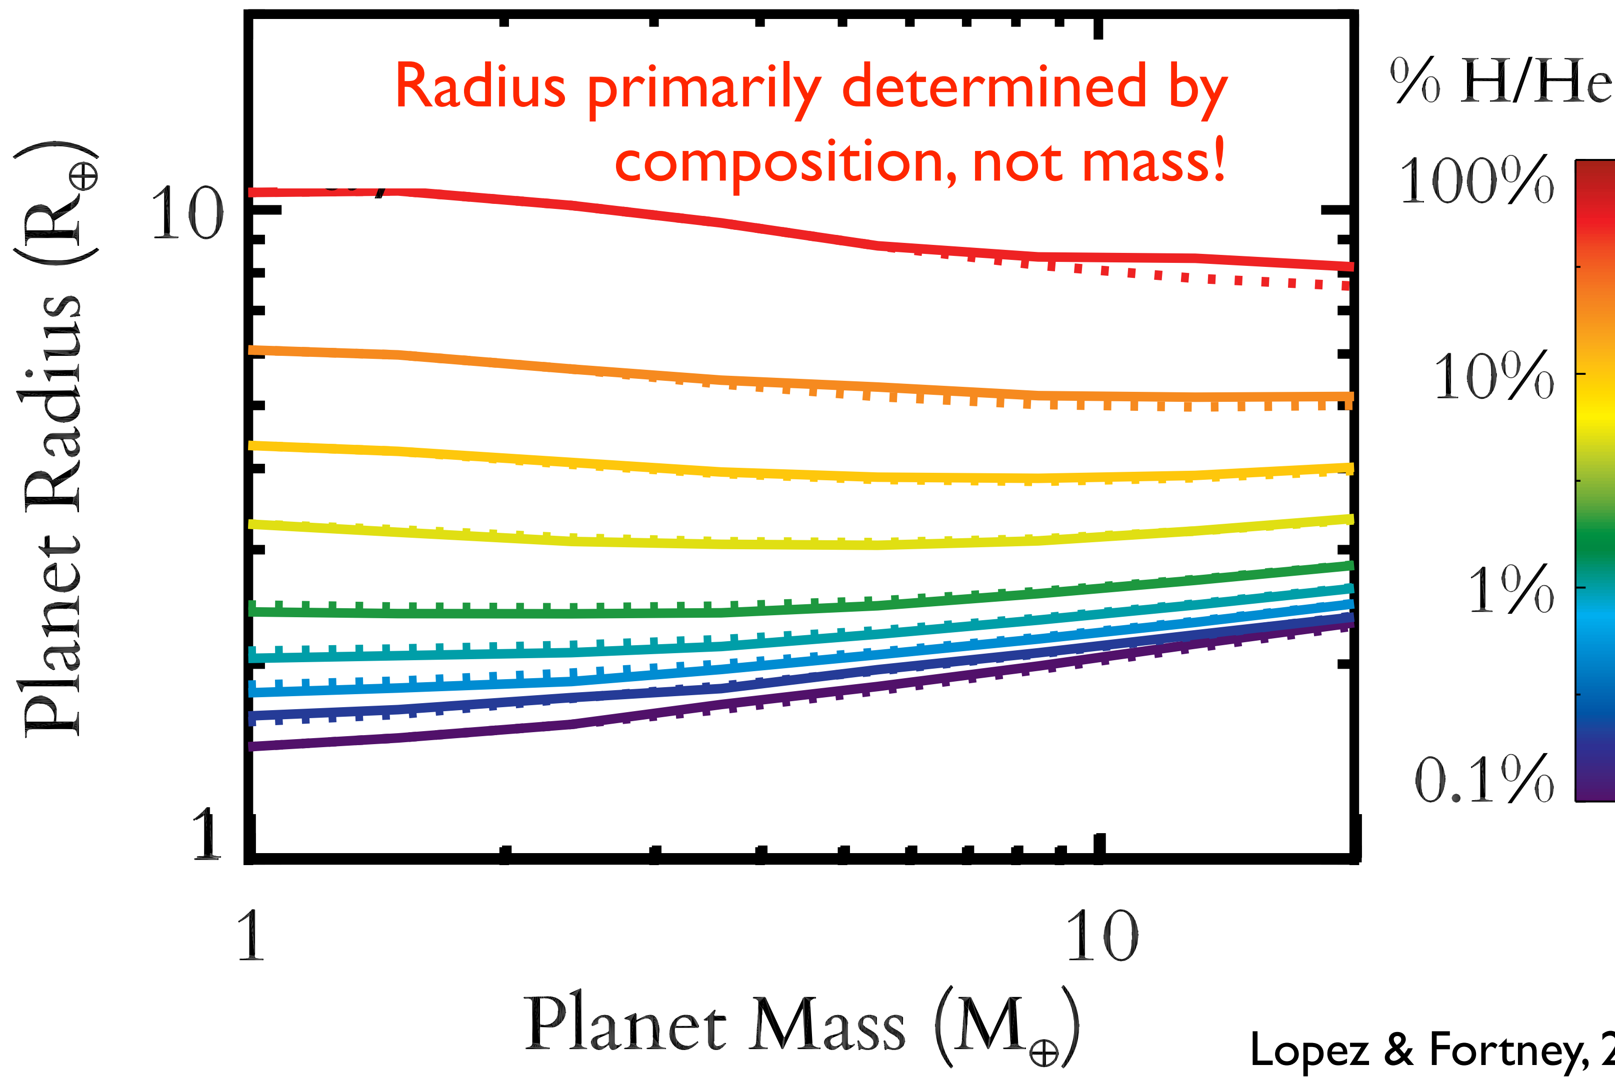

Lopez & Fortney, 2014

Radius Proxy for Composition

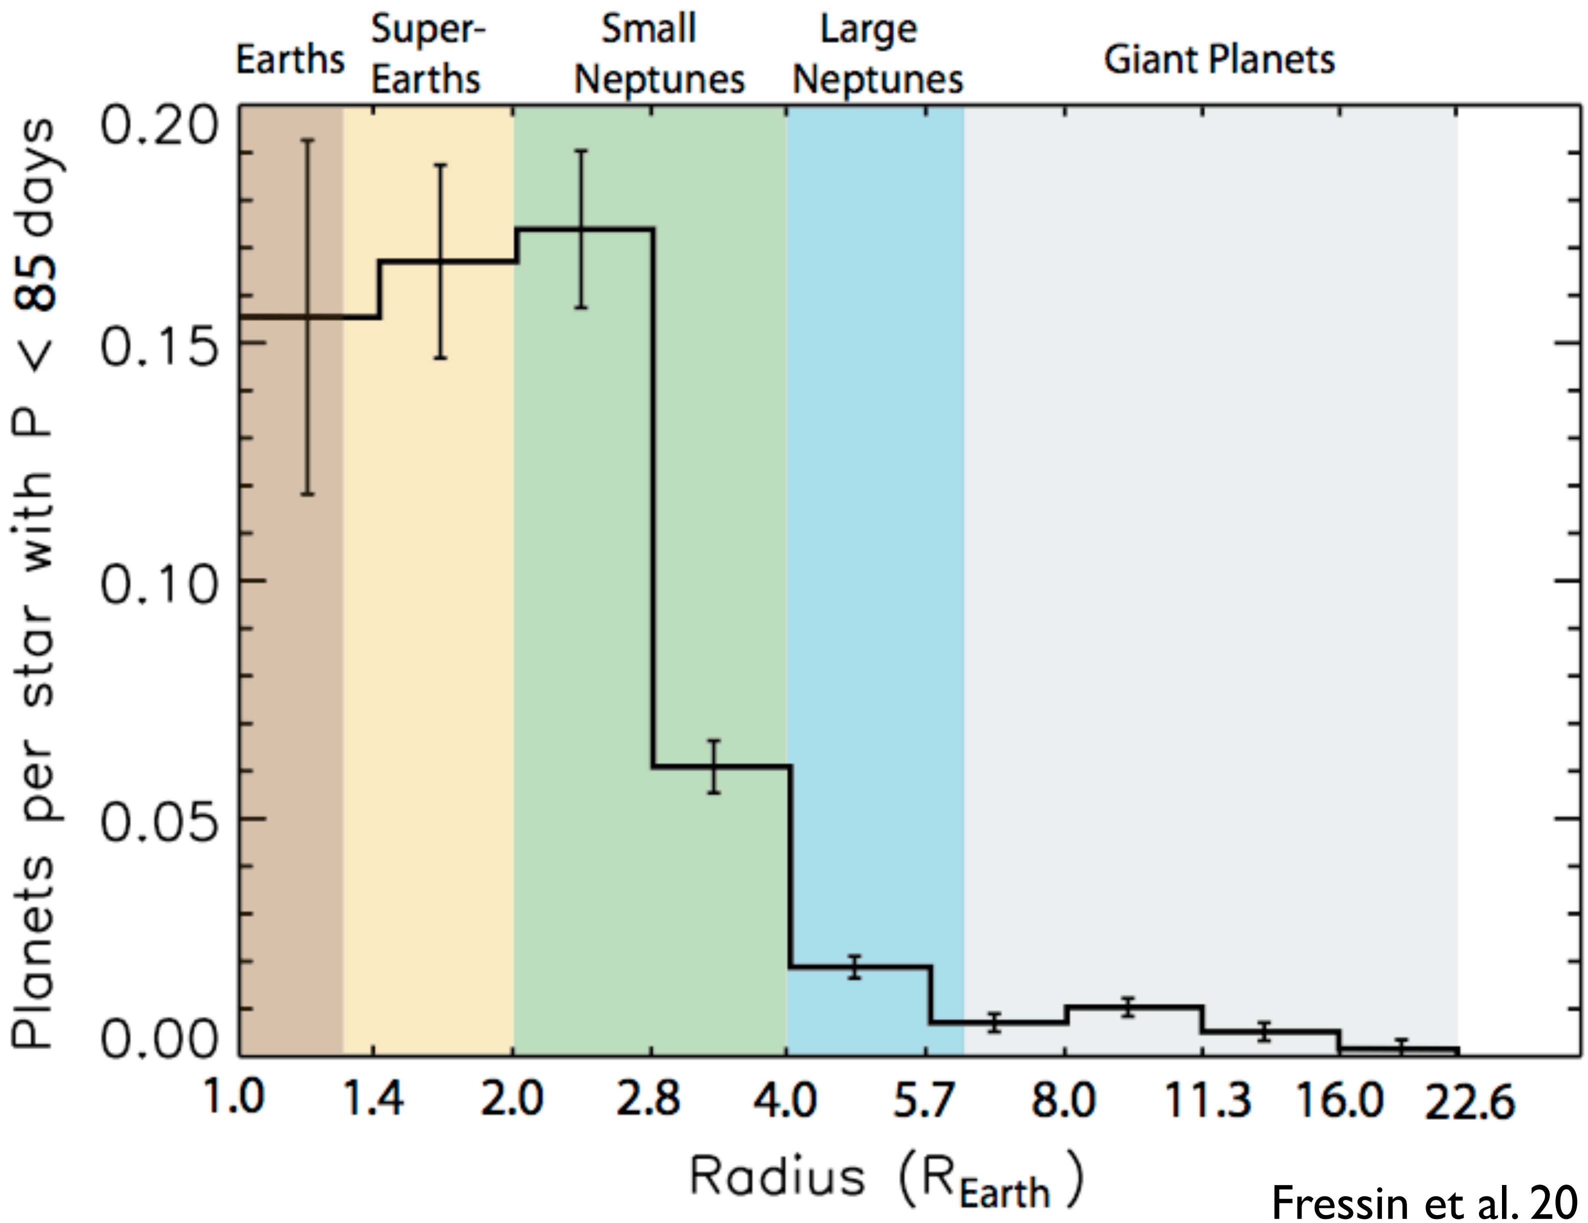

Fressin et al. 2013

Radius Proxy for Composition

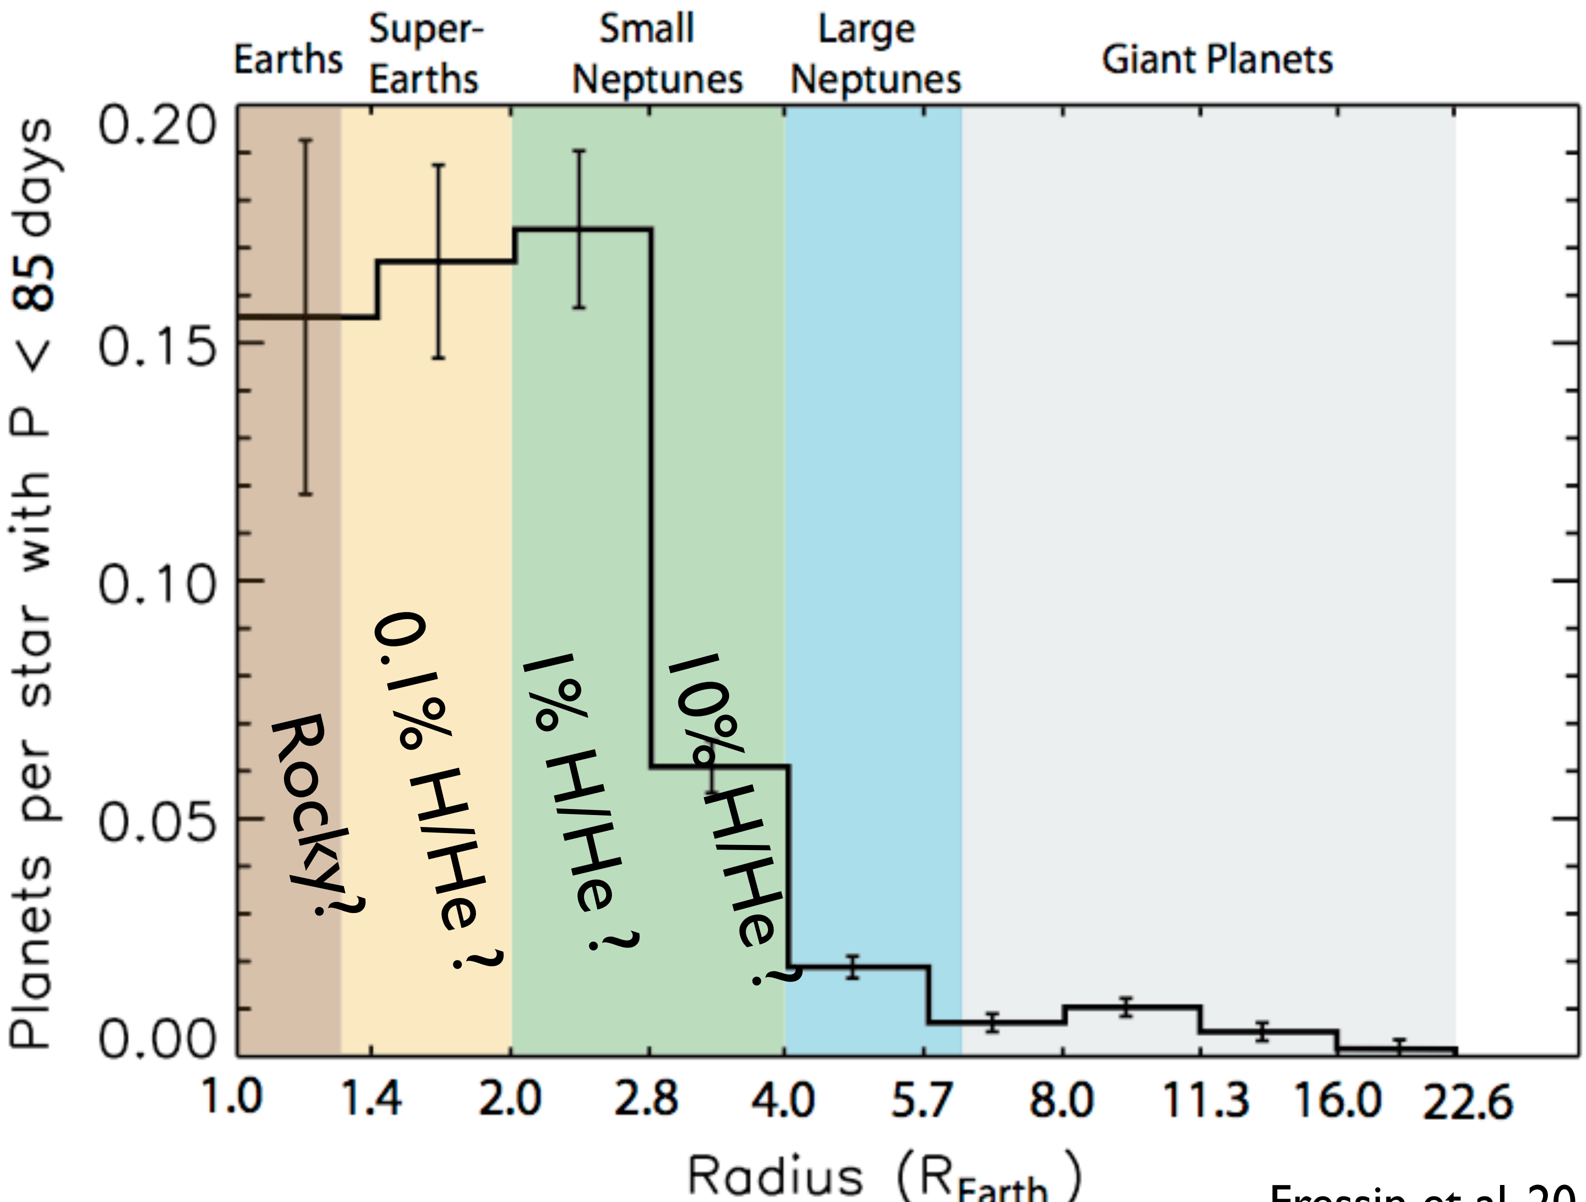

Fressin et al. 2013

Radius Proxy for Composition

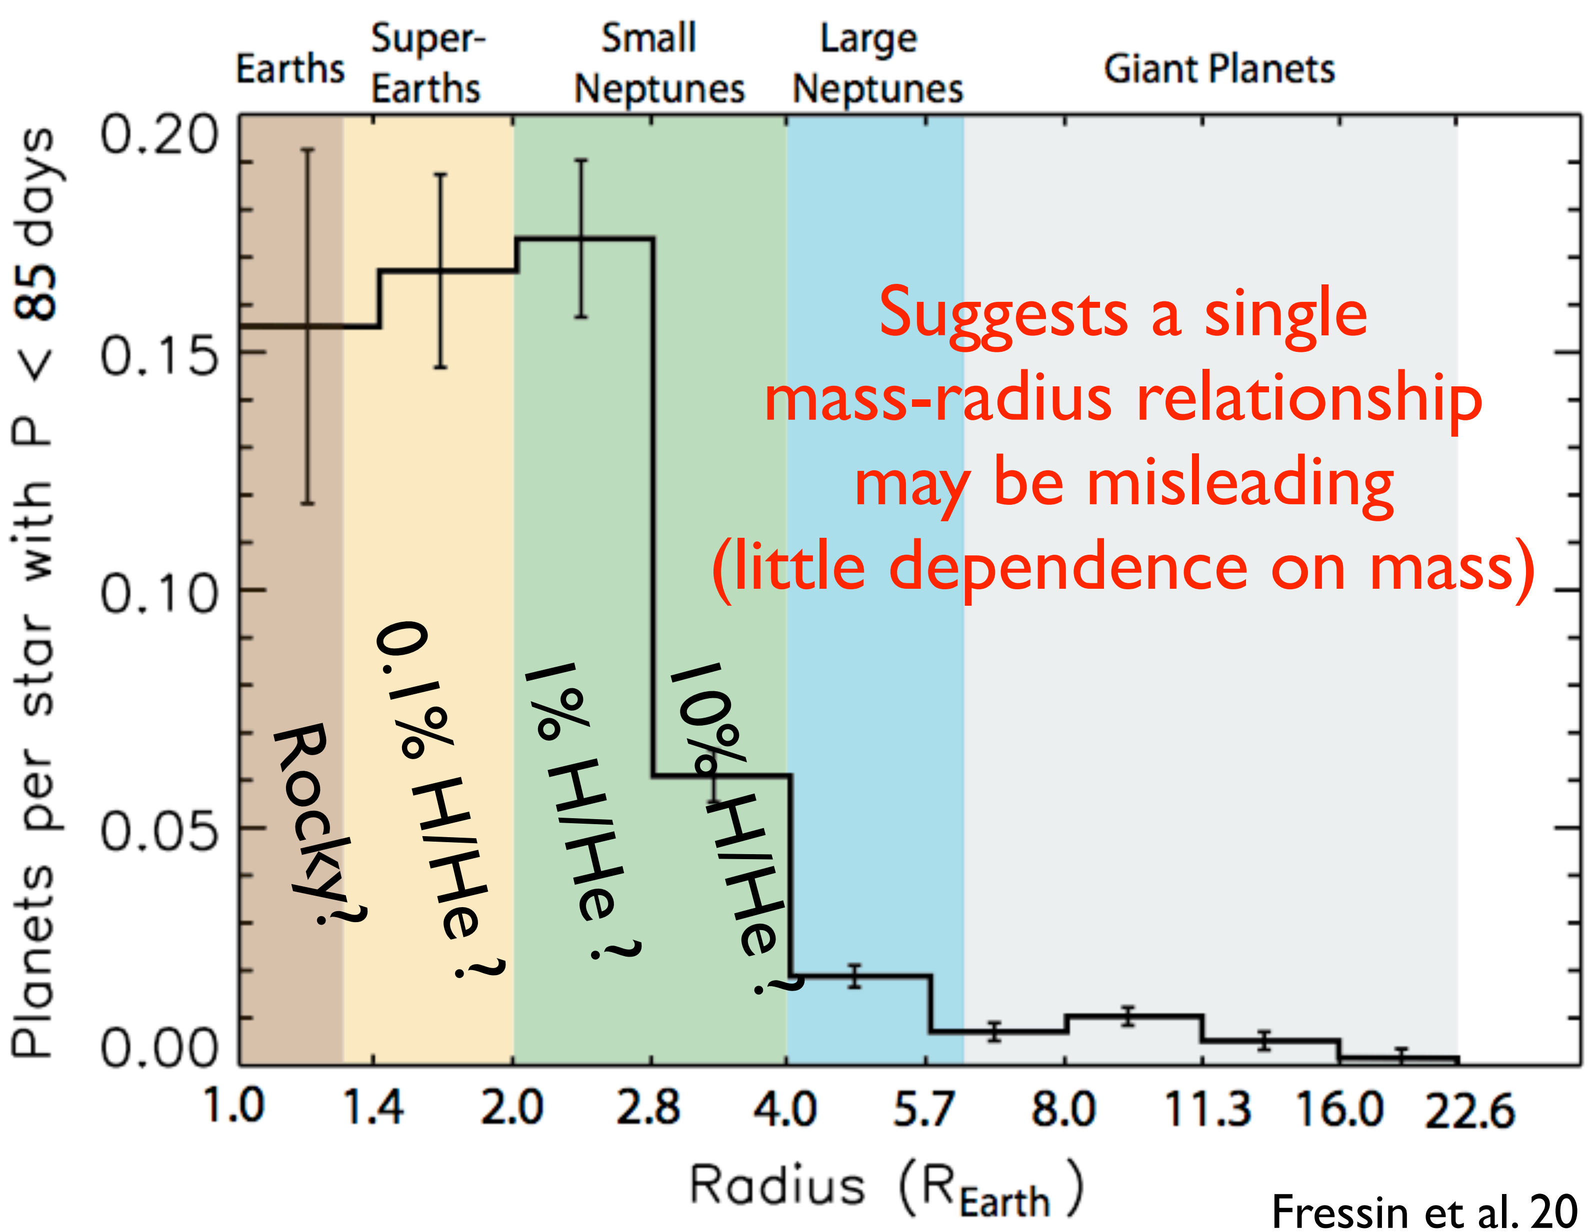

Fressin et al. 2013

... But Inferences are Hard

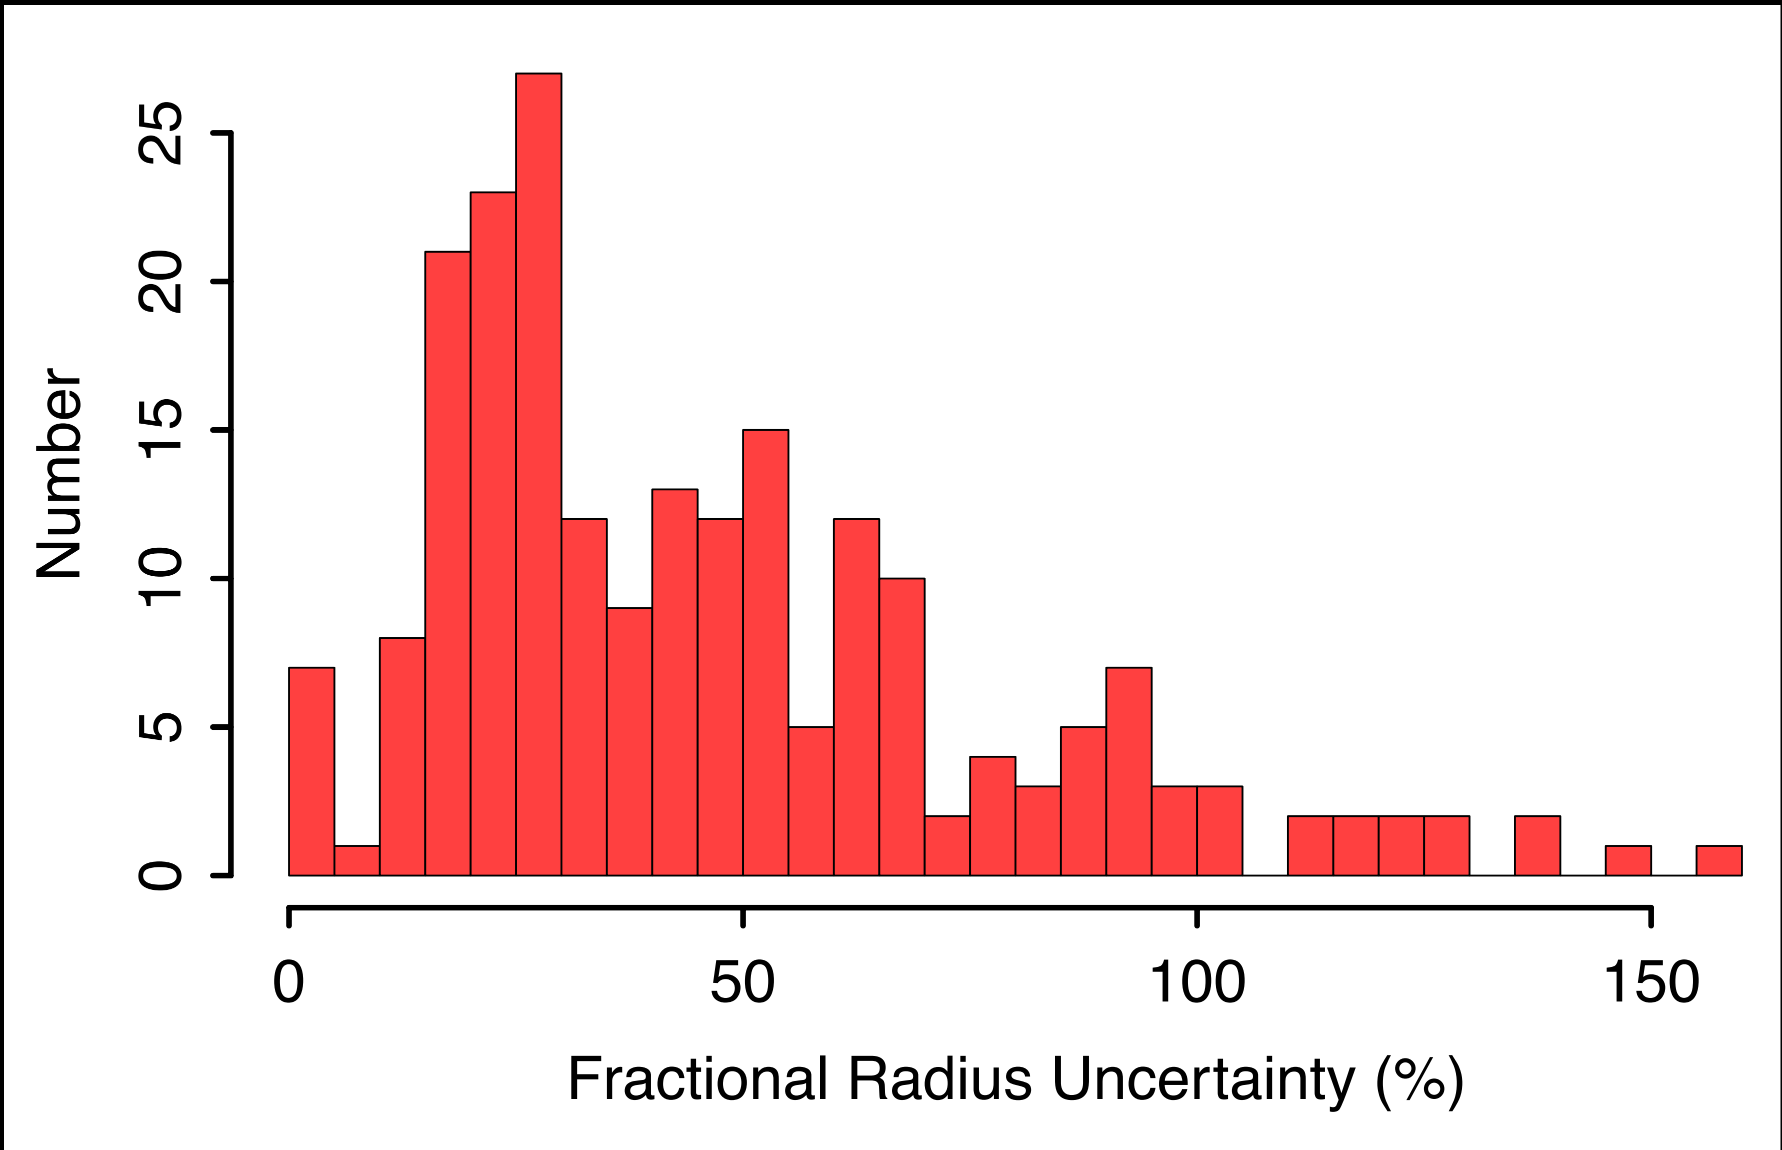

...But Inferences are Hard

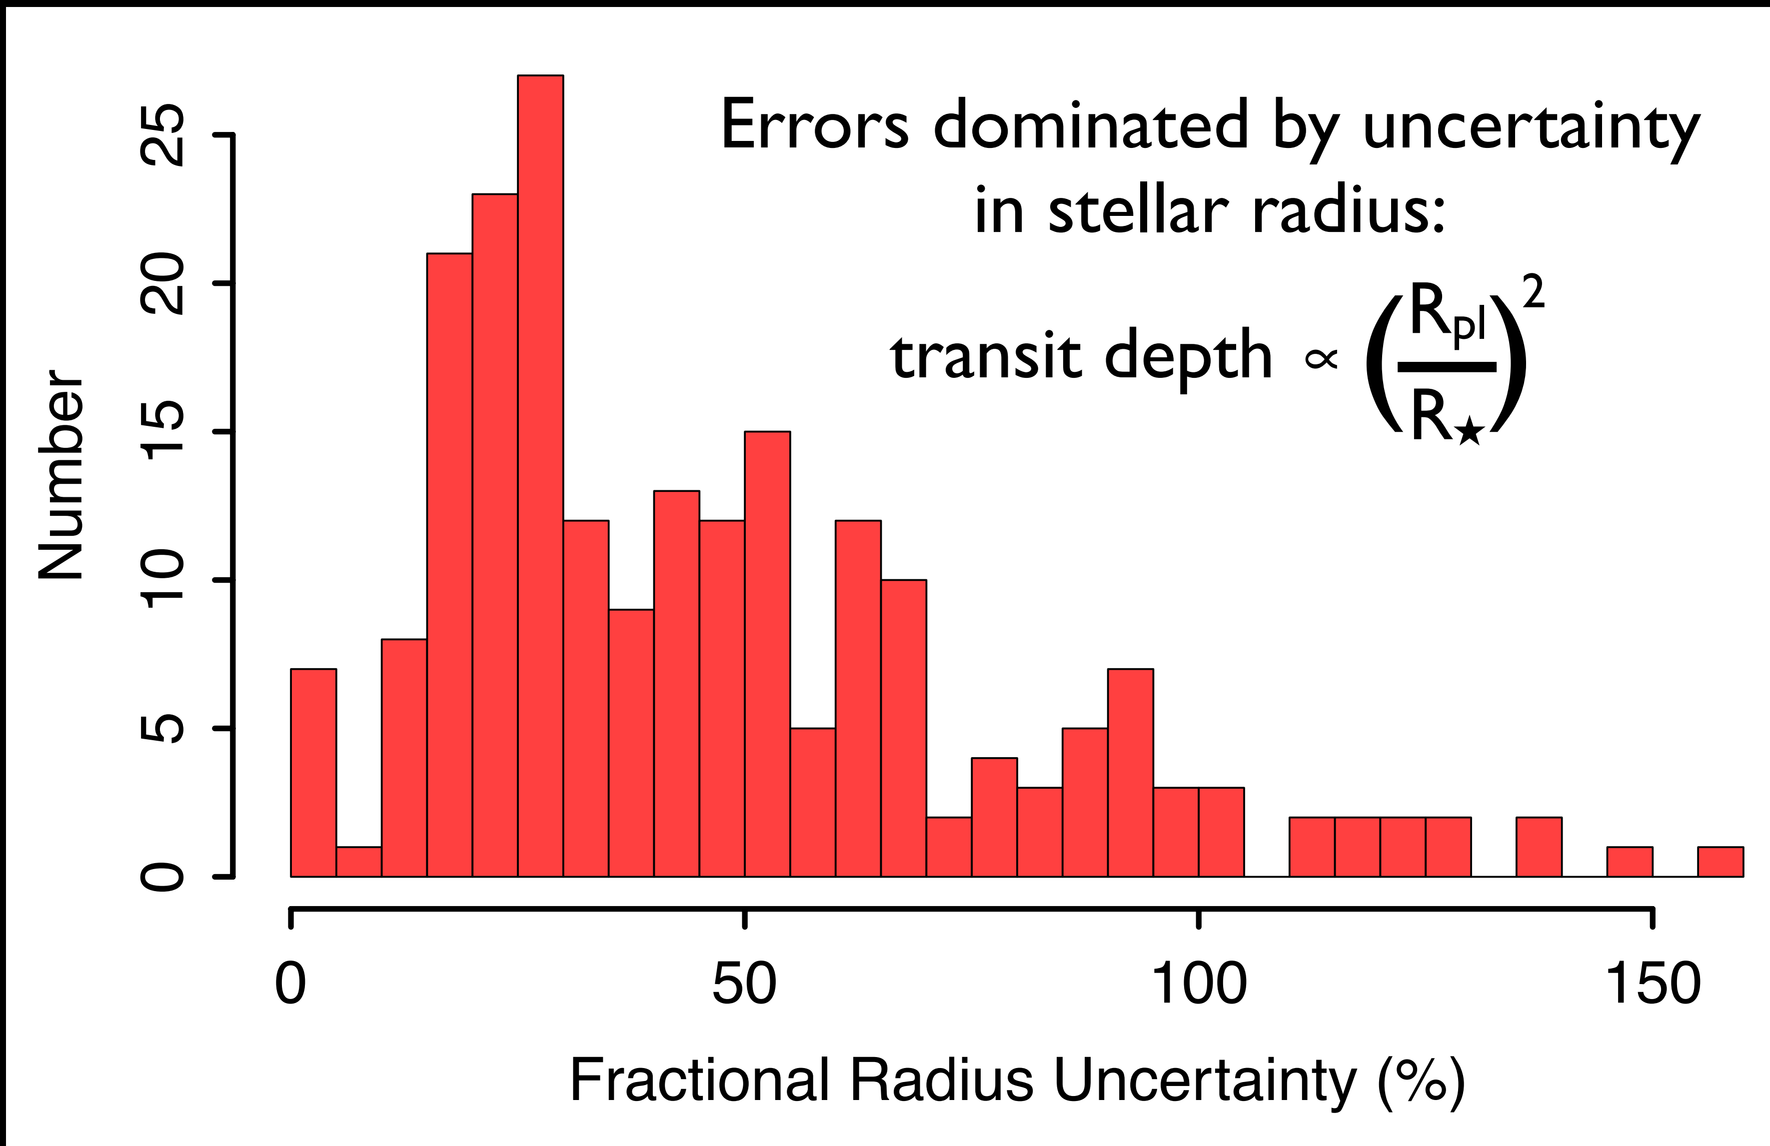

Hierarchical Bayesian Model

Hierarchical Bayesian Model

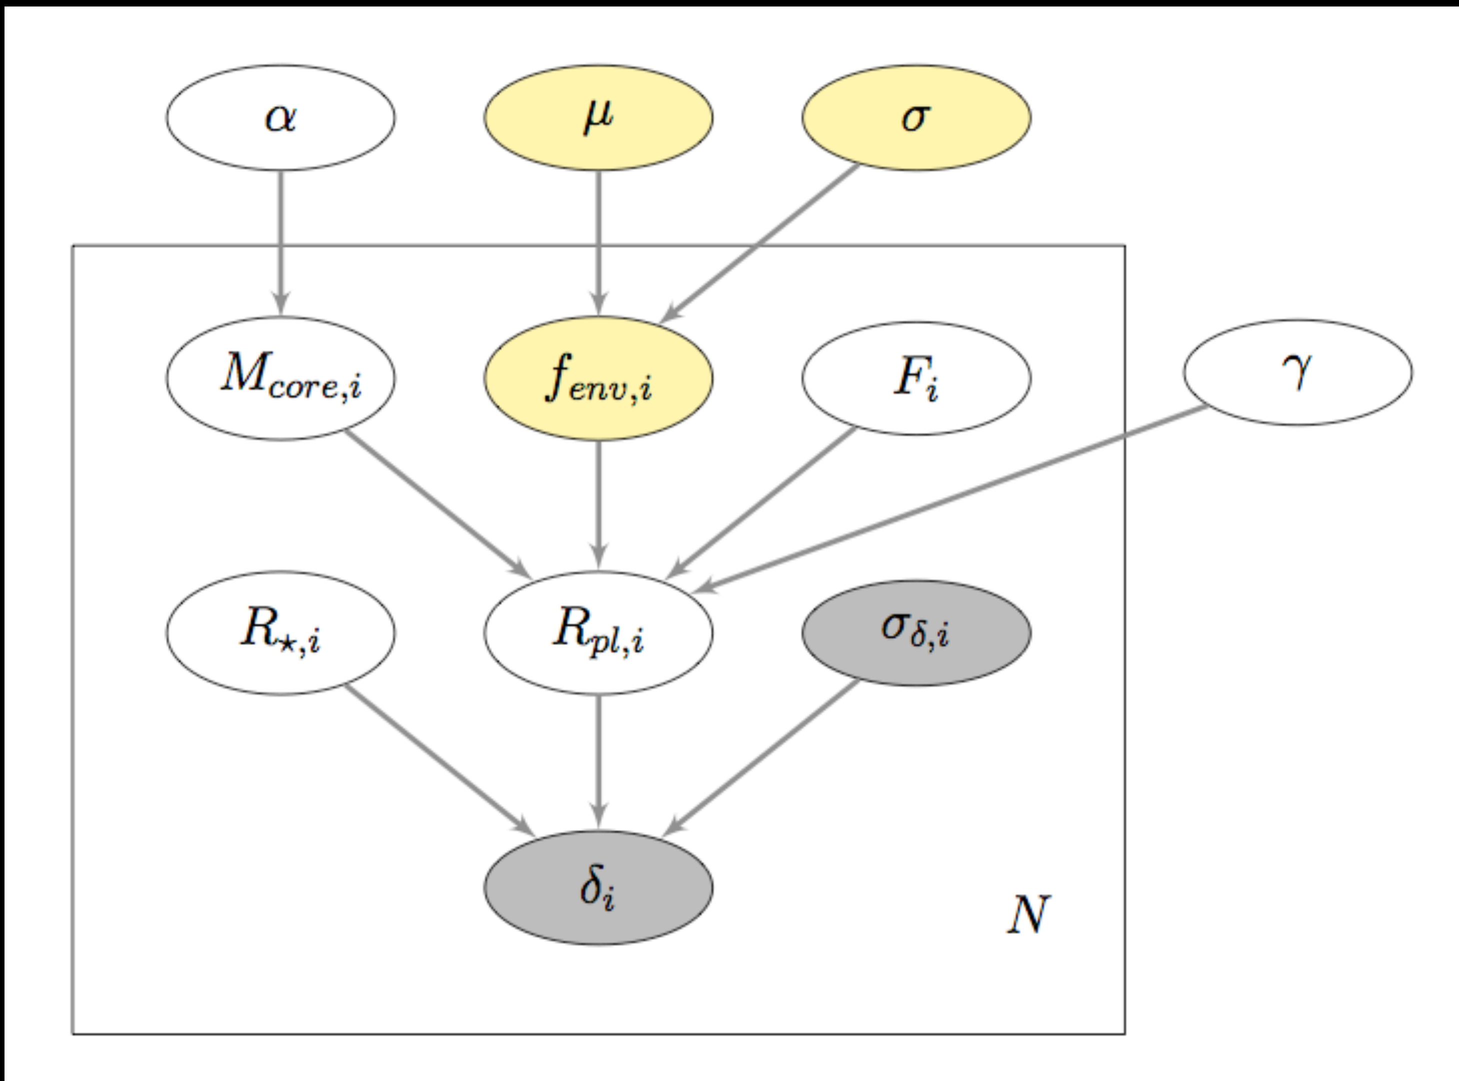

Hierarchical Bayesian Model

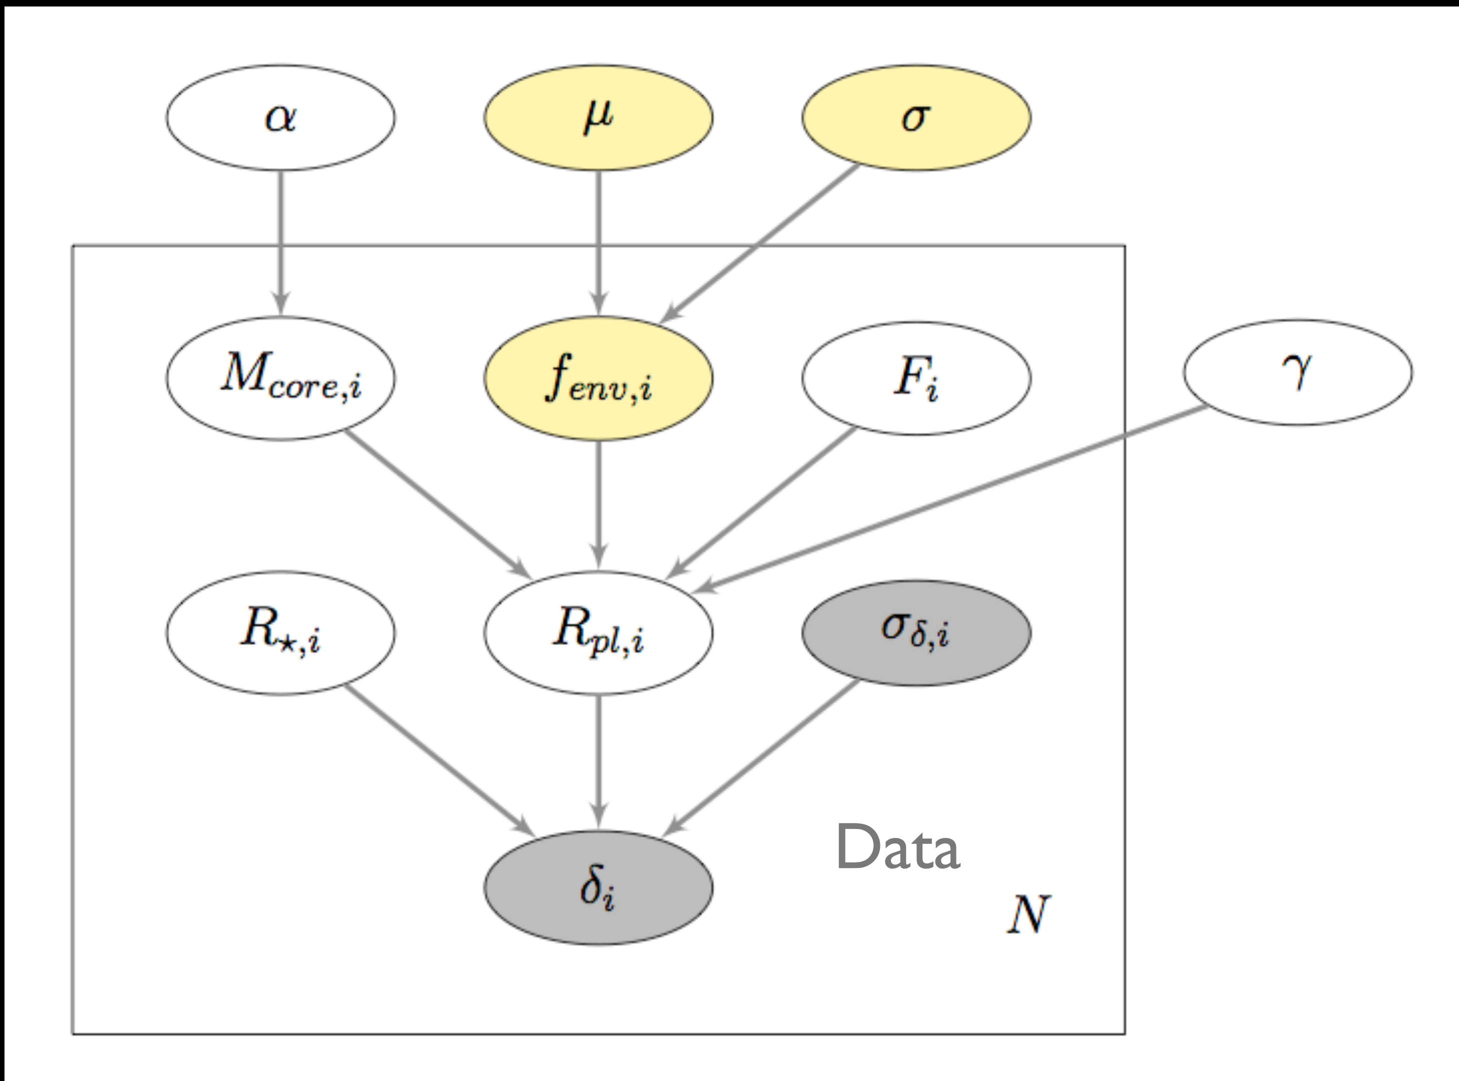

Hierarchical Bayesian Model

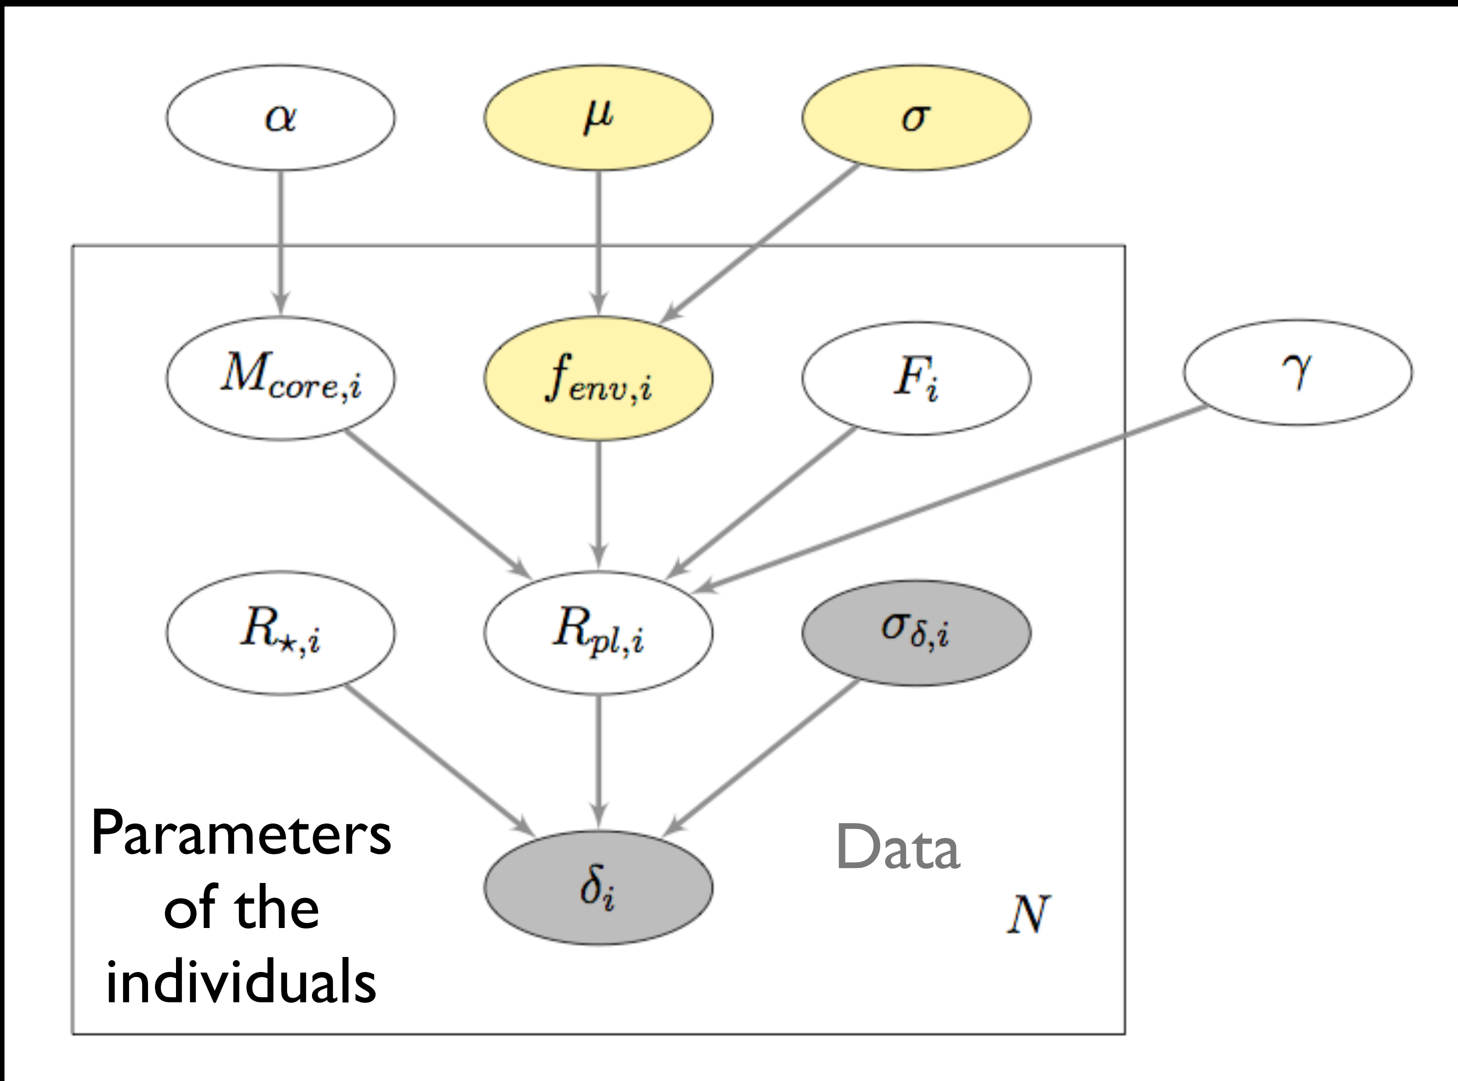

Hierarchical Bayesian Model

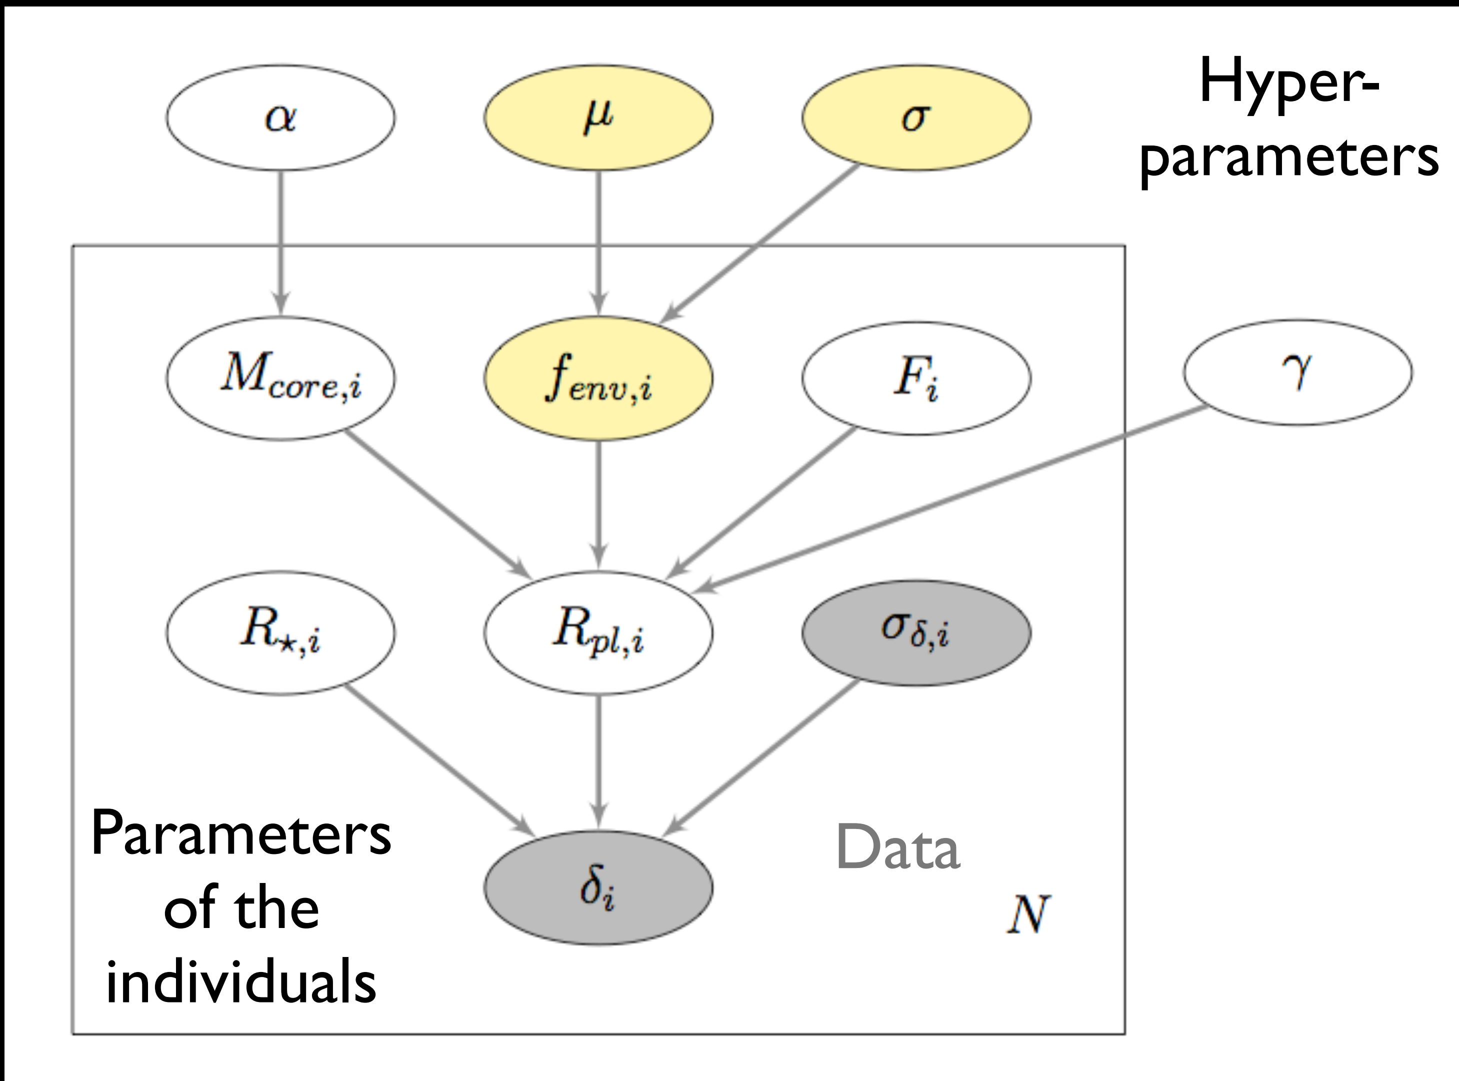

Hierarchical Bayesian Model

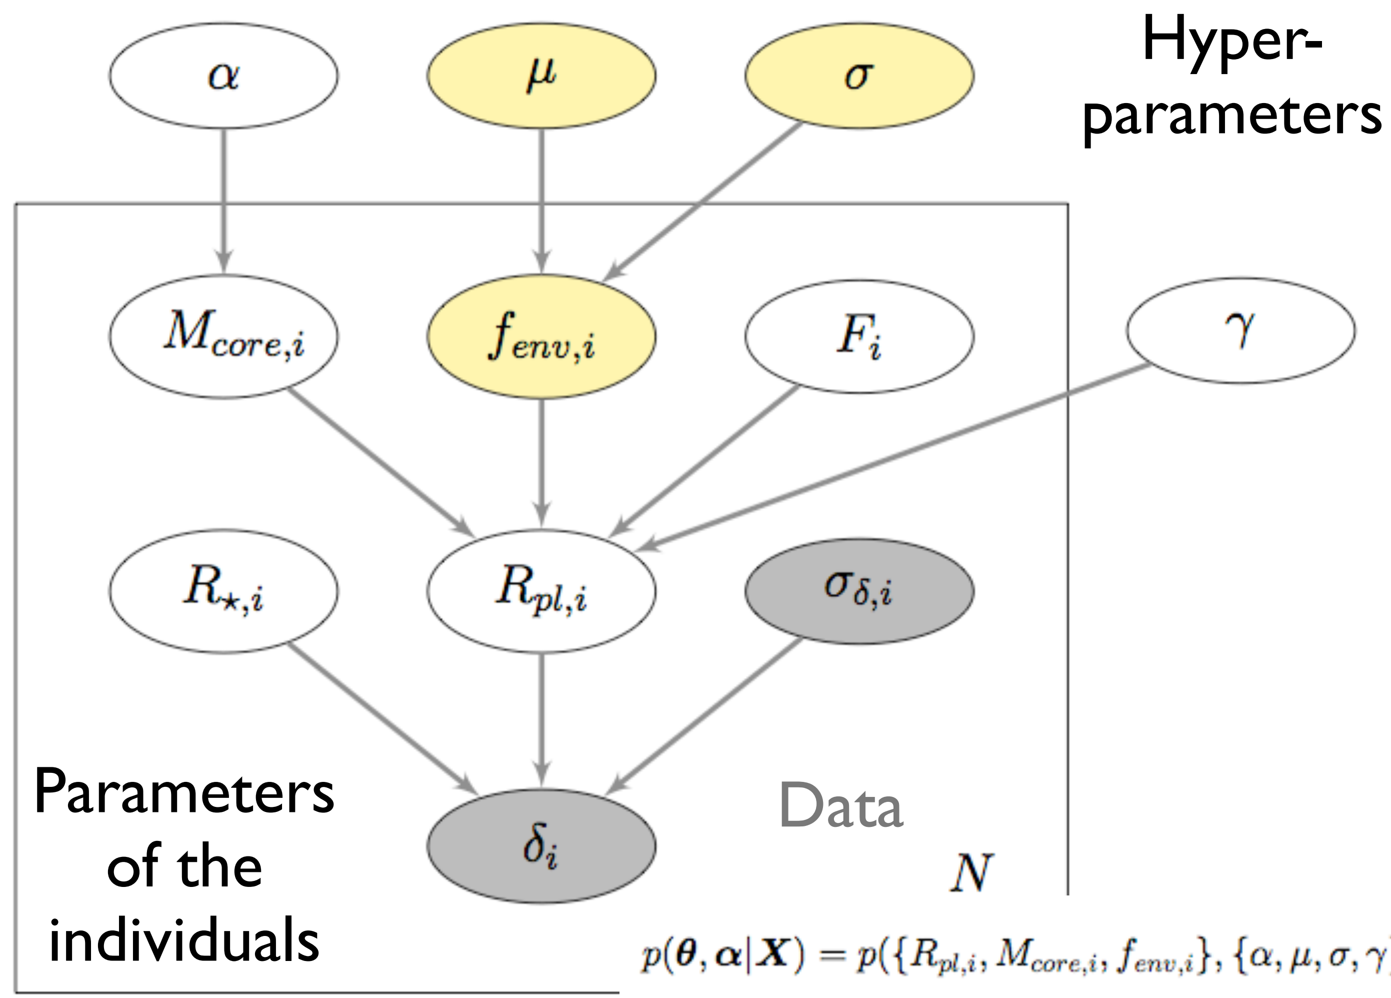

$$p(\boldsymbol{\theta}, \boldsymbol{\alpha}|\boldsymbol{X}) = p(\{R_{pl,i}, M_{core,i}, f_{env,i}\}, \{\alpha, \mu, \sigma, \gamma\}|\delta_i, \sigma_{\delta,i}, F_i) \propto$$

$$\prod_{i=1}^{N} \left\{ p(\delta_i|\sigma_{\delta,i}, R_{pl,i}, R_{\star,i}, M_{core,i}, f_{env,i}, F_i, \alpha, \mu, \sigma, \gamma) \right\}$$

$$\times \prod_{i=1}^{N} \left\{ p(R_{\star,i}) p(M_{pl,i}|\alpha) p(f_{env,i}|\mu, \sigma) \right\} p(\alpha) p(\mu) p(\sigma) p(\gamma)$$

Hierarchical Bayesian Model

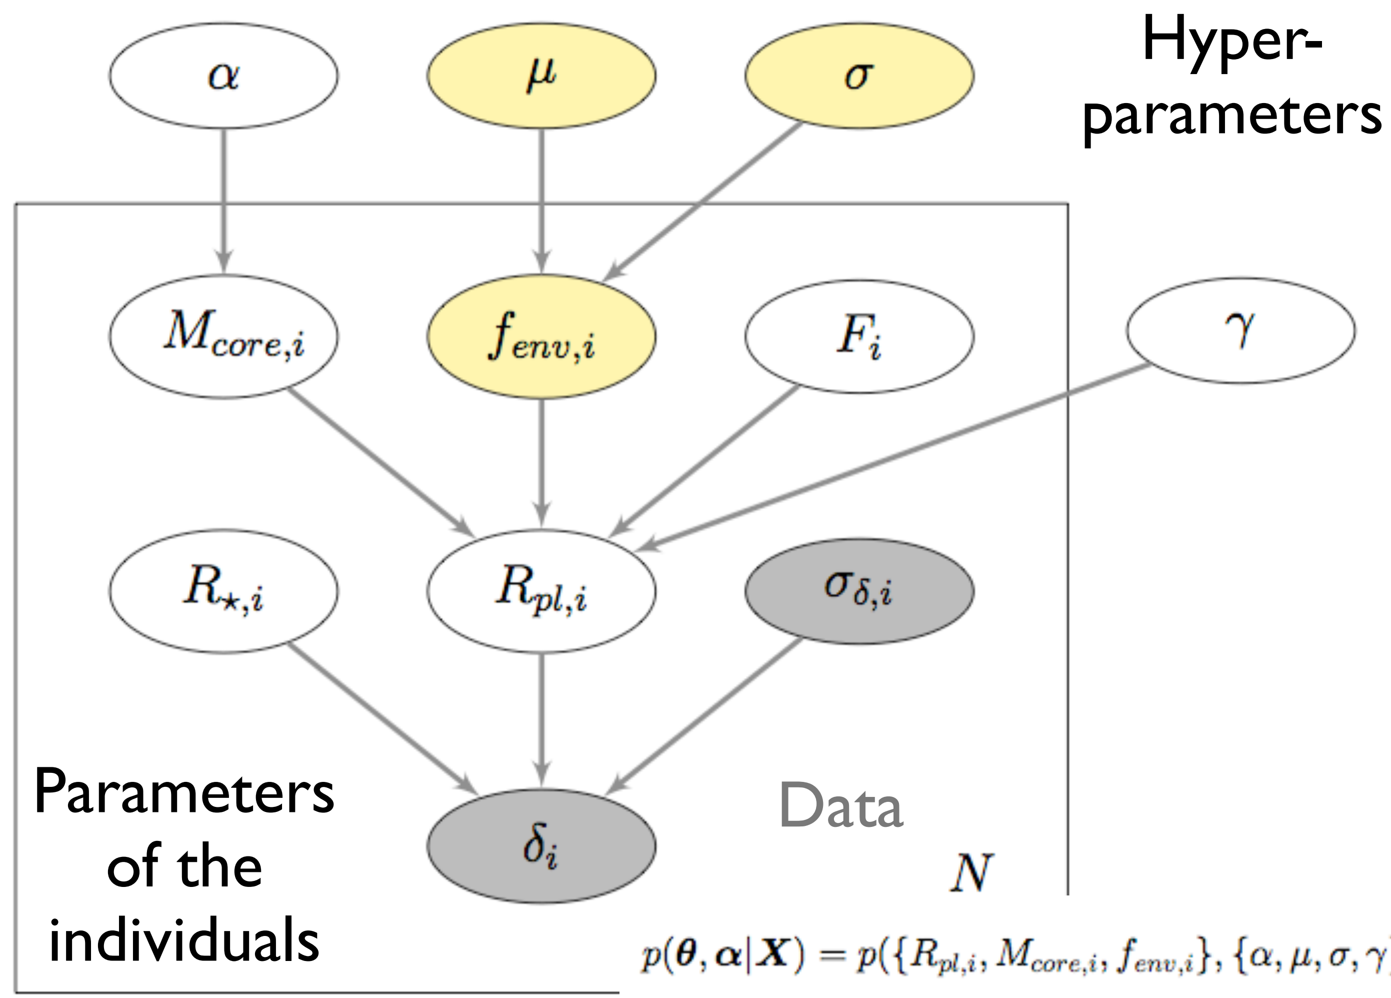

We are interested in the parameters defining planet compositions.

$$
\begin{aligned}
p(\boldsymbol{\theta}, \boldsymbol{\alpha}|\boldsymbol{X}) = p(\{R_{pl,i}, M_{core,i}, f_{env,i}\}, \{\alpha, \mu, \sigma, \gamma\}|\delta_i, \sigma_{\delta,i}, F_i) \propto \\
\prod_{i=1}^{N} \Big\{ p(\delta_i|\sigma_{\delta,i}, R_{pl,i}, R_{\star,i}, M_{core,i}, f_{env,i}, F_i, \alpha, \mu, \sigma, \gamma) \Big\} \\
\times \prod_{i=1}^{N} \Big\{ p(R_{\star,i}) p(M_{pl,i}|\alpha) p(f_{env,i}|\mu, \sigma) \Big\} p(\alpha) p(\mu) p(\sigma) p(\gamma)
\end{aligned}
$$

Hierarchical Bayesian Model

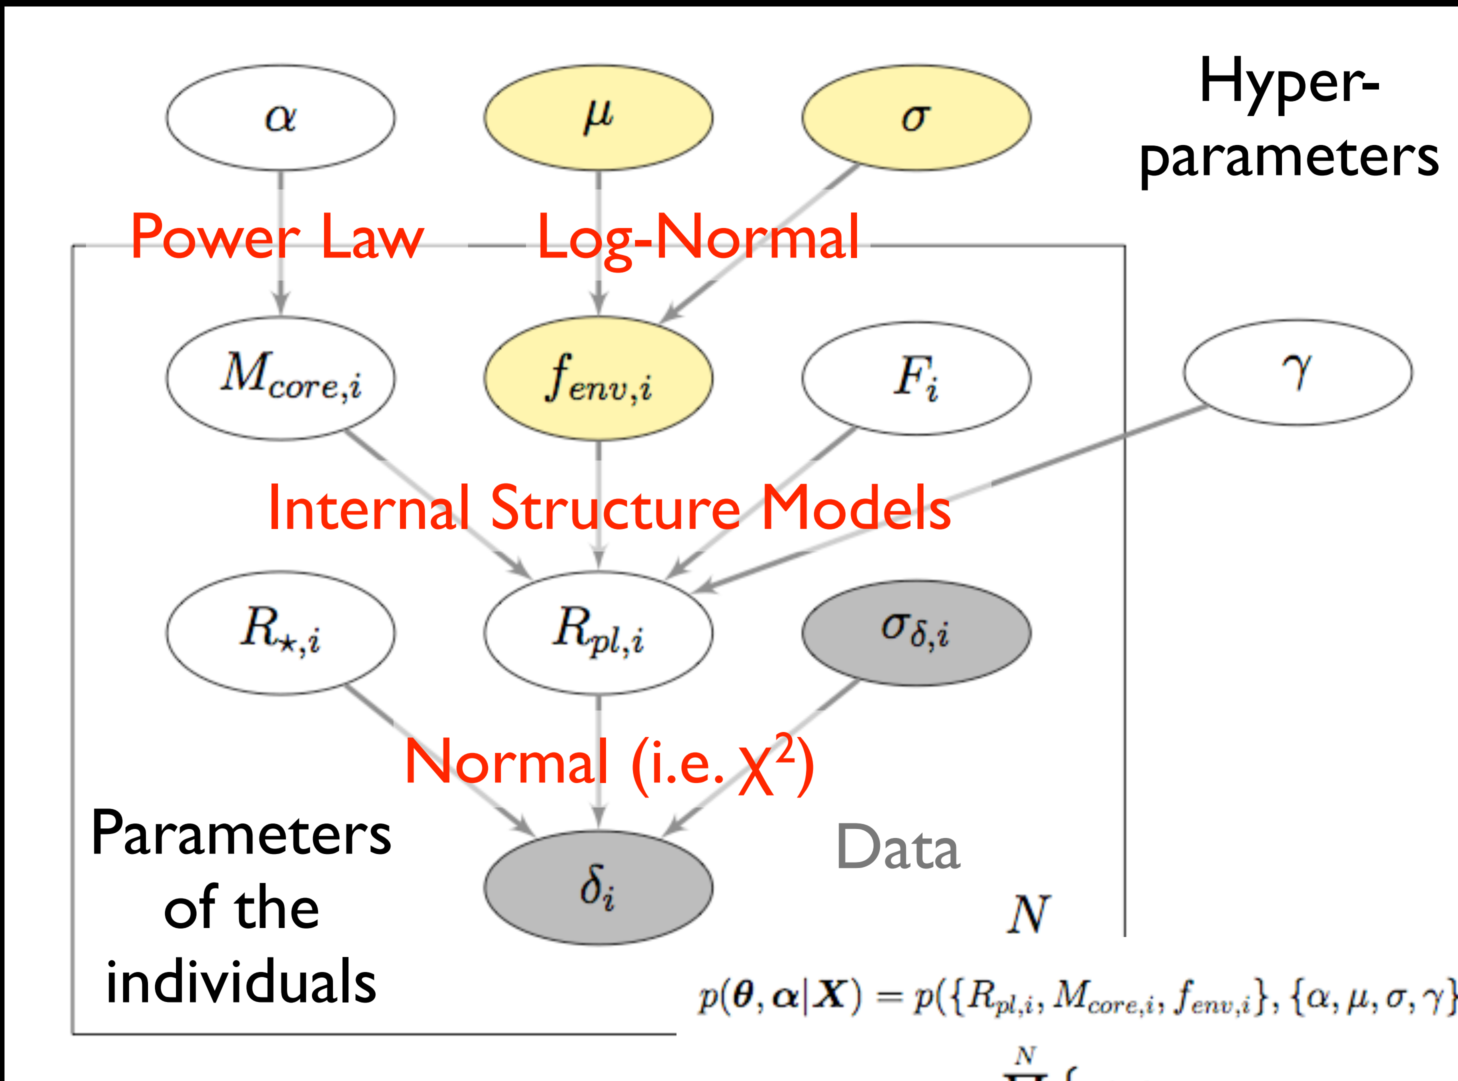

We are interested in the parameters defining planet compositions.

$$p(\boldsymbol{\theta},\boldsymbol{\alpha}|\boldsymbol{X}) = p(\{R_{pl,i}, M_{core,i}, f_{env,i}\},\{\alpha,\mu,\sigma,\gamma\}|\delta_i,\sigma_{\delta,i},F_i) \propto$$

$$\prod_{i=1}^{N}\left\{p(\delta_i|\sigma_{\delta,i},R_{pl,i},R_{\star,i},M_{core,i},f_{env,i},F_i,\alpha,\mu,\sigma,\gamma)\right\}$$

$$\times \prod_{i=1}^{N}\left\{p(R_{\star,i})p(M_{pl,i}|\alpha)p(f_{env,i}|\mu,\sigma)\right\}p(\alpha)p(\mu)p(\sigma)p(\gamma)$$

Hierarchical Bayesian Model

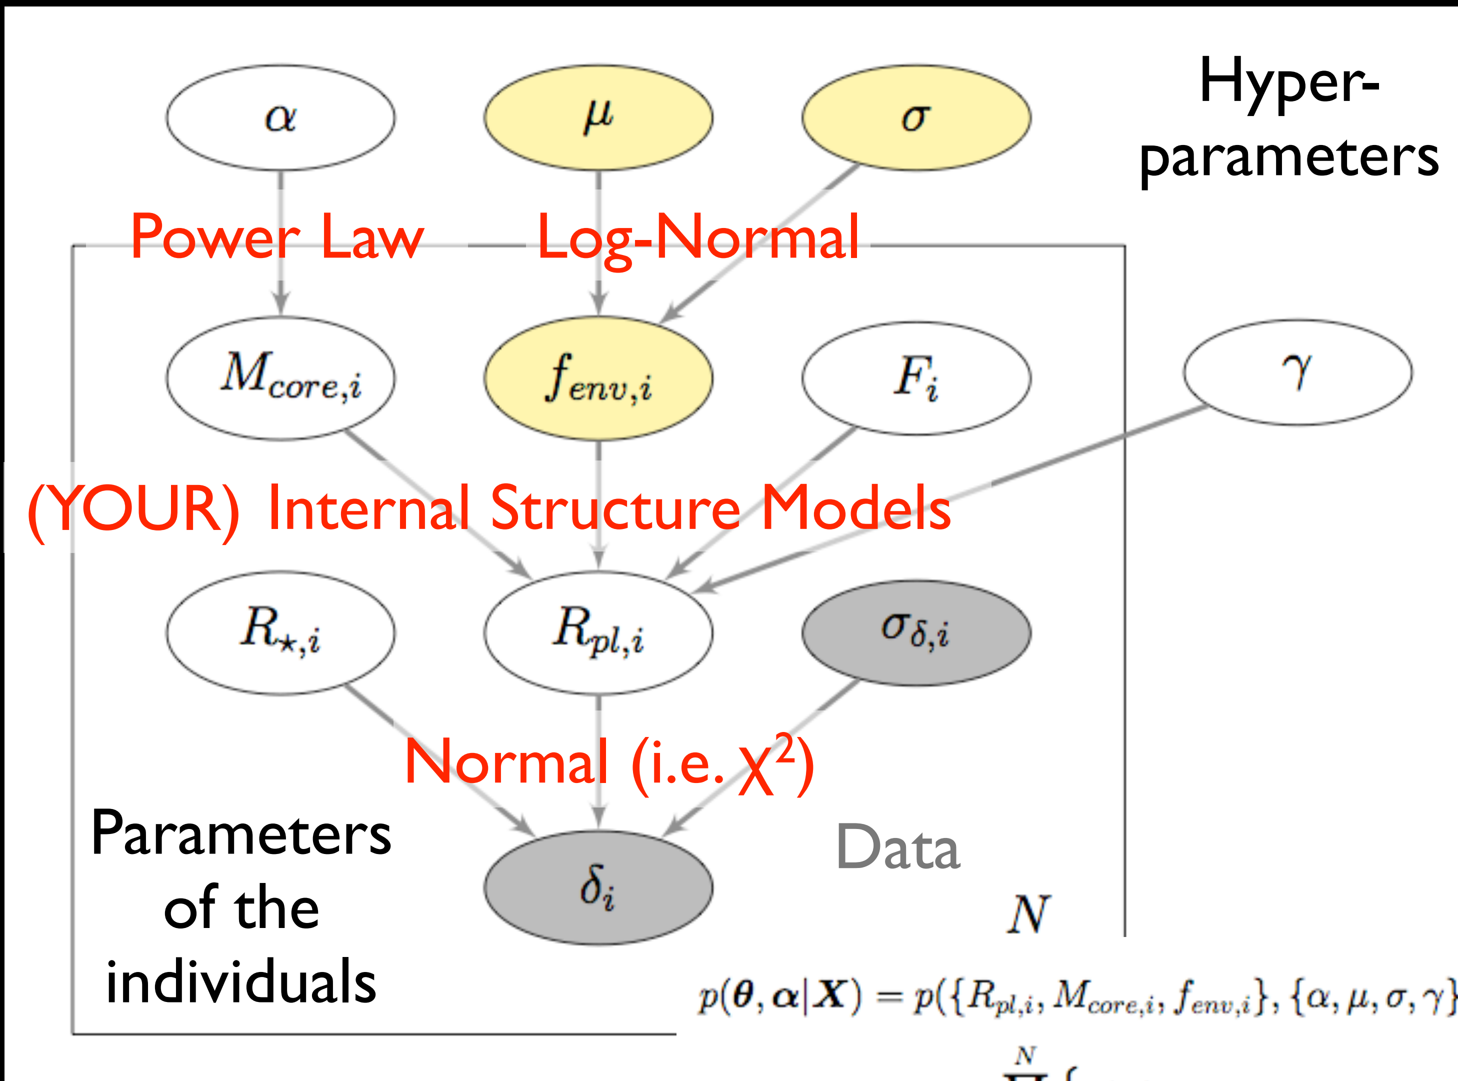

We are interested in the parameters defining planet compositions.

$$p(\boldsymbol{\theta}, \boldsymbol{\alpha}|\boldsymbol{X}) = p(\{R_{pl,i}, M_{core,i}, f_{env,i}\}, \{\alpha, \mu, \sigma, \gamma\}|\delta_i, \sigma_{\delta,i}, F_i) \propto$$

$$\prod_{i=1}^{N} \left\{ p(\delta_i|\sigma_{\delta,i}, R_{pl,i}, R_{\star,i}, M_{core,i}, f_{env,i}, F_i, \alpha, \mu, \sigma, \gamma) \right\}$$

$$\times \prod_{i=1}^{N} \left\{ p(R_{\star,i}) p(M_{pl,i}|\alpha) p(f_{env,i}|\mu, \sigma) \right\} p(\alpha) p(\mu) p(\sigma) p(\gamma)$$

Hierarchical Bayesian Model

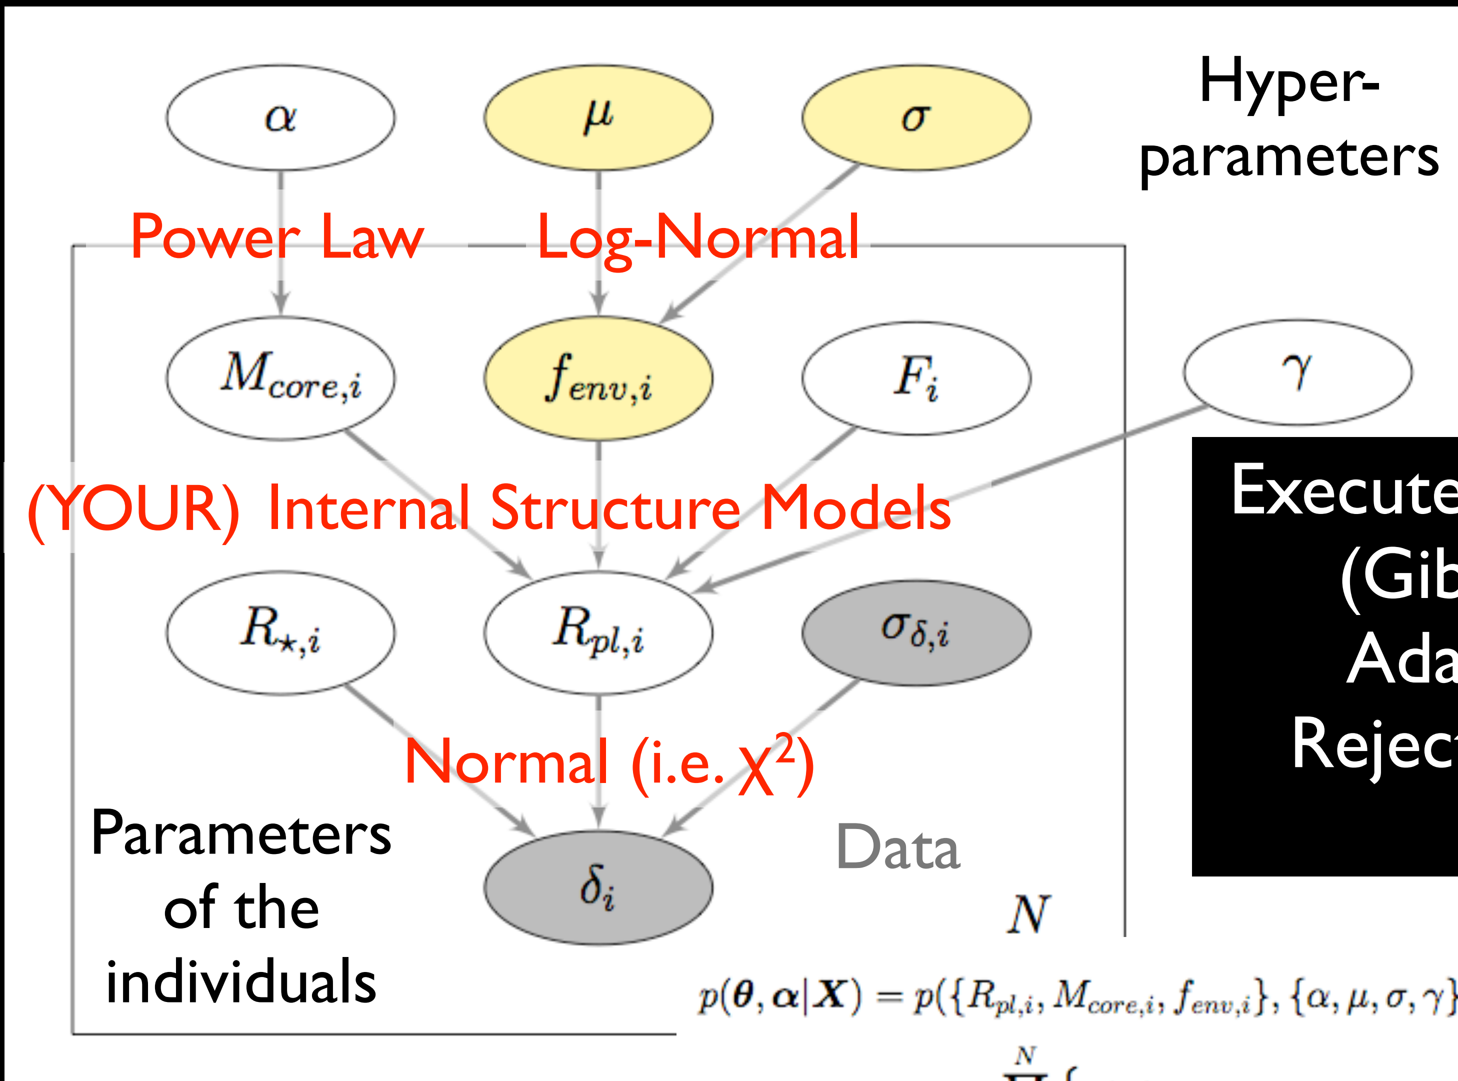

We are interested in the parameters defining planet compositions.

Execute hierarchical MCMC (Gibbs sampling with Adaptive Metropolis Rejection via JAGS in R)

$$p(\boldsymbol{\theta}, \boldsymbol{\alpha} | \boldsymbol{X}) = p(\{R_{pl,i}, M_{core,i}, f_{env,i}\}, \{\alpha, \mu, \sigma, \gamma\} | \delta_i, \sigma_{\delta,i}, F_i) \propto$$

$$\prod_{i=1}^{N} \left\{ p(\delta_i | \sigma_{\delta,i}, R_{pl,i}, R_{\star,i}, M_{core,i}, f_{env,i}, F_i, \alpha, \mu, \sigma, \gamma) \right\}$$

$$\times \prod_{i=1}^{N} \left\{ p(R_{\star,i}) p(M_{pl,i} | \alpha) p(f_{env,i} | \mu, \sigma) \right\} p(\alpha) p(\mu) p(\sigma) p(\gamma)$$

Hierarchical Bayesian Model

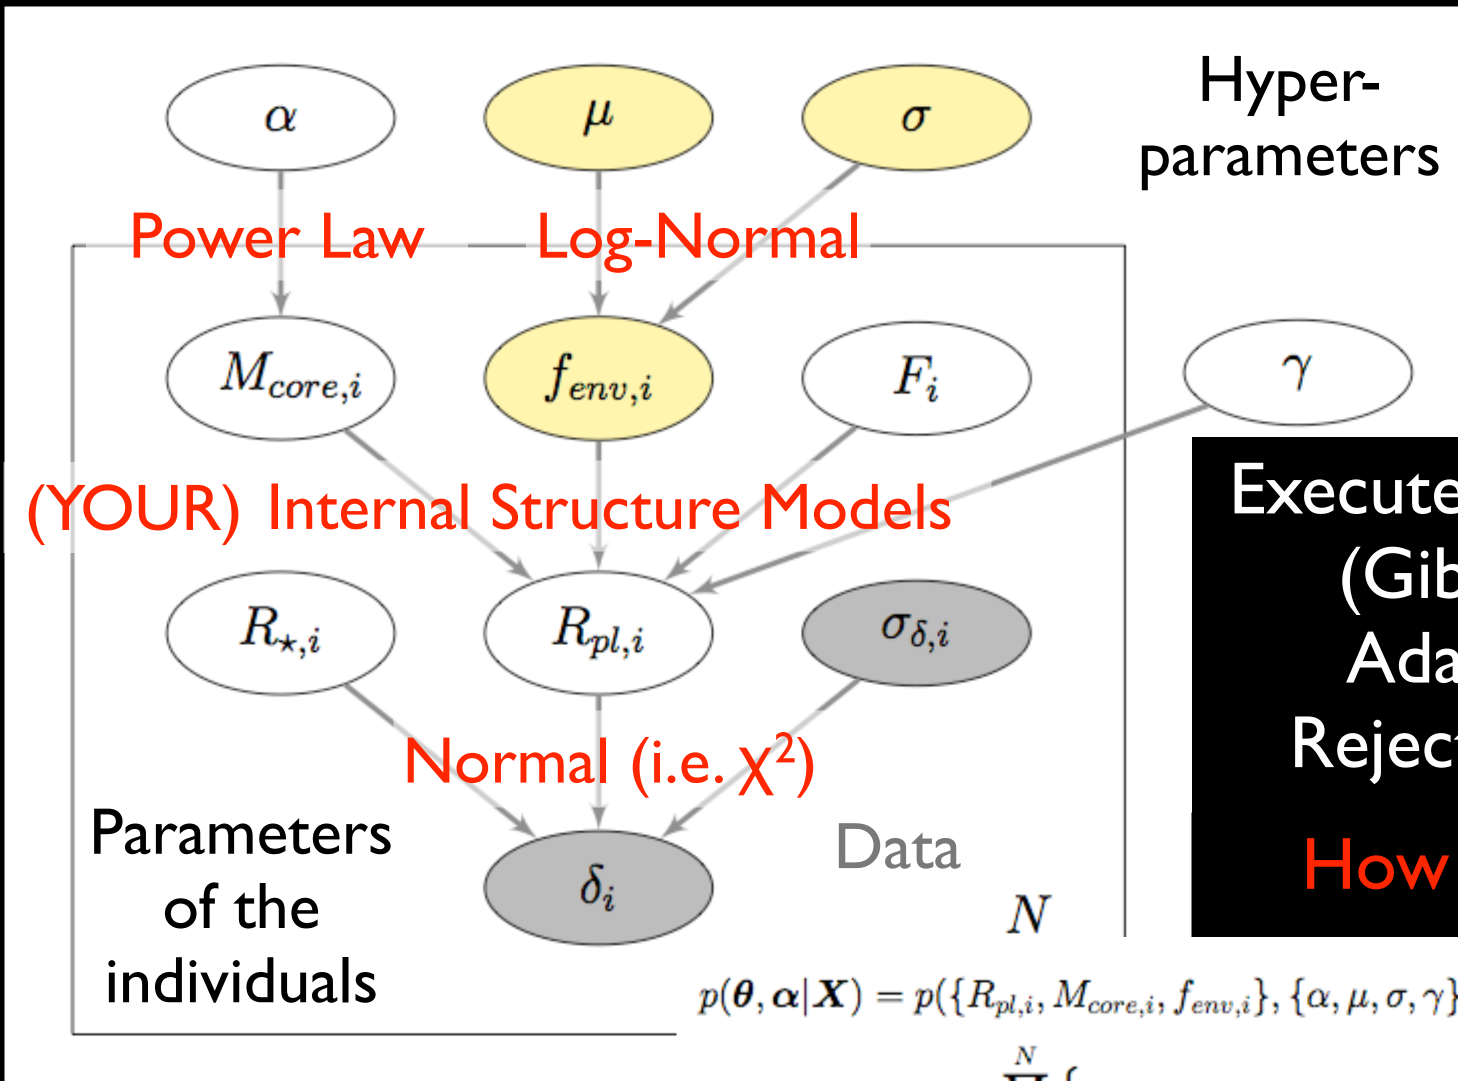

We are interested in the parameters defining planet compositions.

Execute hierarchical MCMC (Gibbs sampling with Adaptive Metropolis Rejection via JAGS in R)

How to define dataset?

$$p(\boldsymbol{\theta}, \boldsymbol{\alpha}|\boldsymbol{X}) = p(\{R_{pl,i}, M_{core,i}, f_{env,i}\}, \{\alpha, \mu, \sigma, \gamma\}|\delta_i, \sigma_{\delta,i}, F_i) \propto$$

$$\prod_{i=1}^{N} \Big\{ p(\delta_i|\sigma_{\delta,i}, R_{pl,i}, R_{\star,i}, M_{core,i}, f_{env,i}, F_i, \alpha, \mu, \sigma, \gamma) \Big\}$$

$$\times \prod_{i=1}^{N} \Big\{ p(R_{\star,i}) p(M_{pl,i}|\alpha) p(f_{env,i}|\mu, \sigma) \Big\} p(\alpha) p(\mu) p(\sigma) p(\gamma)$$

Applying to *Kepler* planets

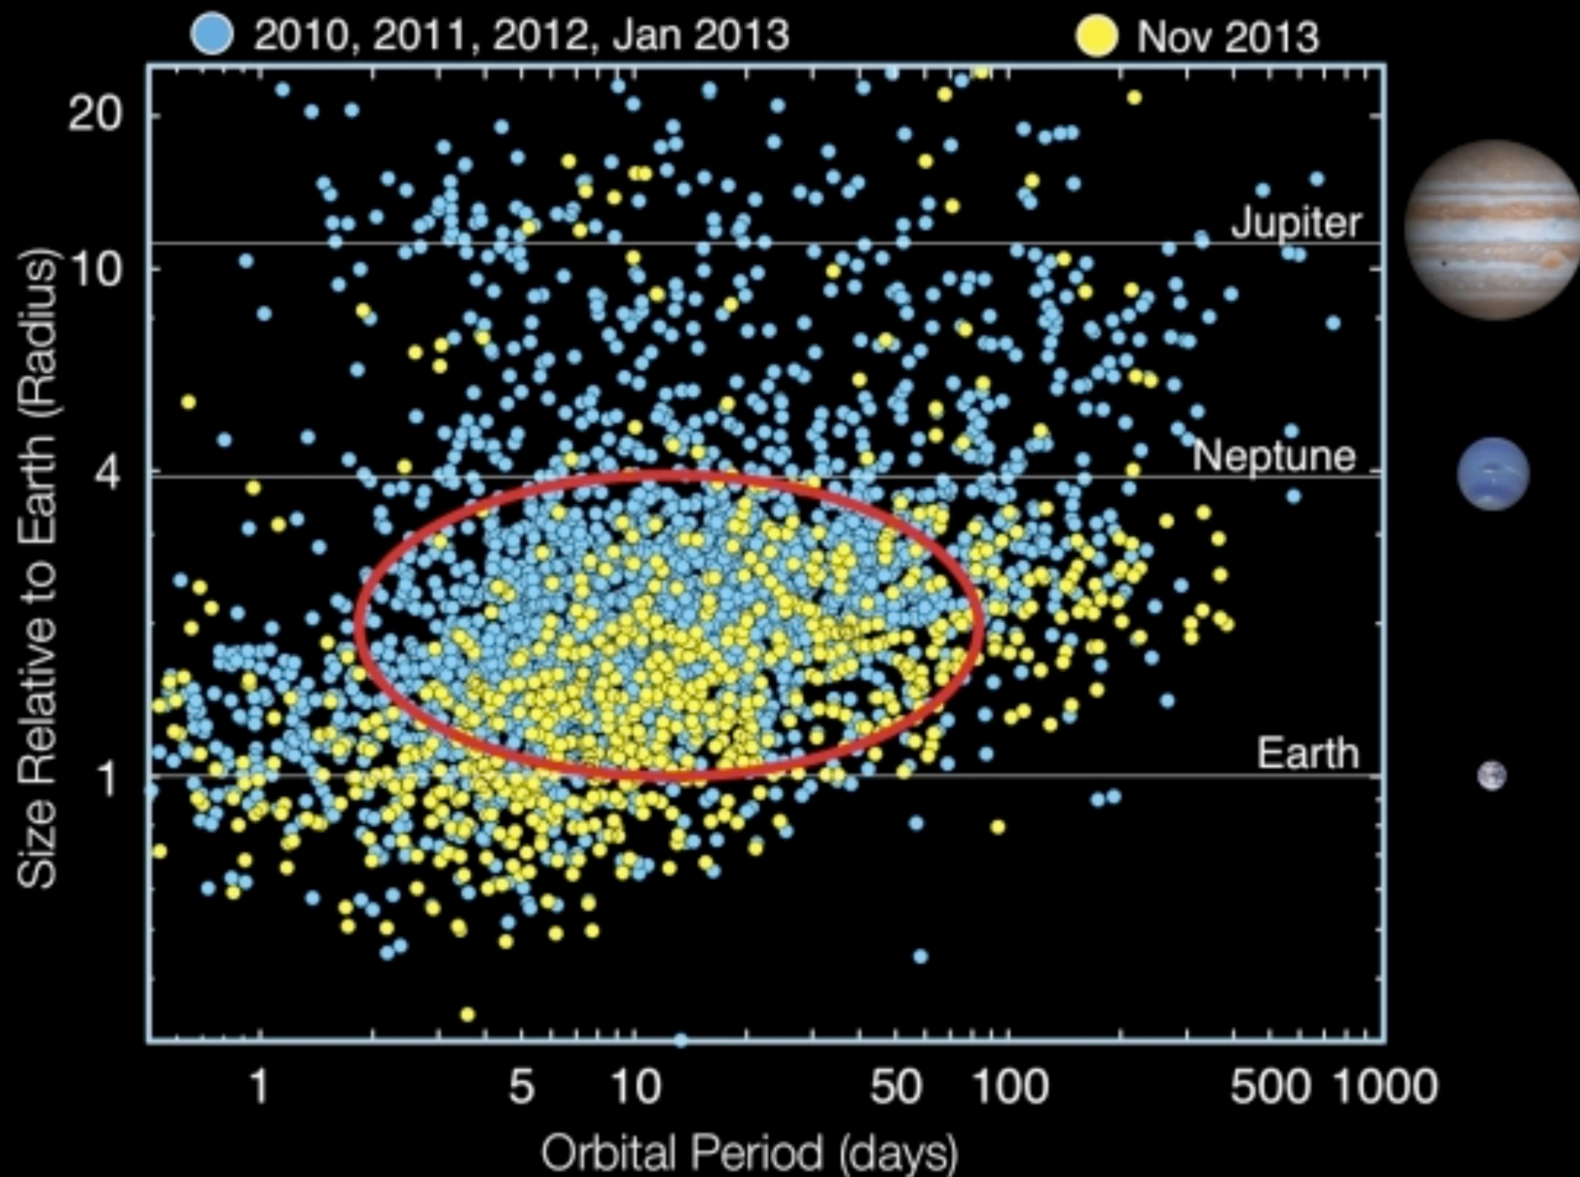

Applying to *Kepler* planets

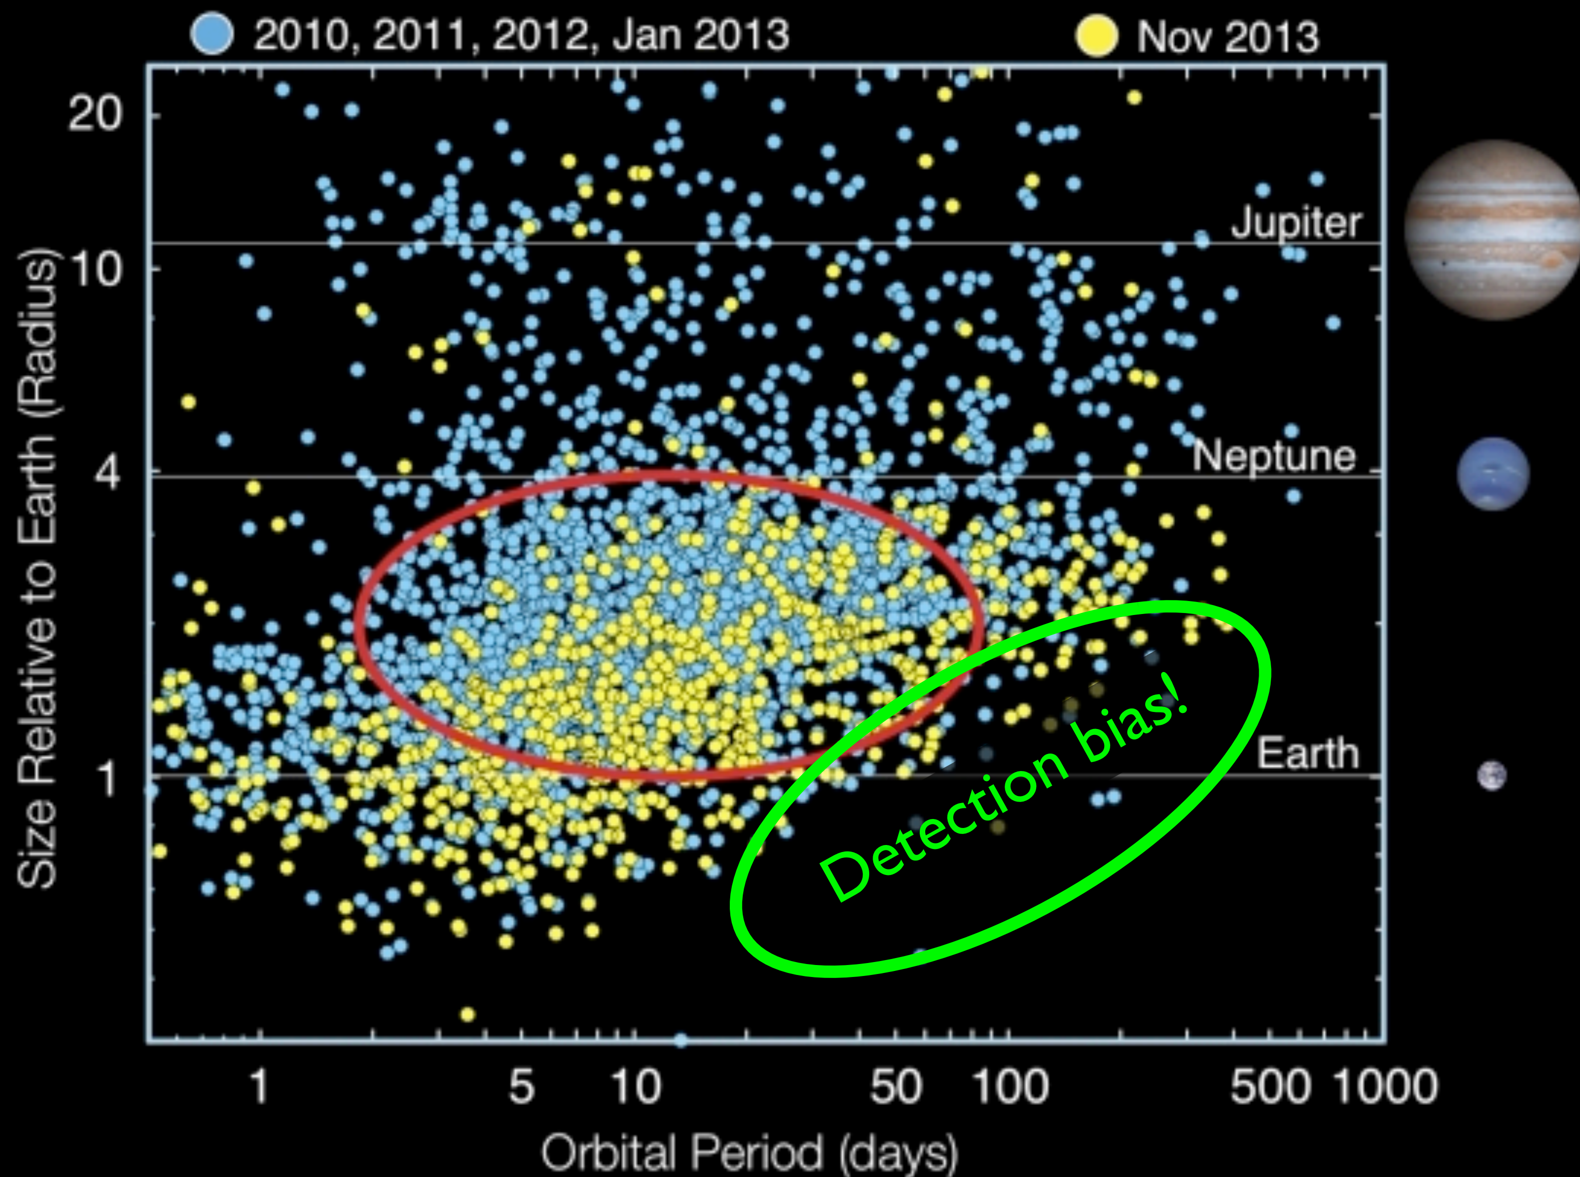

Applying to *Kepler* planets

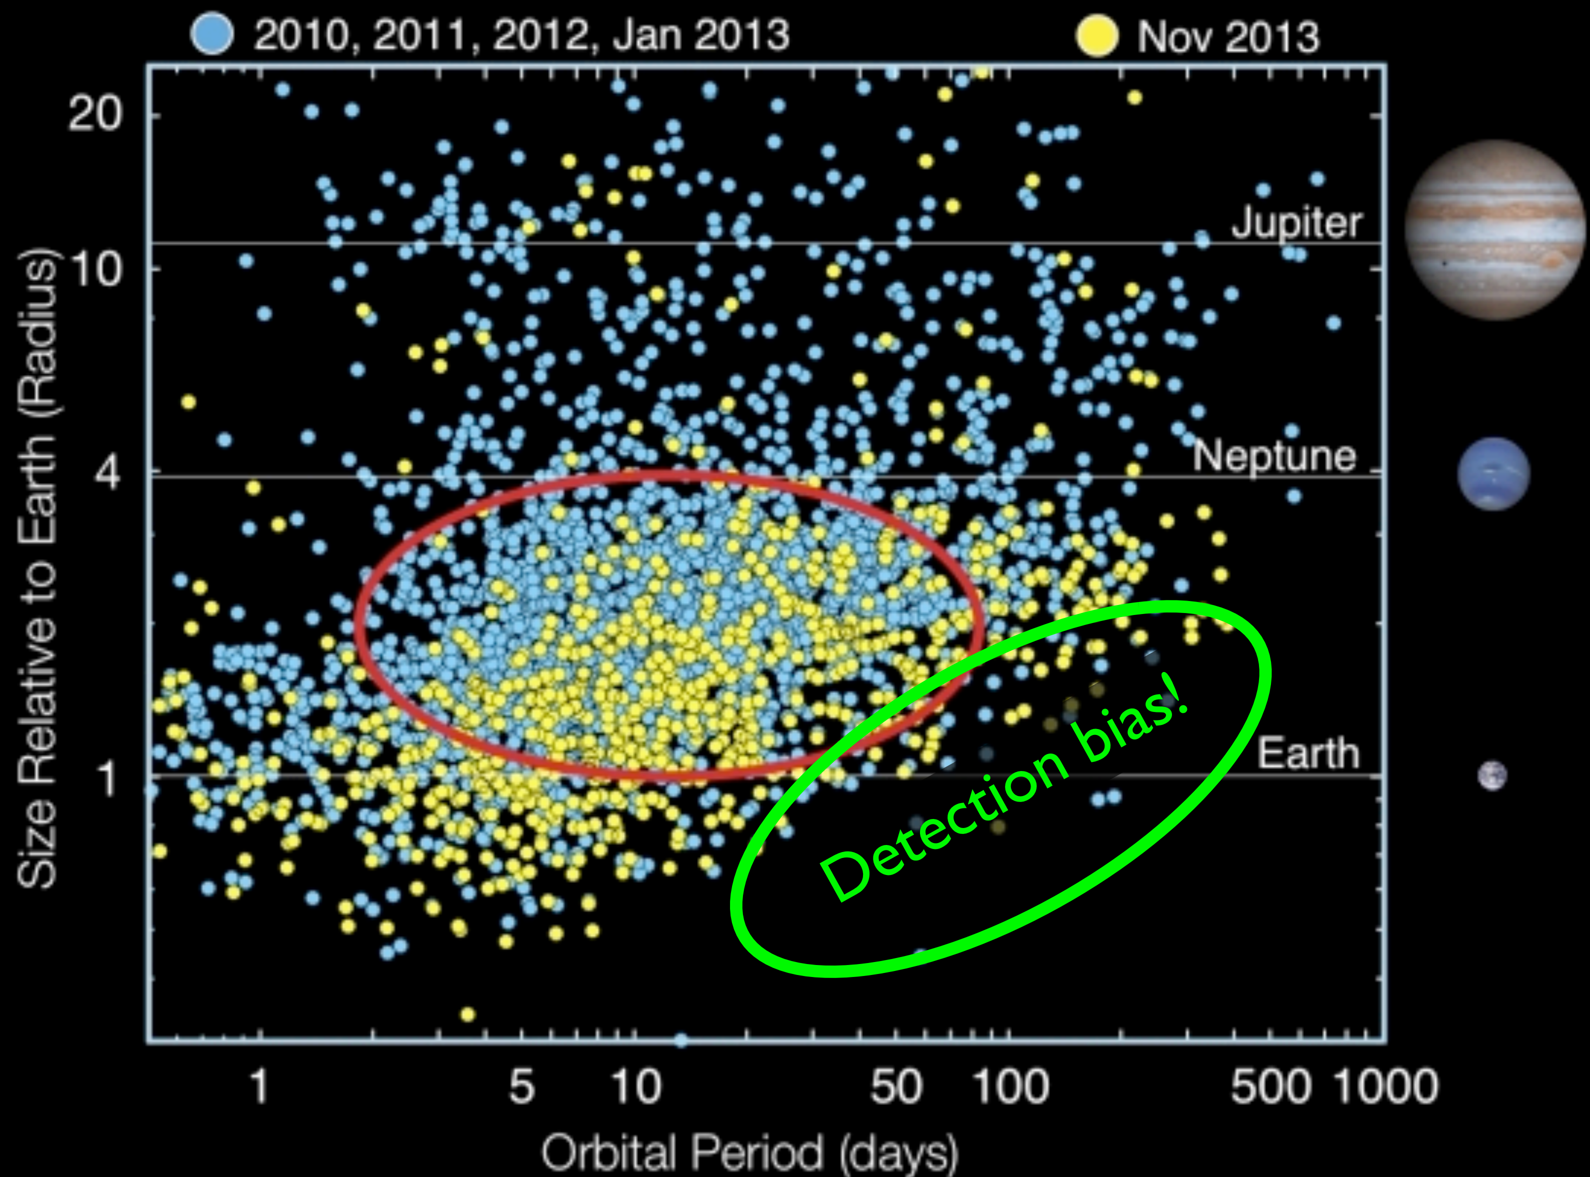

Extreme caution needed in interpreting the observed radius distribution!

Sample Carefully Chosen

Sample Carefully Chosen

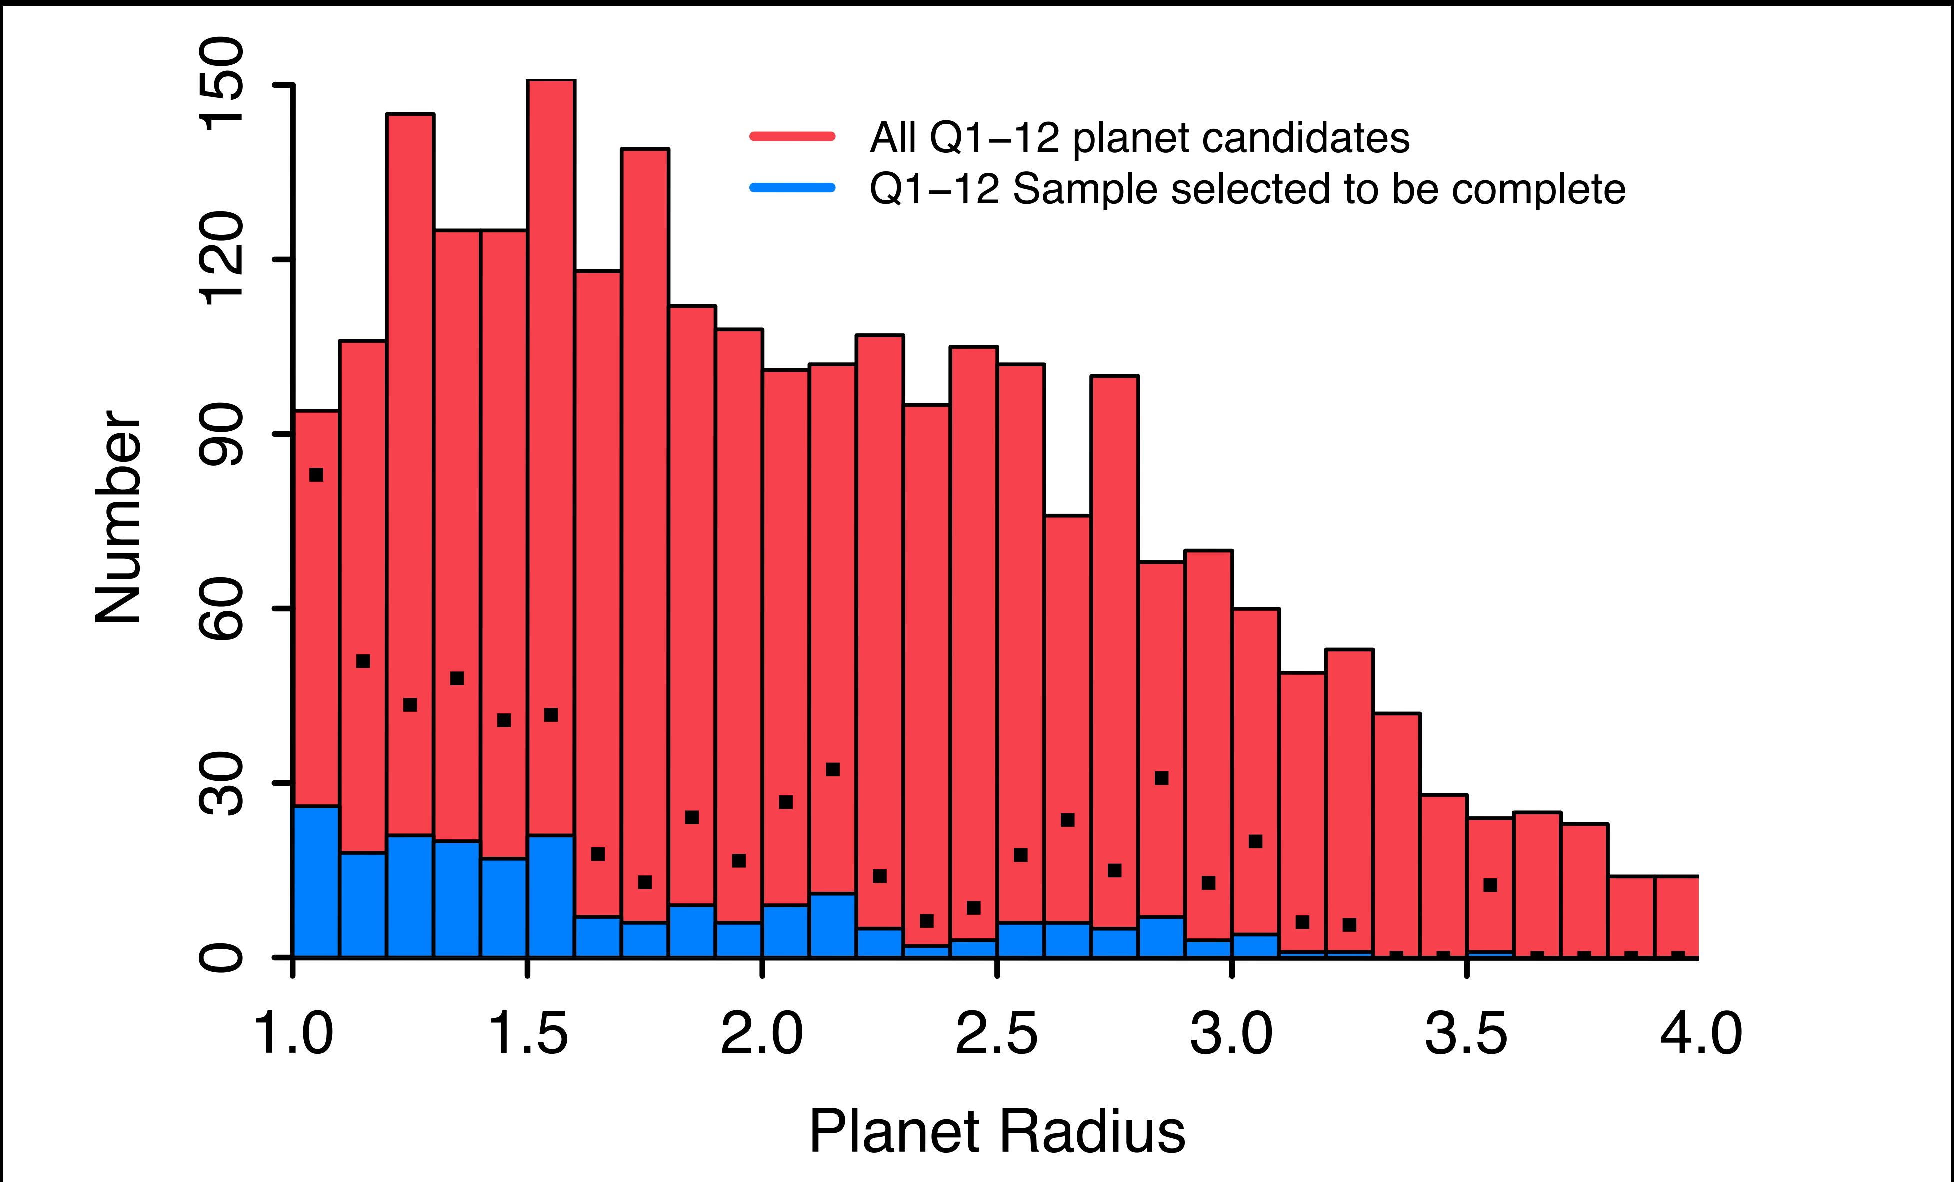

Sample Carefully Chosen

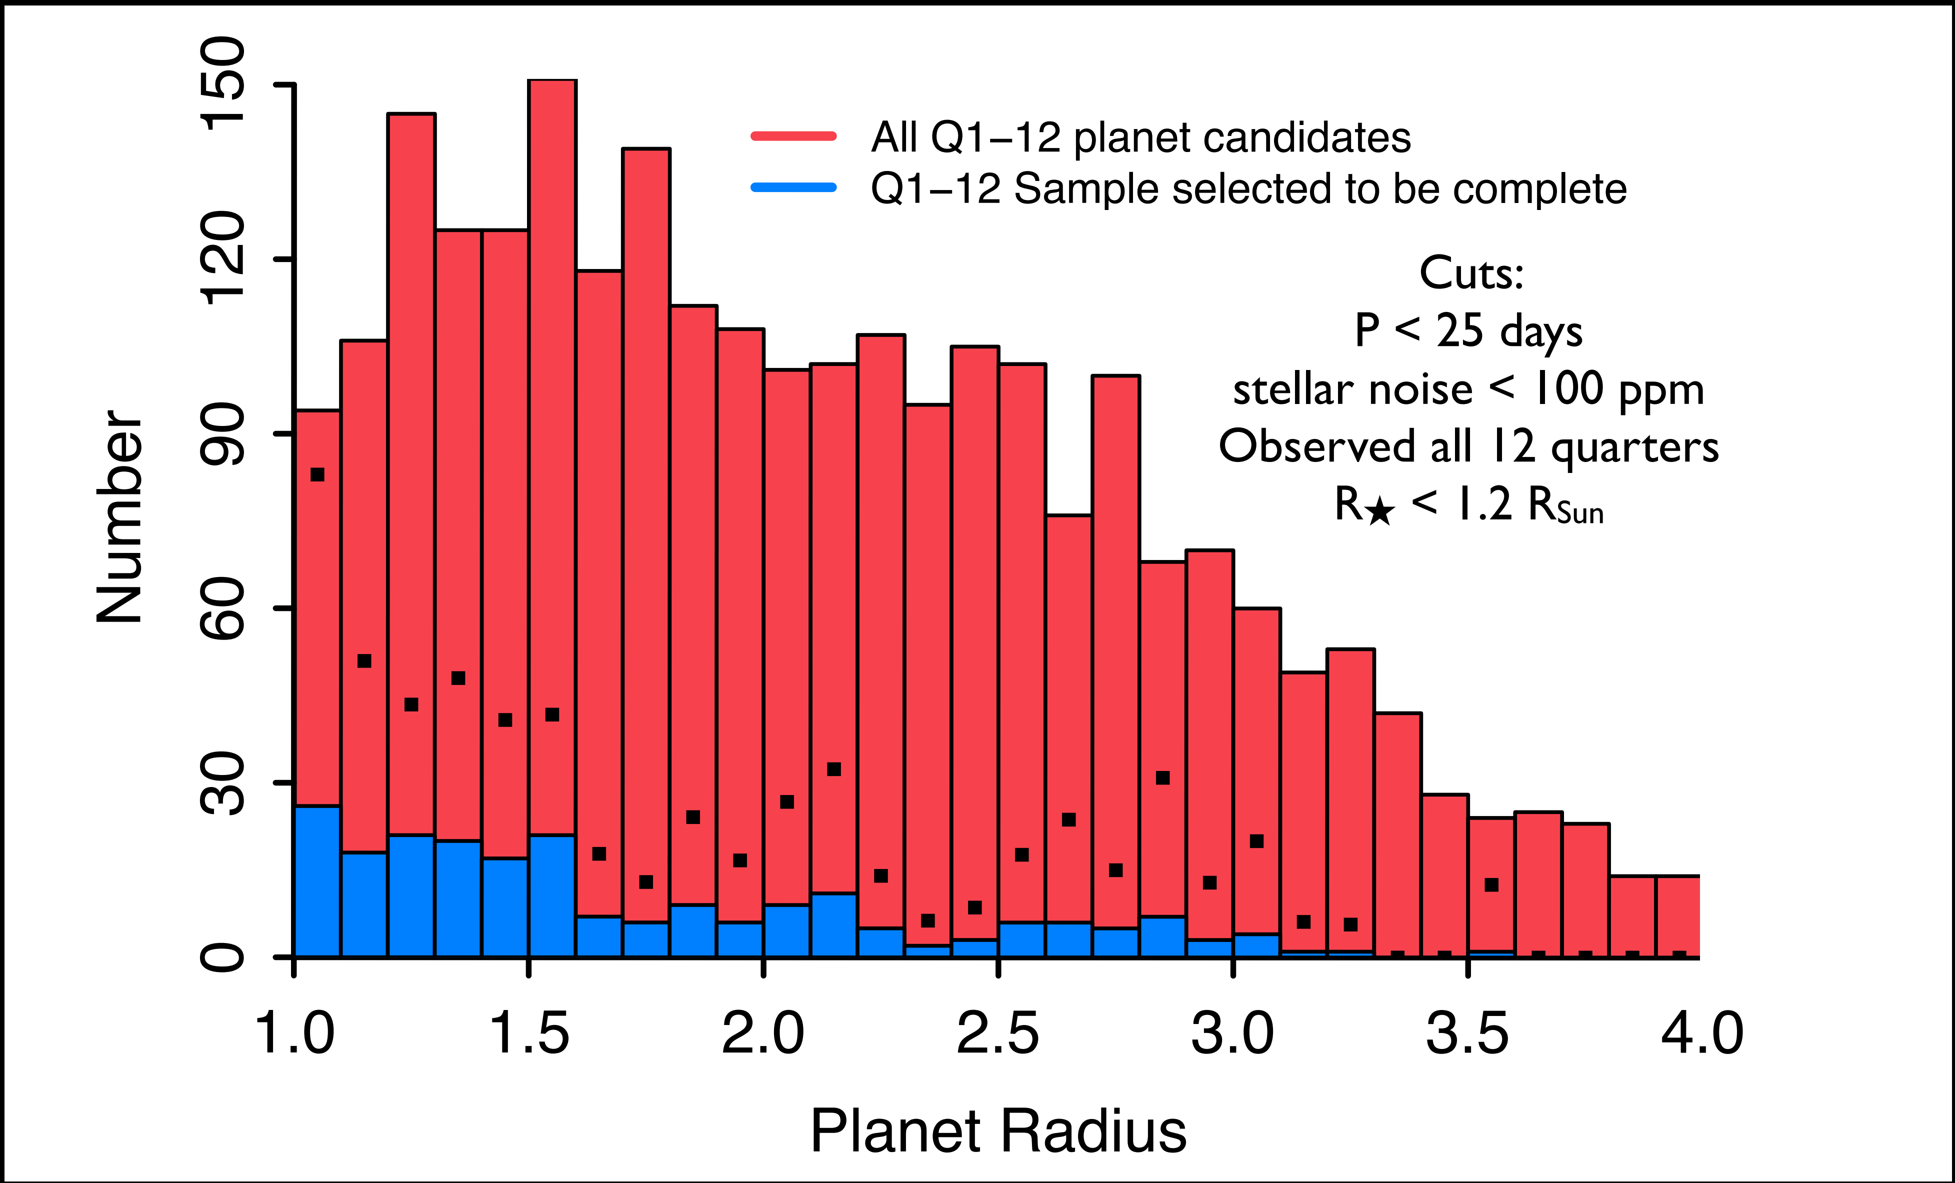

Sample Carefully Chosen

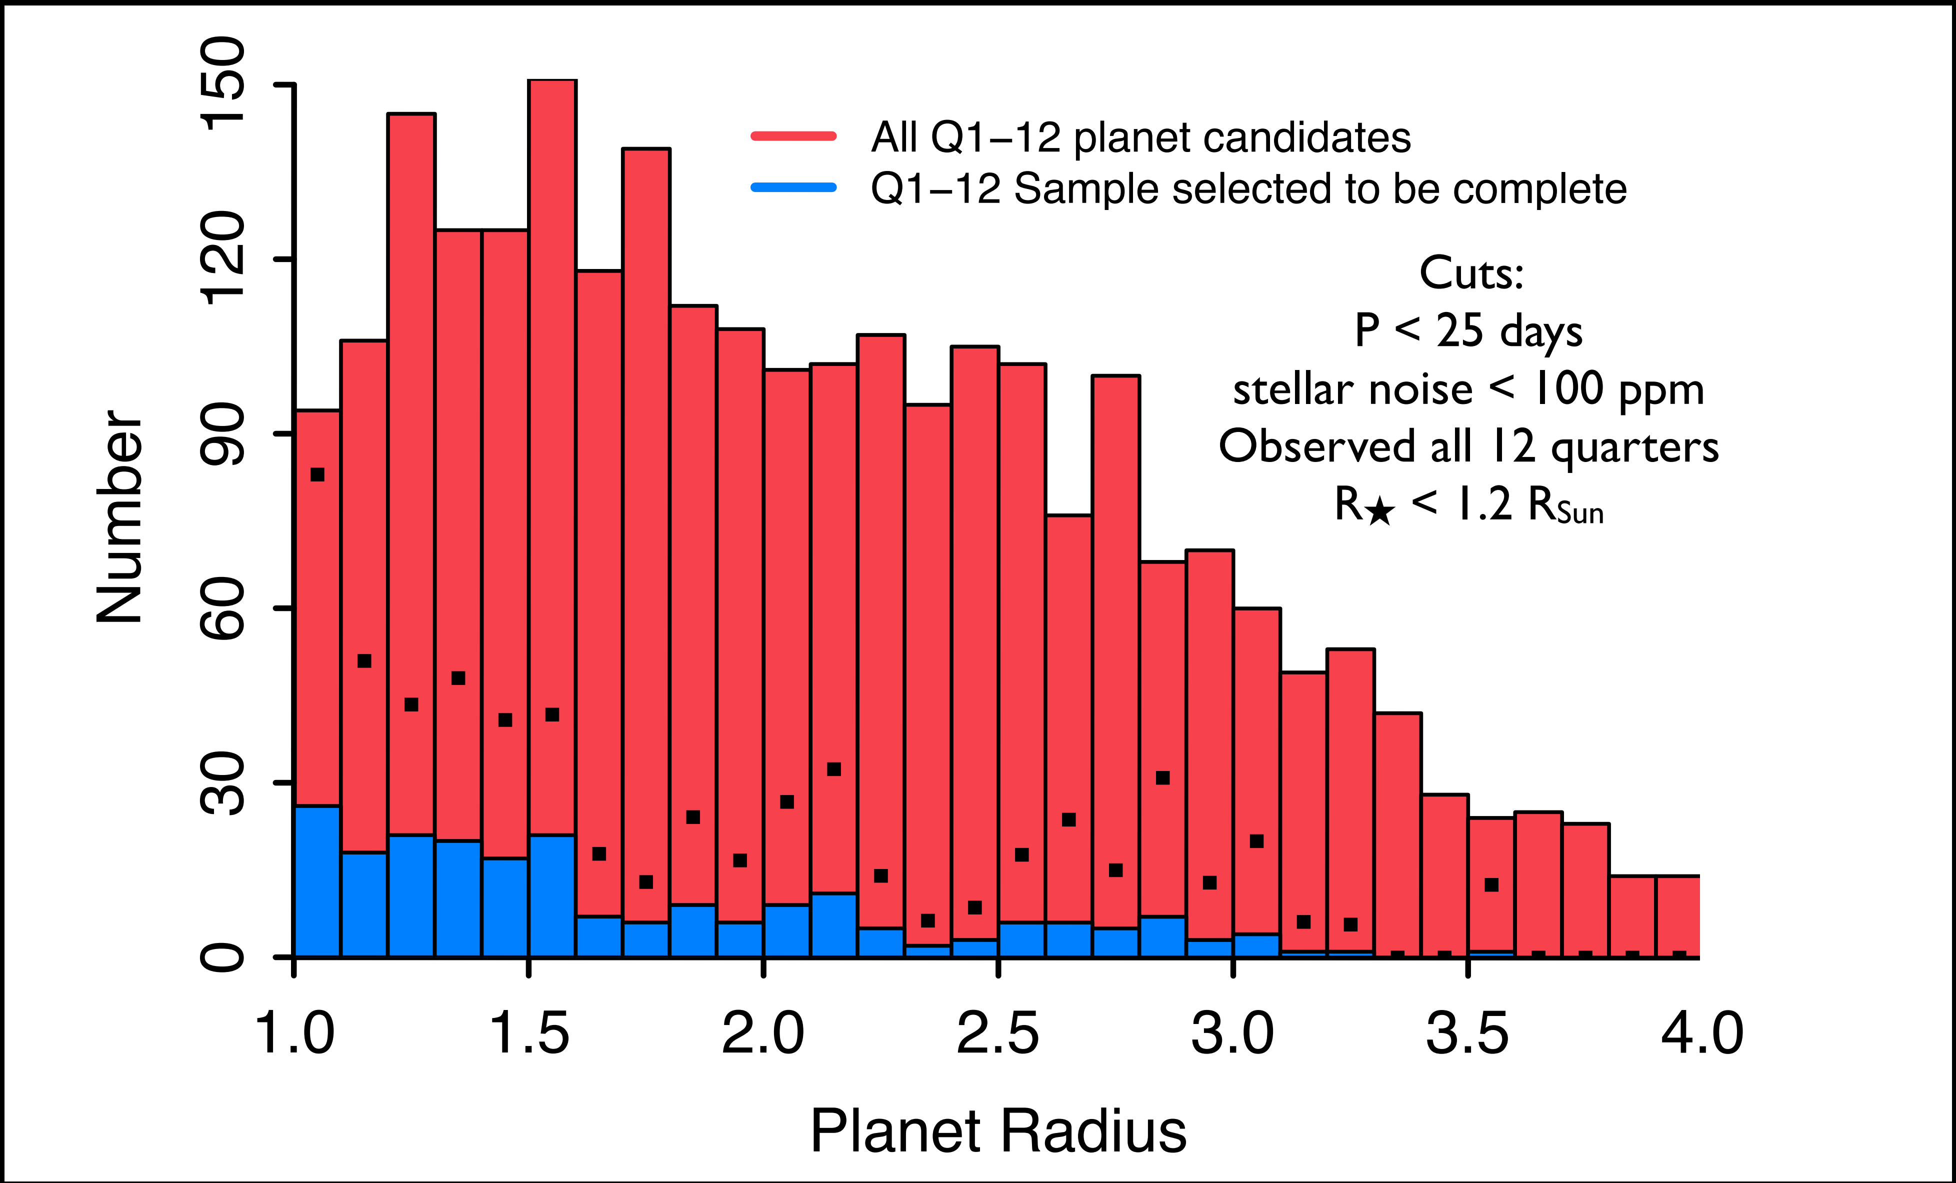

Run the MCMC, and . . .

Results

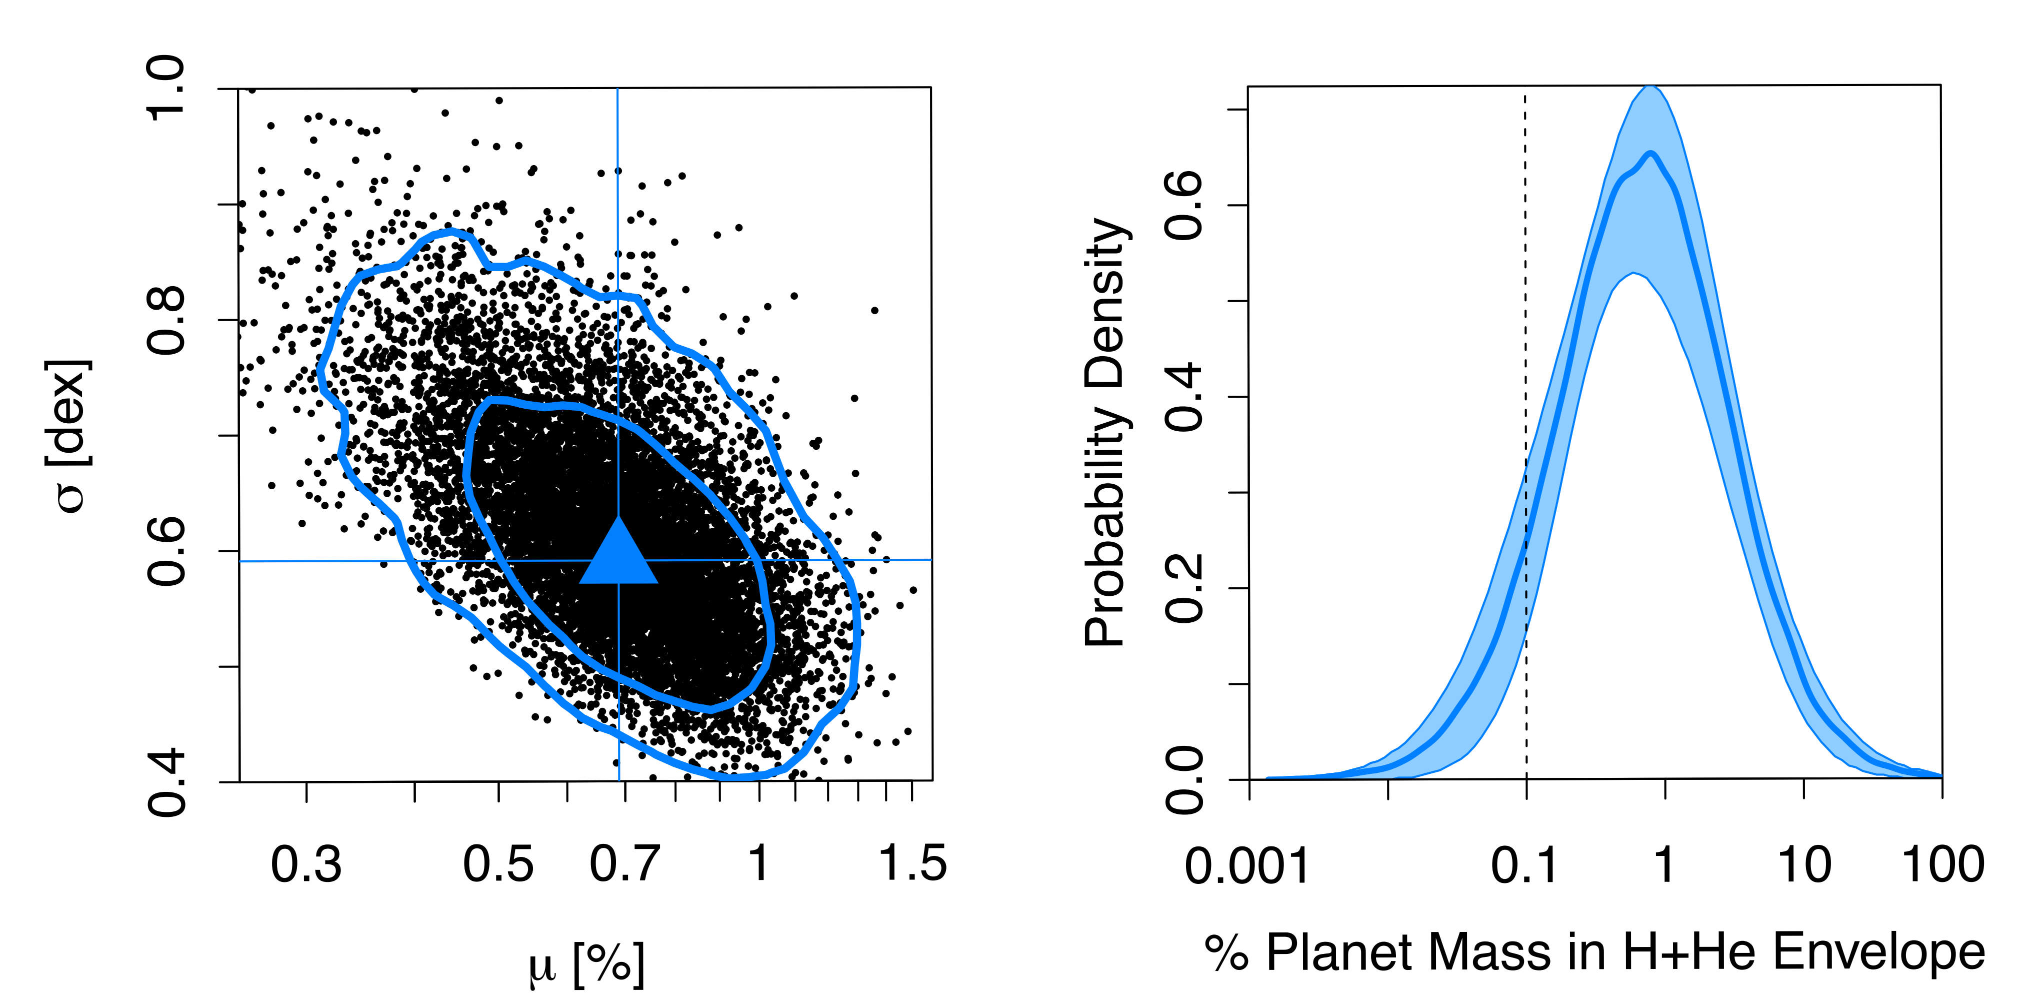

1.0
0.8
0.6
0.4
σ [dex]
0.3
0.5
0.7
1
1.5
μ [%]
Probability Density
0.0
0.2
0.4
0.6
0.001
0.1
1
10
100
% Planet Mass in H+He Envelope

Results

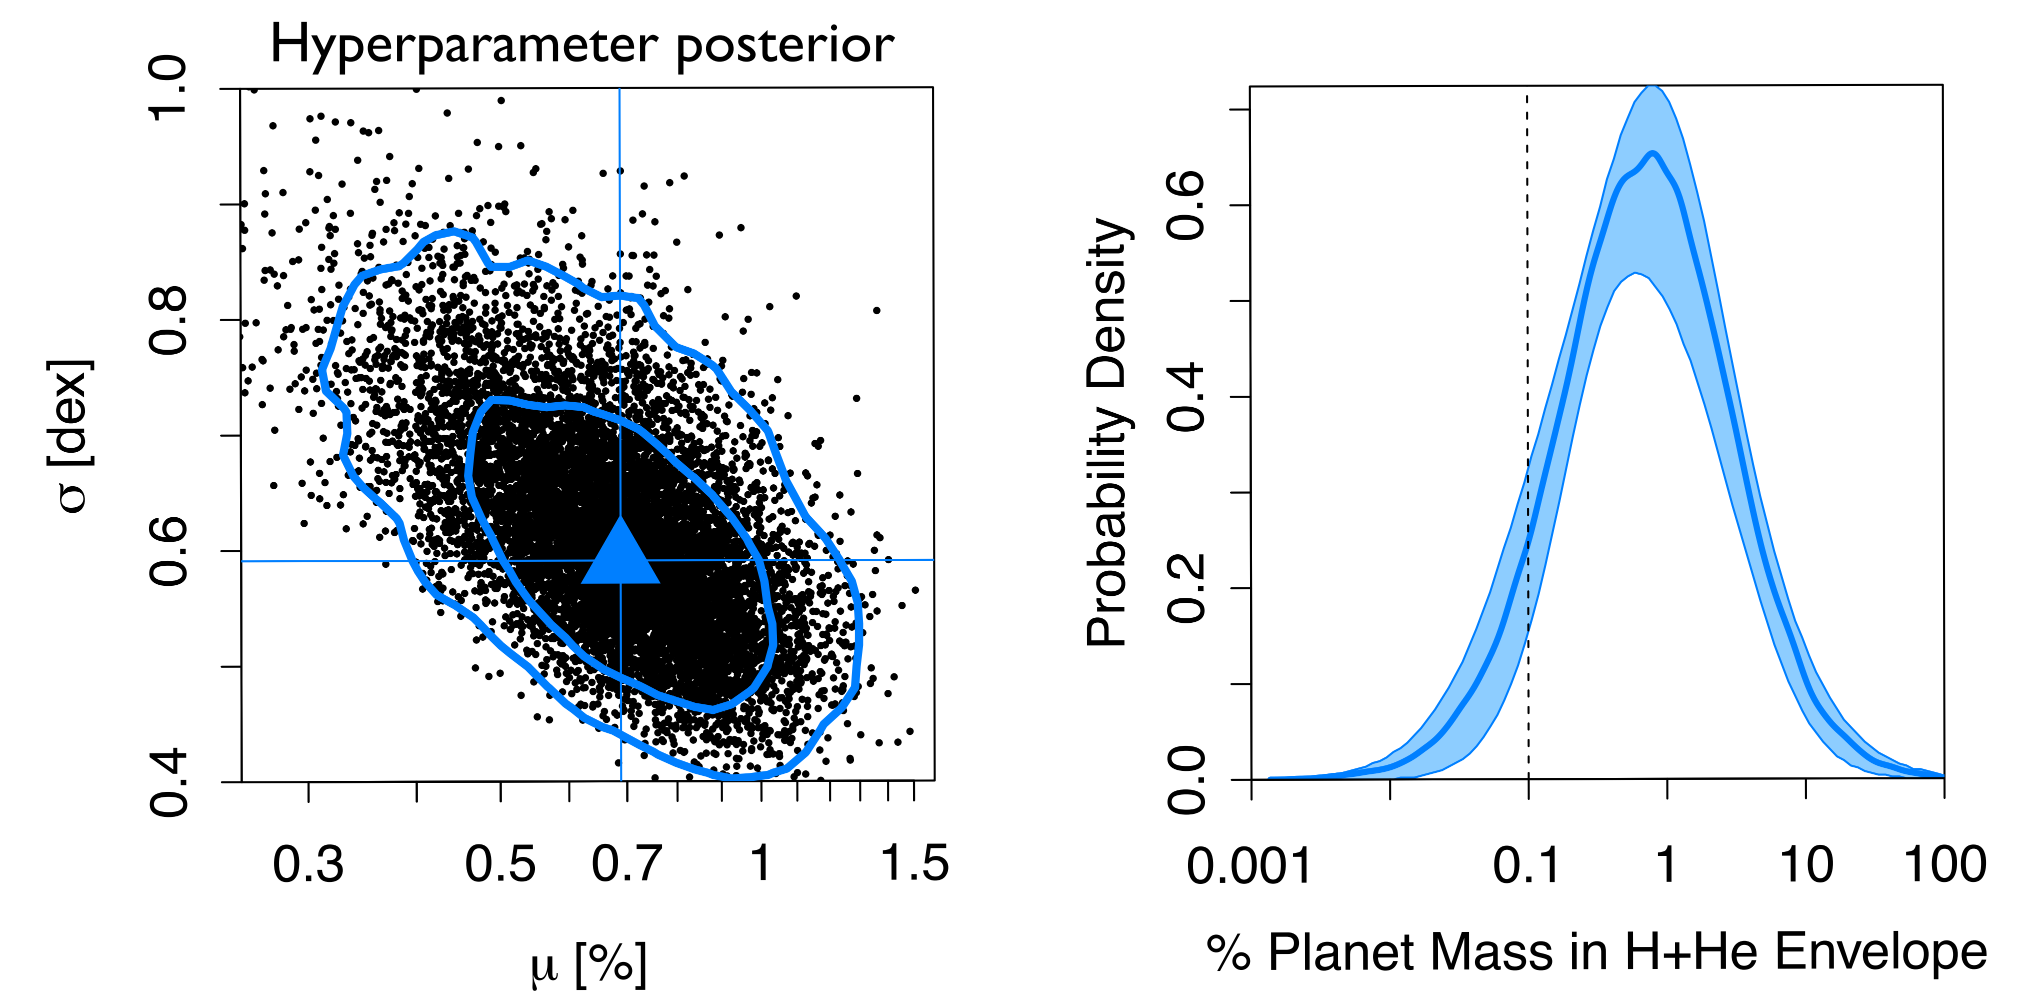
Hyperparameter posterior
σ [dex]
1.0
0.8
0.6
0.4
0.3
0.5
0.7
1
1.5
μ [%]
Probability Density
0.6
0.4
0.2
0.0
0.001
0.1
1
10
100
% Planet Mass in H+He Envelope

Results

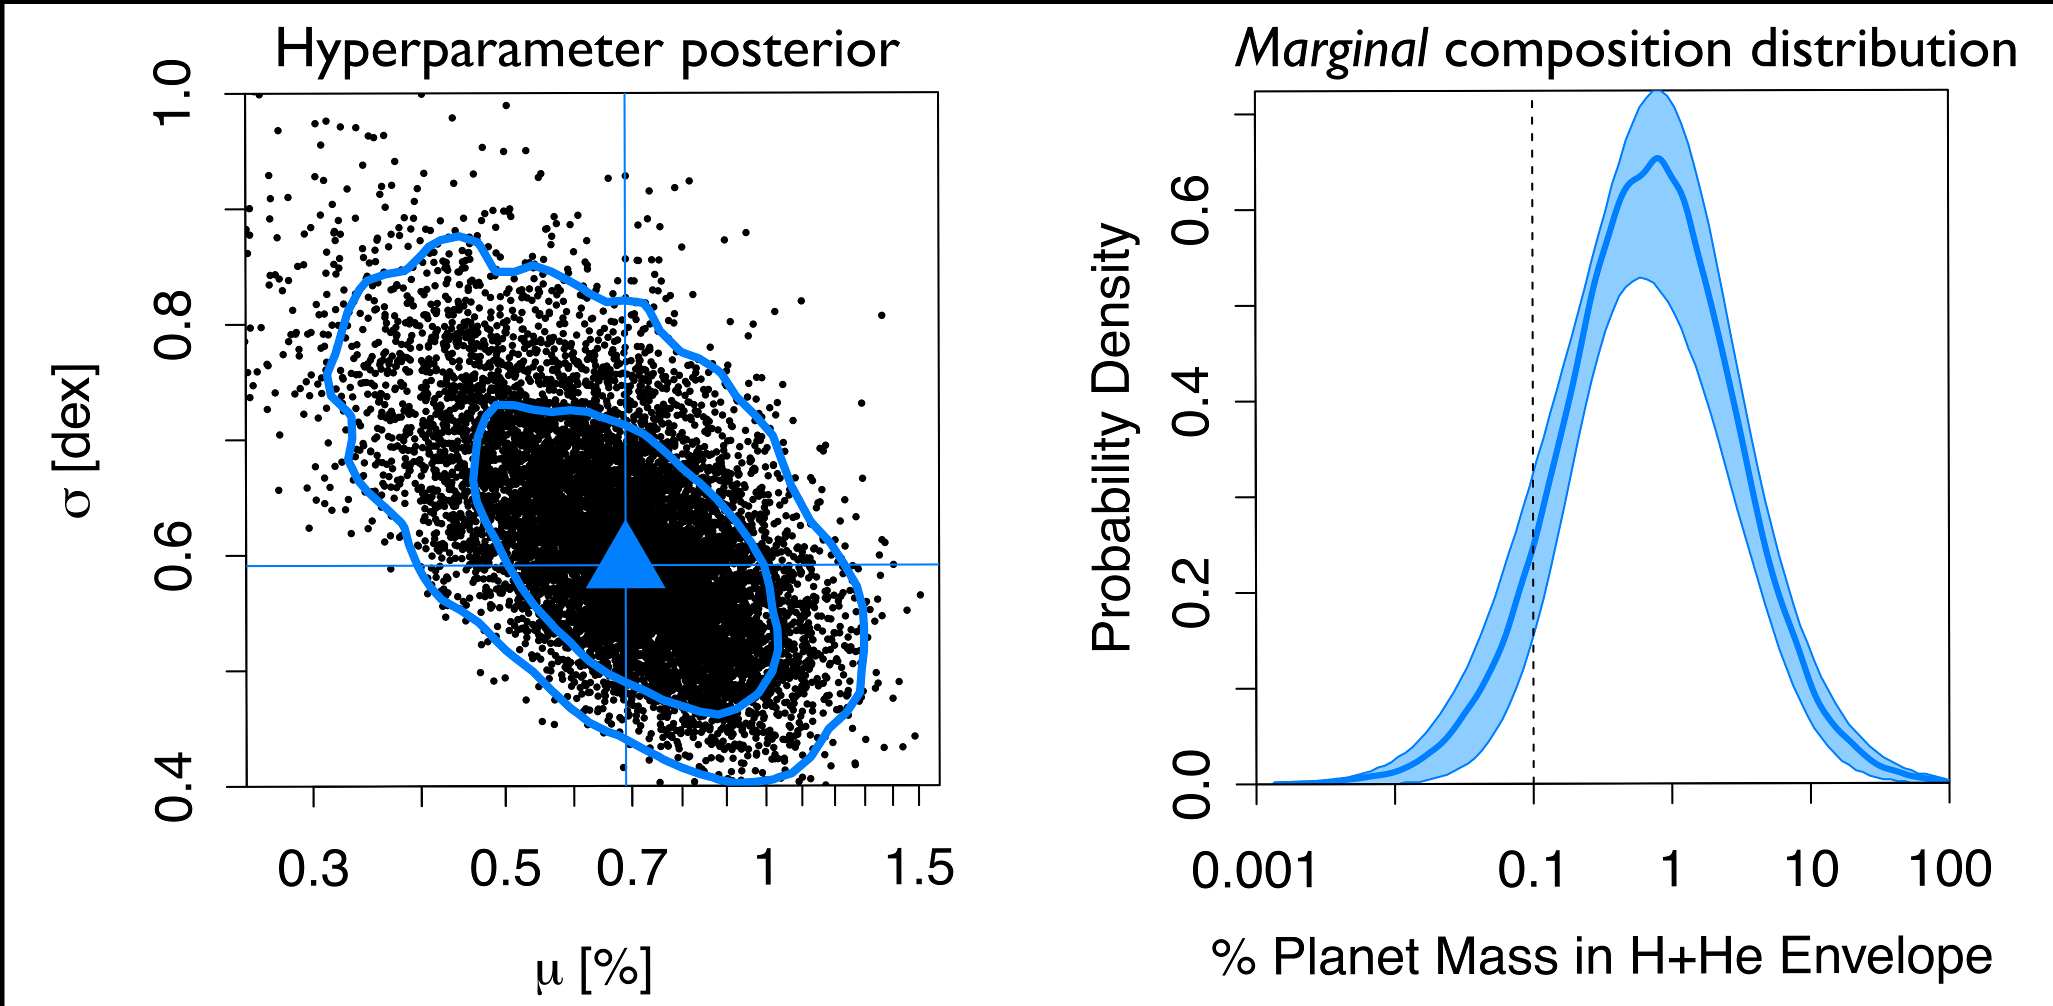

Hyperparameter posterior
σ [dex]
0.4
0.6
0.8
1.0
0.3
0.5
0.7
1
1.5
μ [%]
Marginal composition distribution
Probability Density
0.0
0.2
0.4
0.6
0.001
0.1
1
10
100
% Planet Mass in H+He Envelope

Results

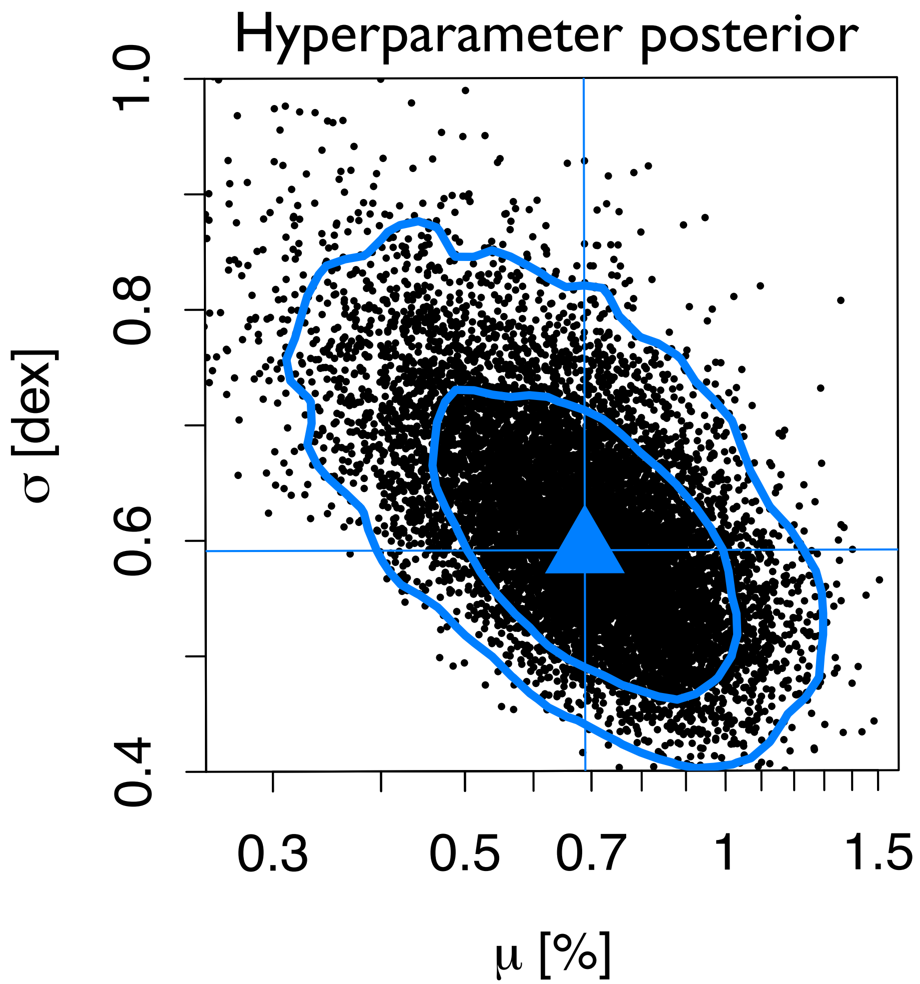

Hyperparameter posterior
σ [dex]
1.0
0.8
0.6
0.4
0.3
0.5
0.7
1
1.5
μ [%]

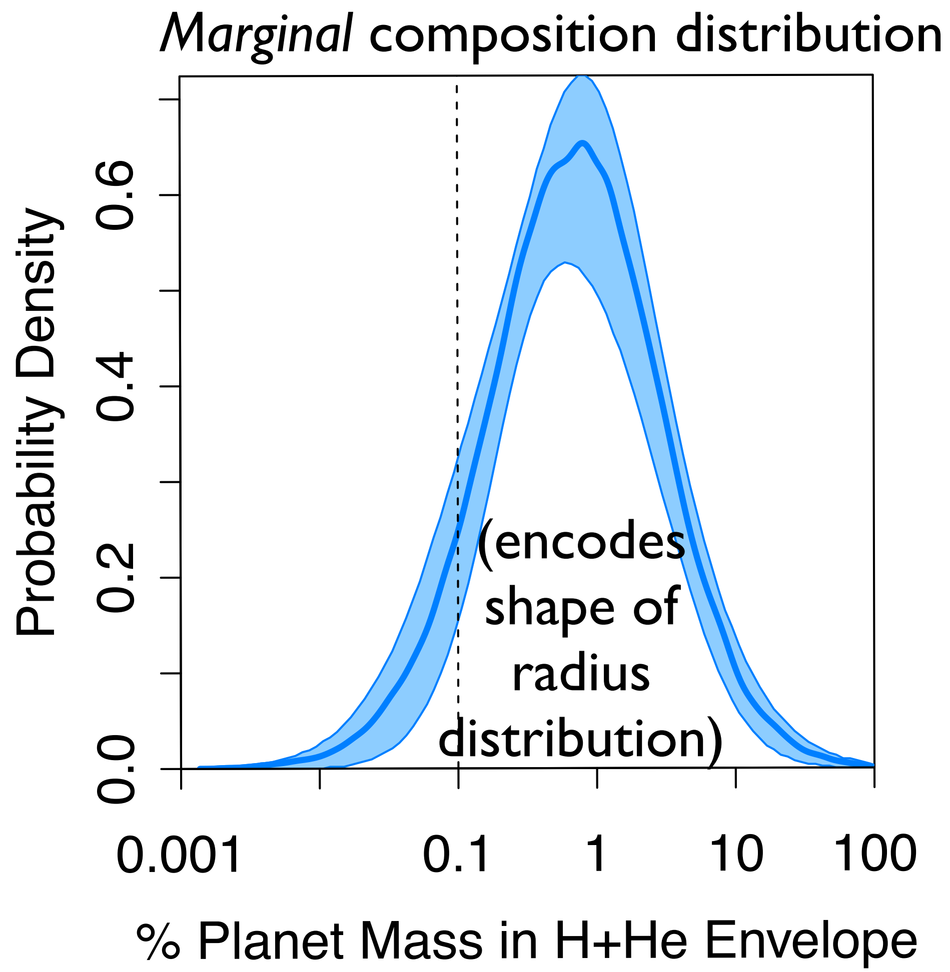

Marginal composition distribution
Probability Density
0.6
0.4
0.2
0.0
(encodes shape of radius distribution)
0.001
0.1
1
10
100
% Planet Mass in H+He Envelope

Results

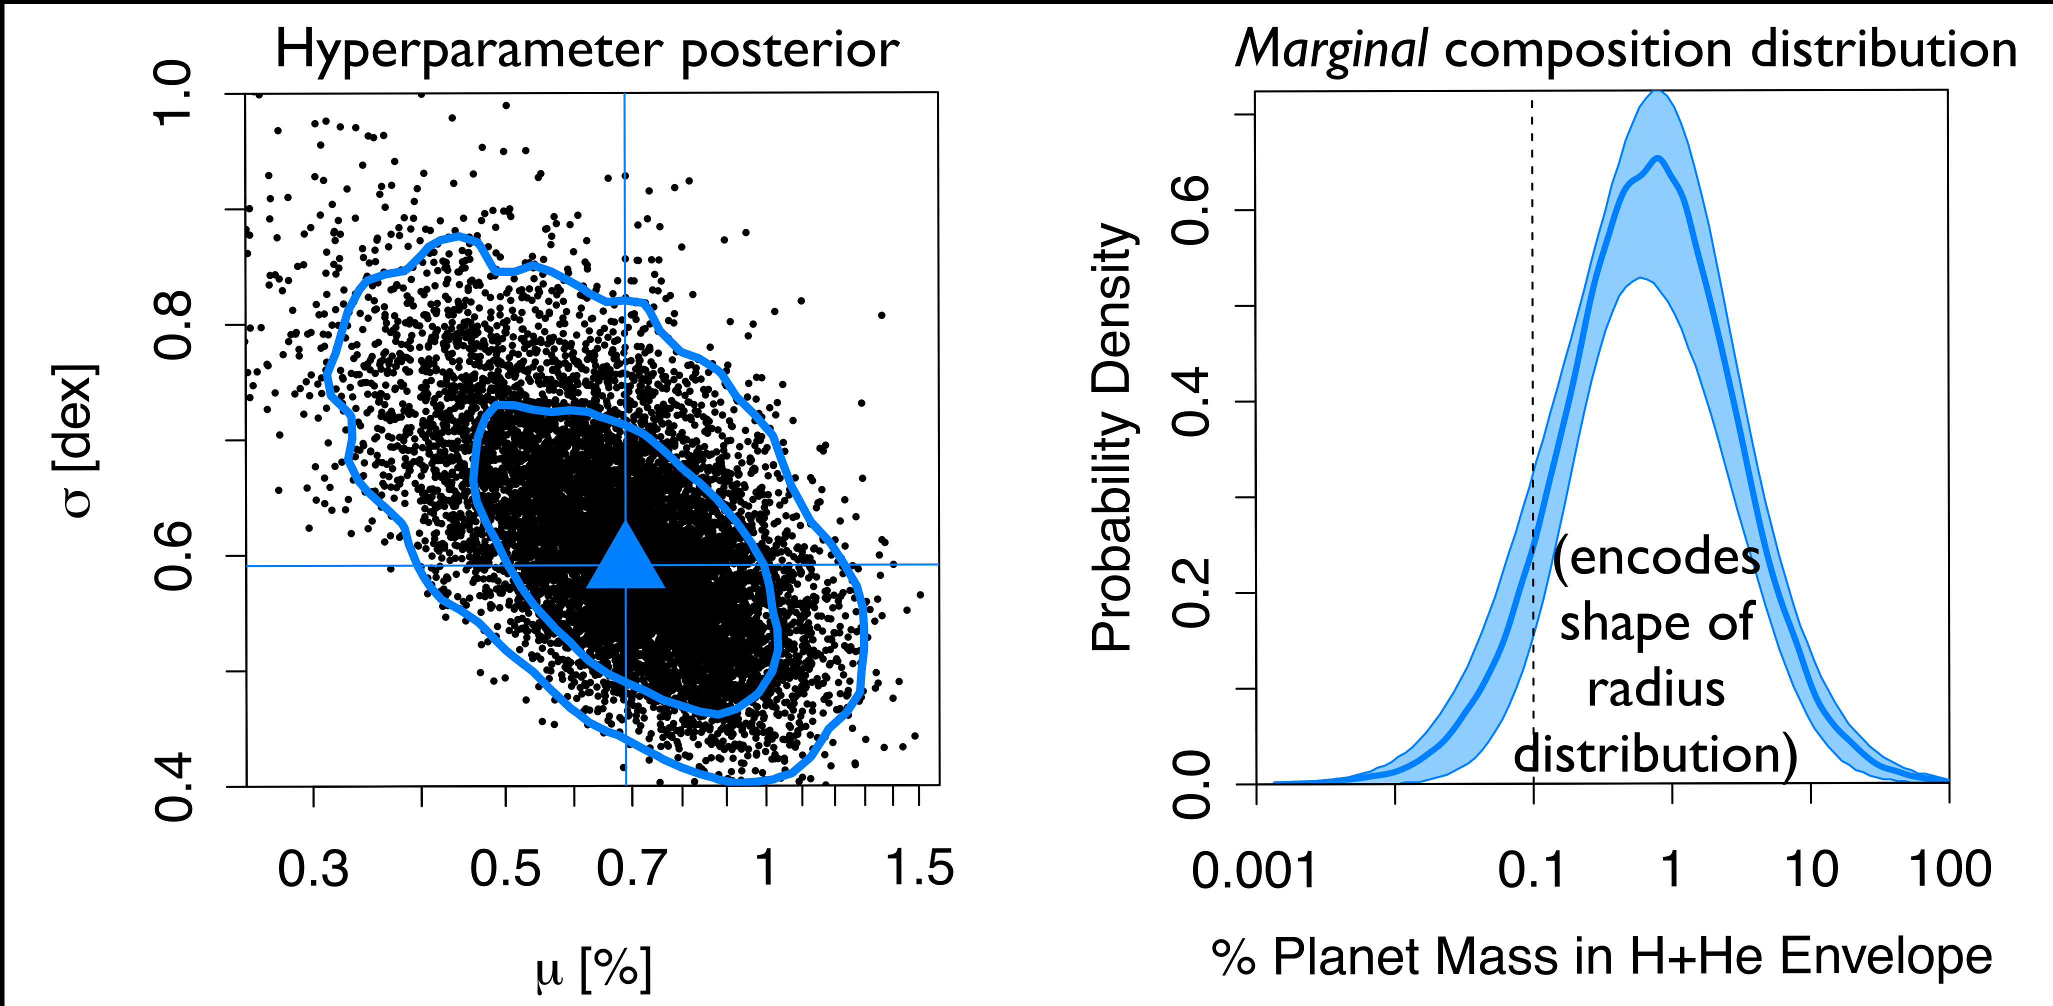

First composition distribution:

Results

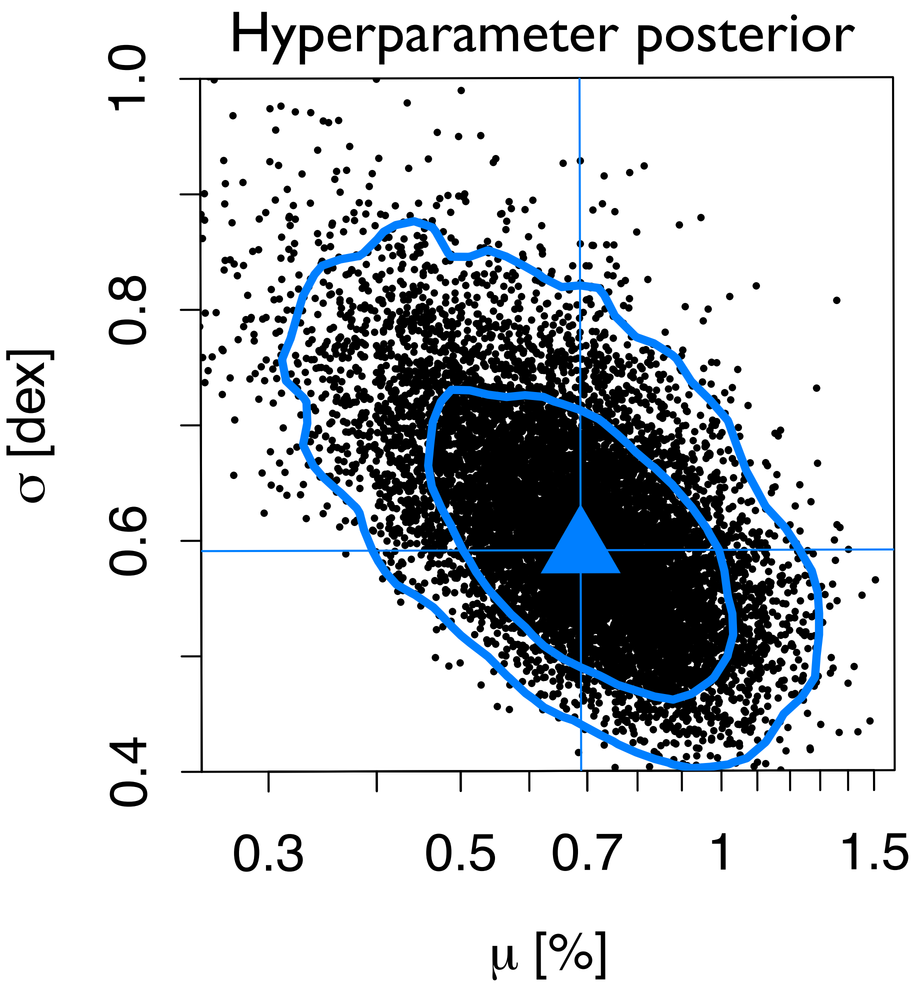

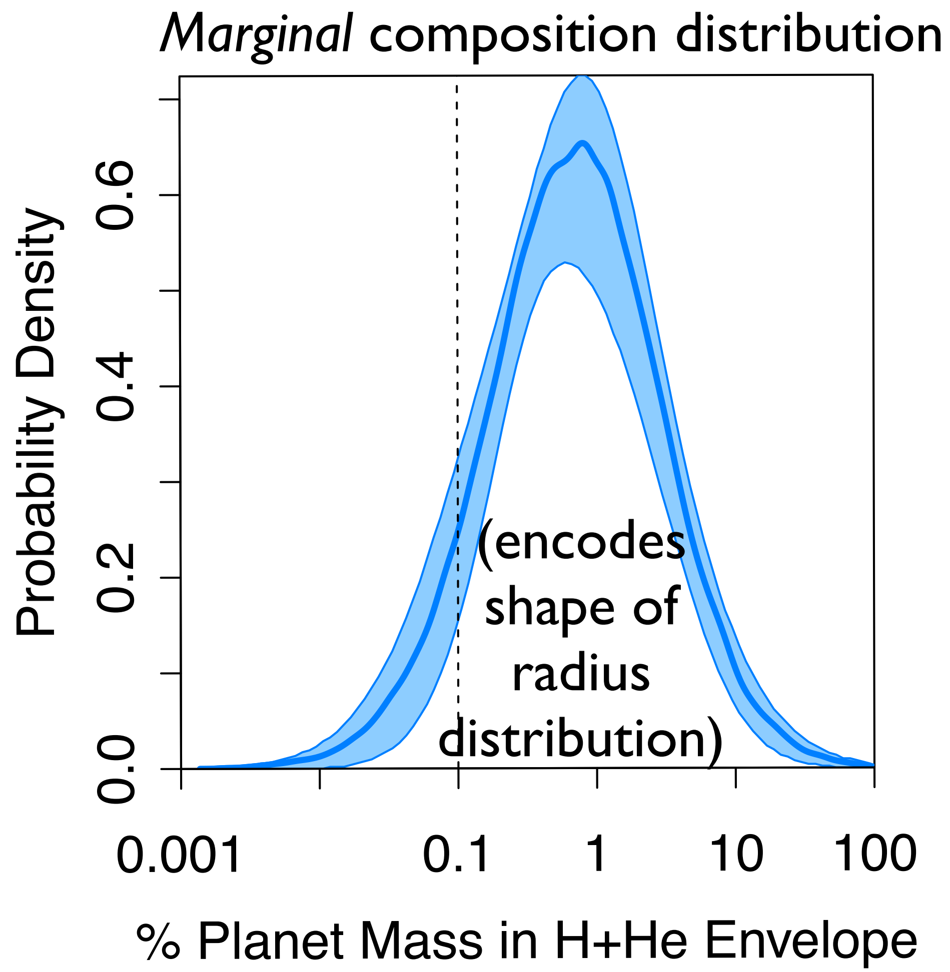

First composition distribution:

~ 1% envelope mass fractions are the most likely

Wolfgang & Lopez, submitted (http://arxiv.org/abs/1409.2982)

Implications

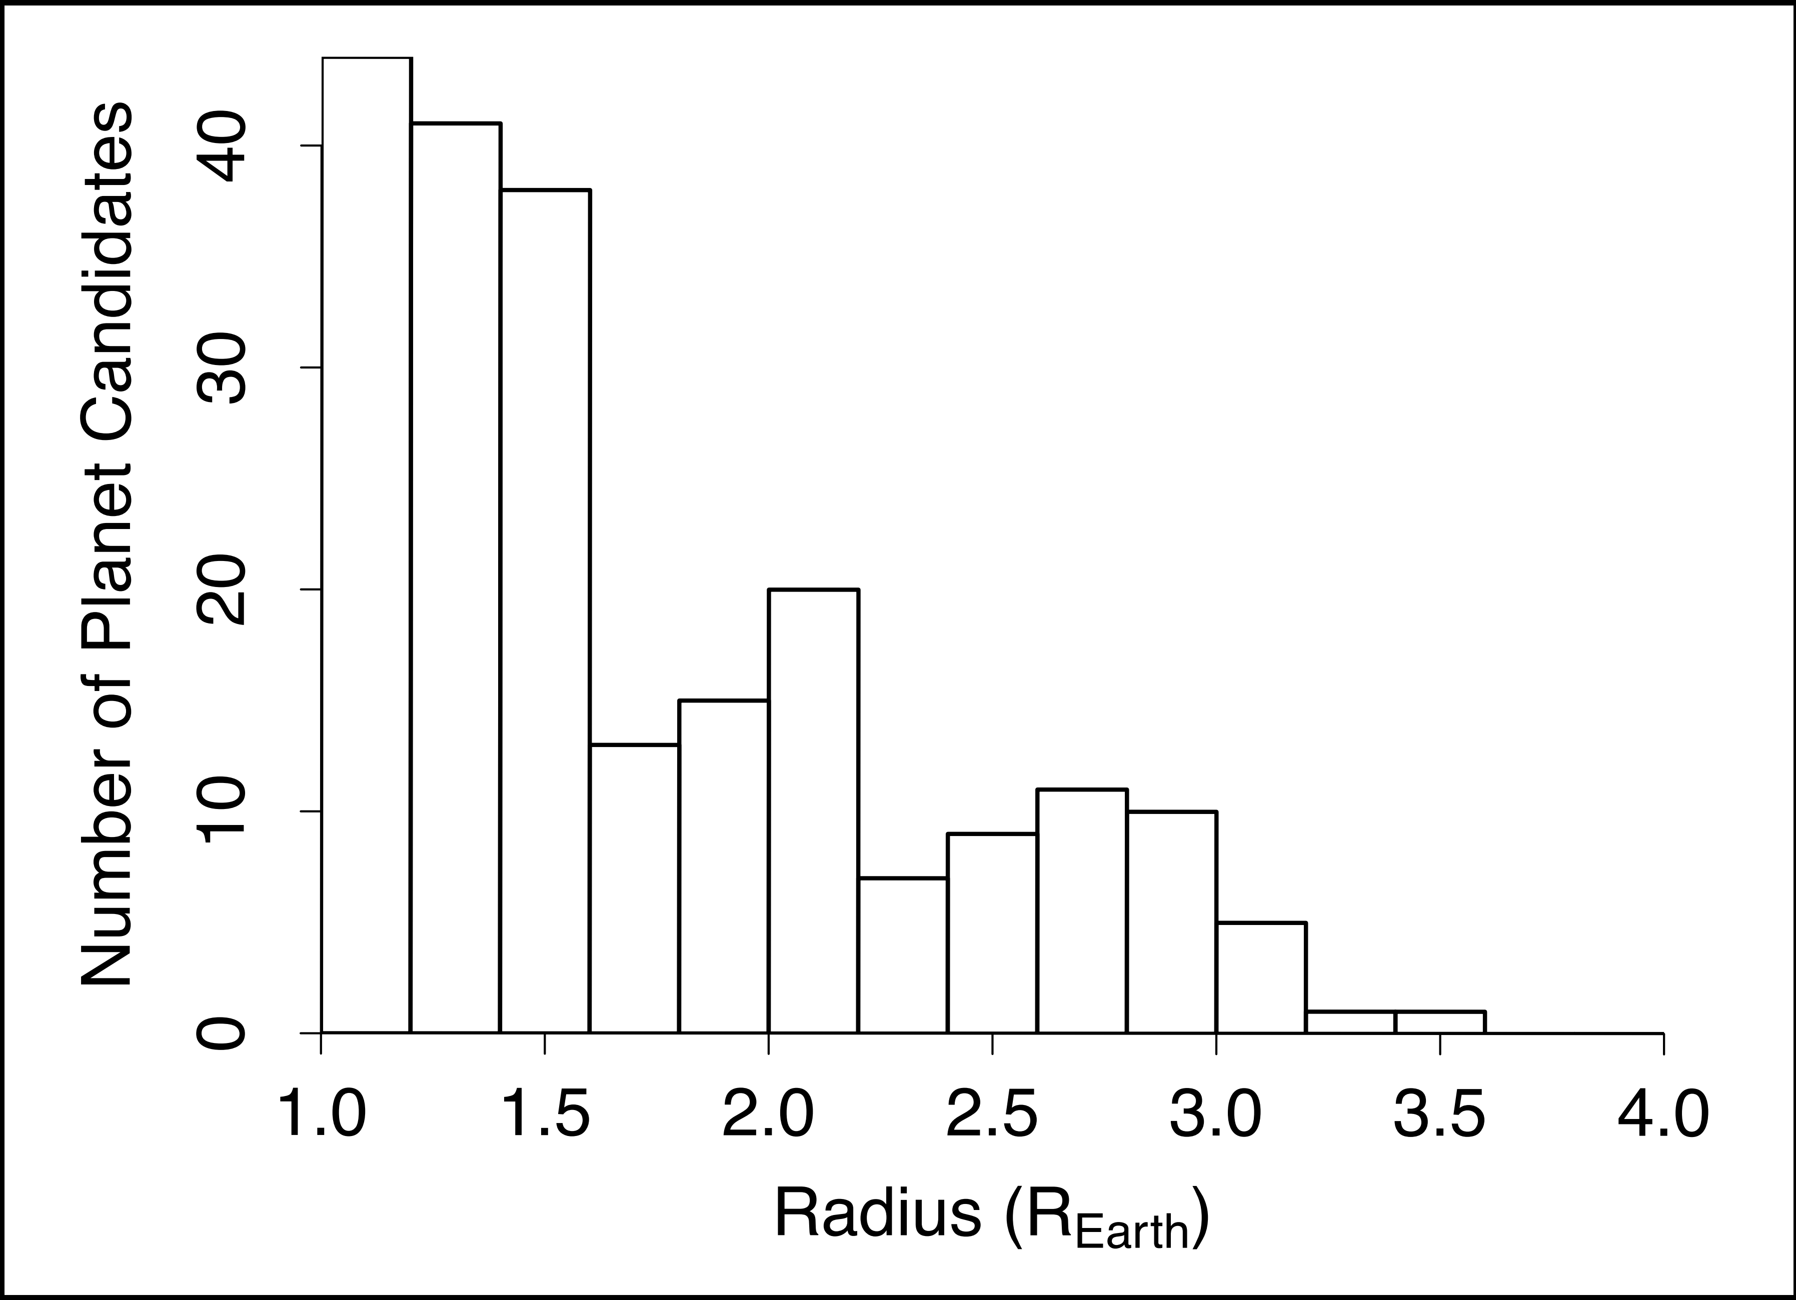

Implications

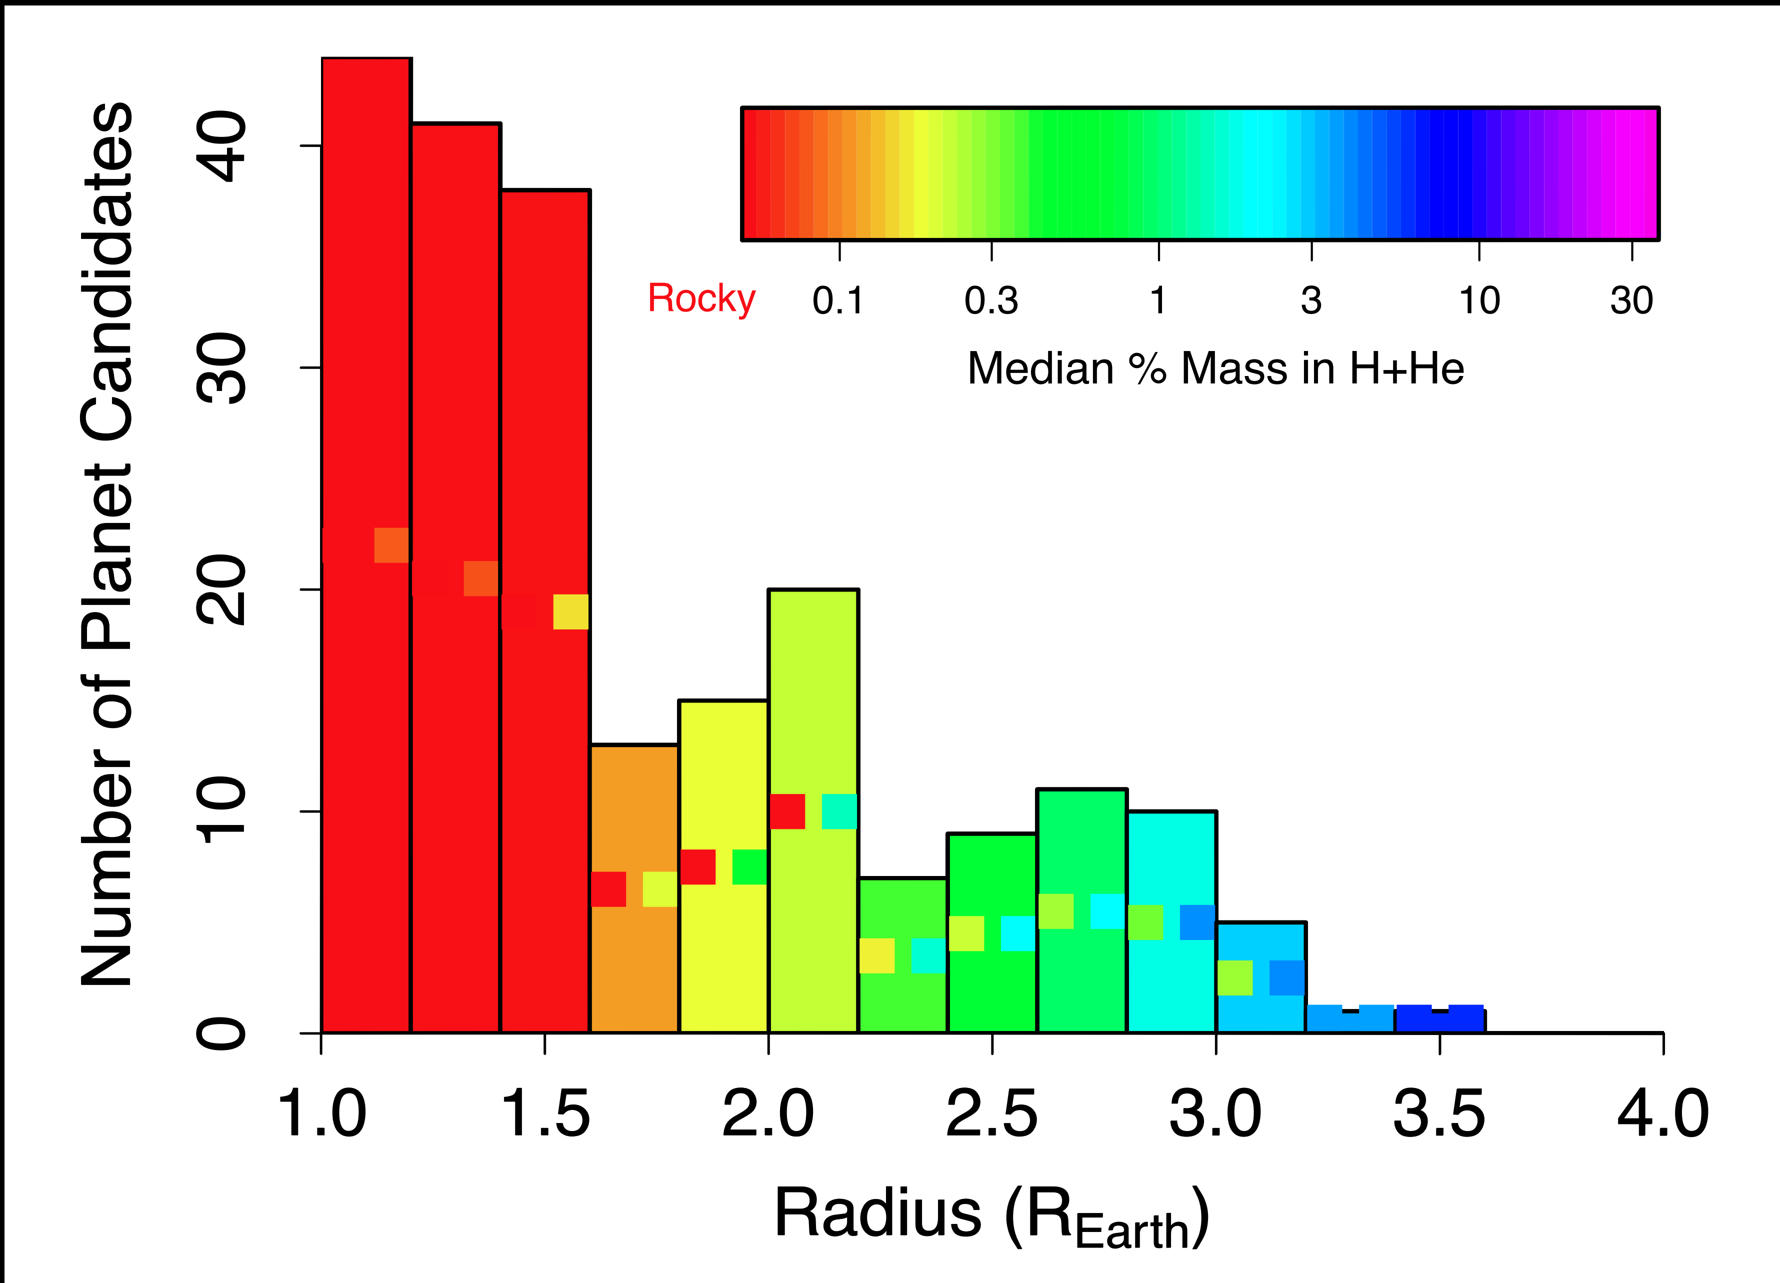

Implications

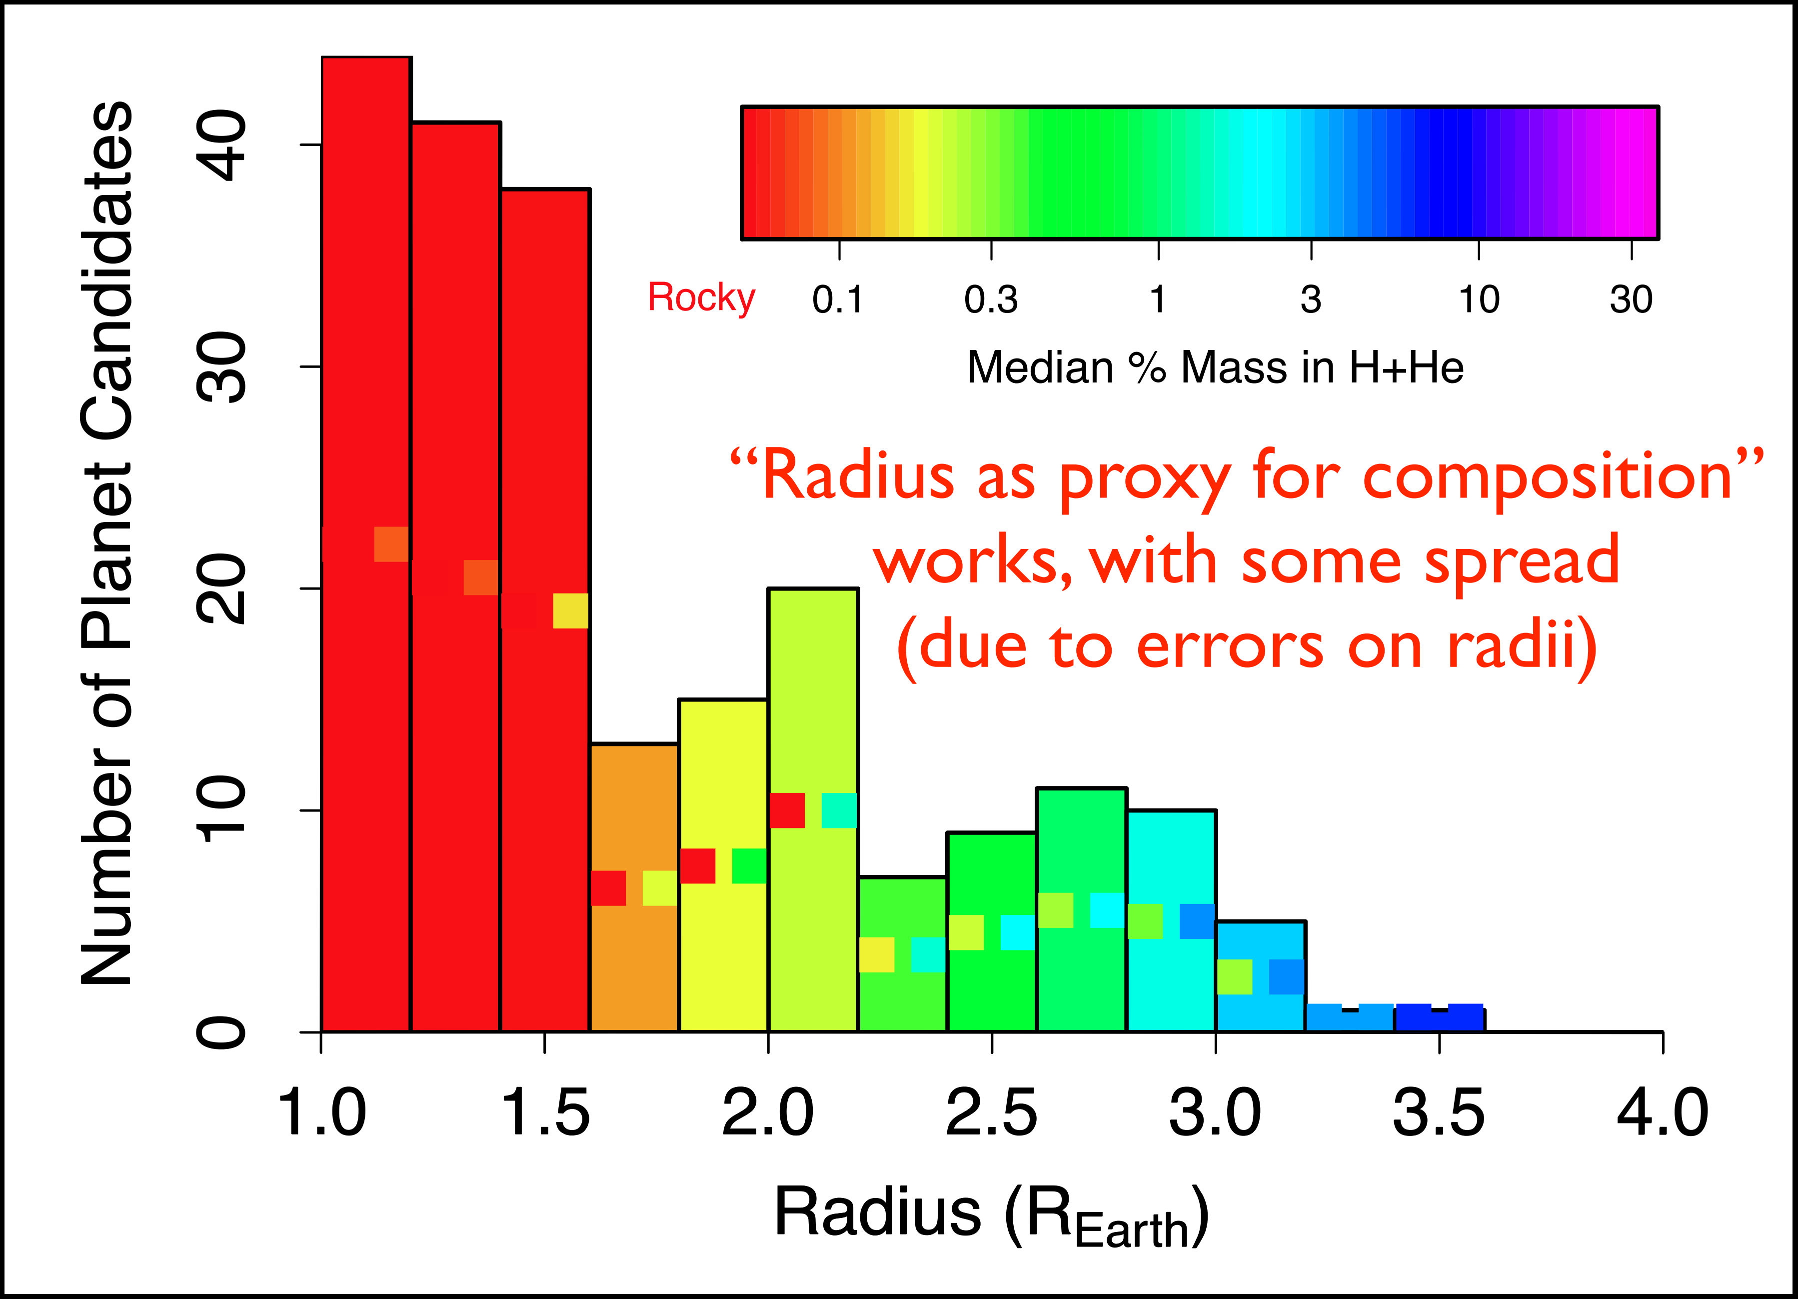

“Radius as proxy for composition” works, with some spread (due to errors on radii)

Implications

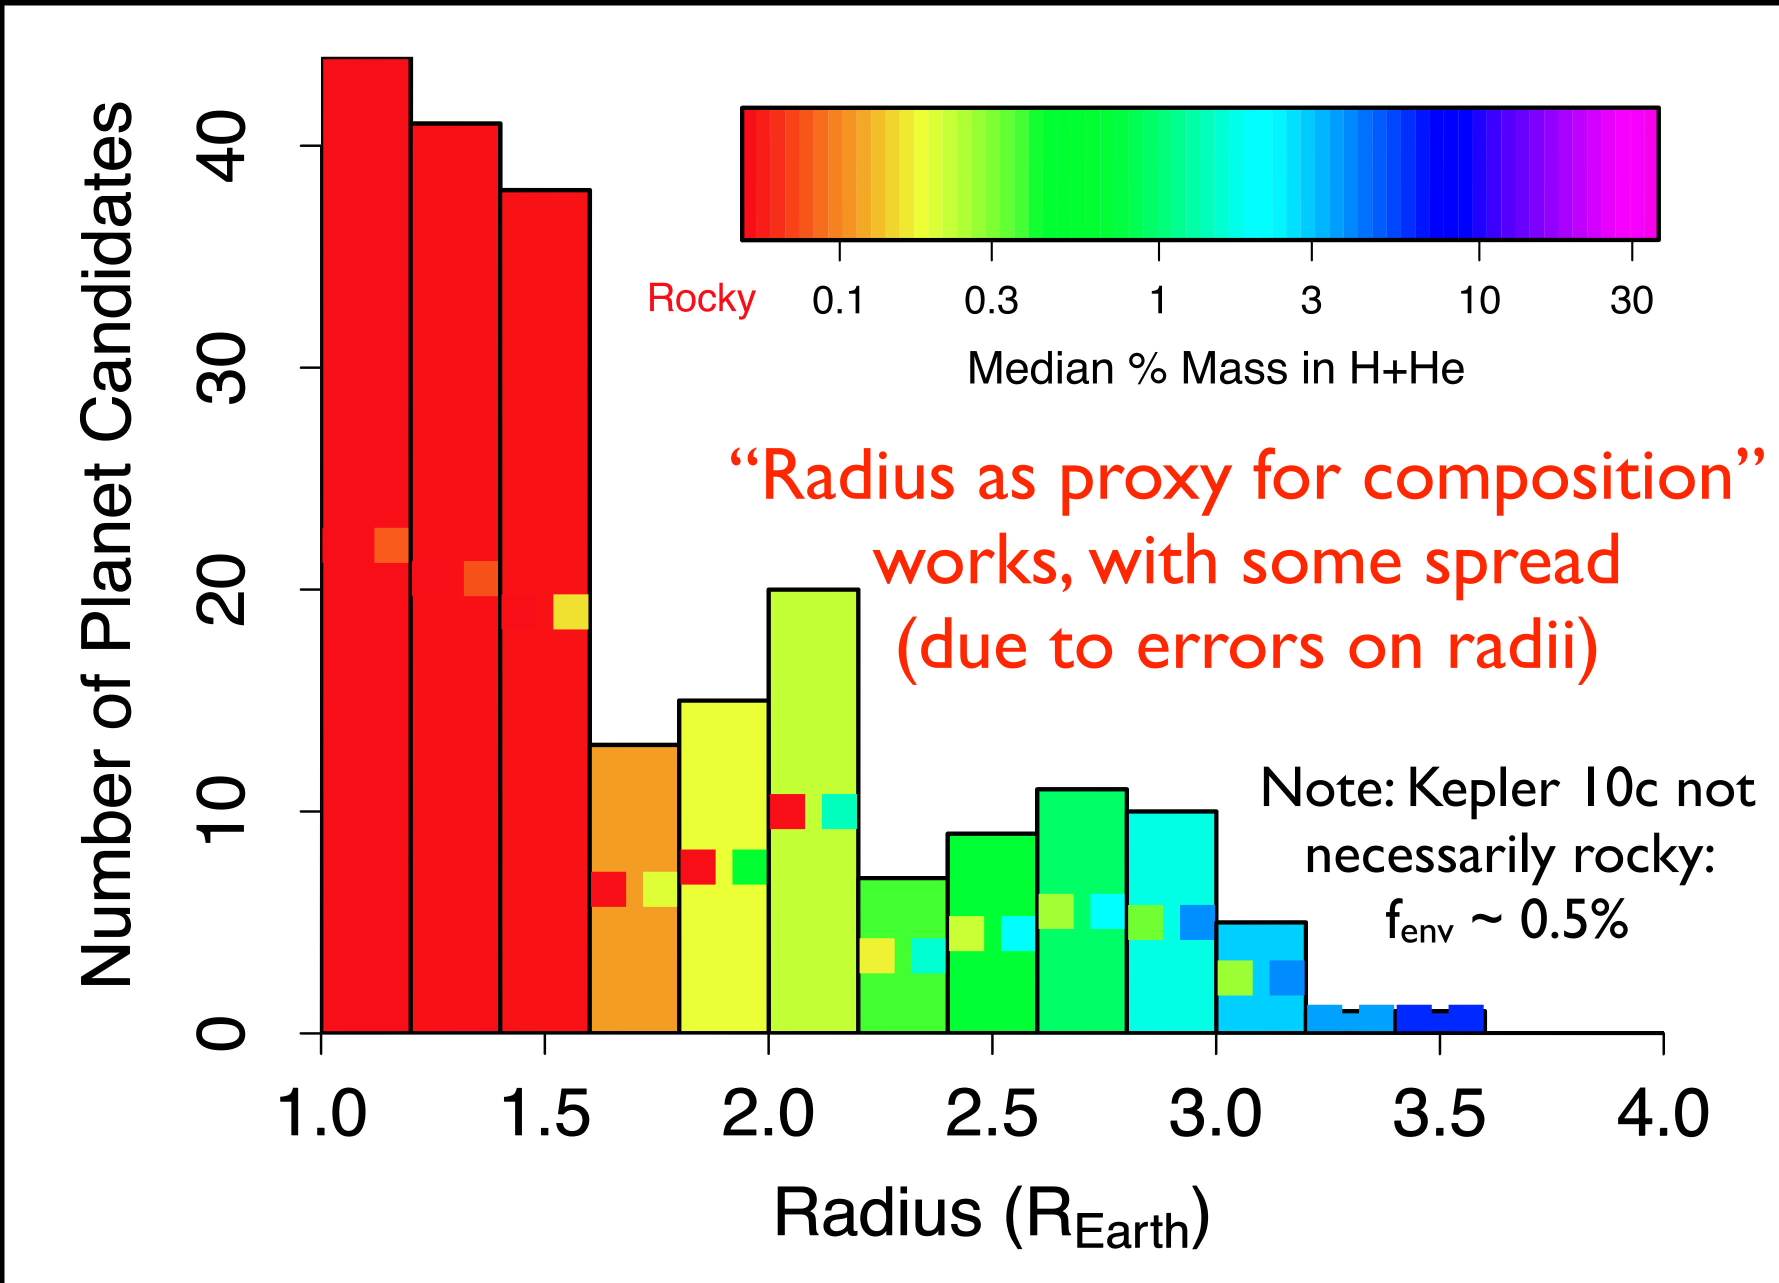

Implications

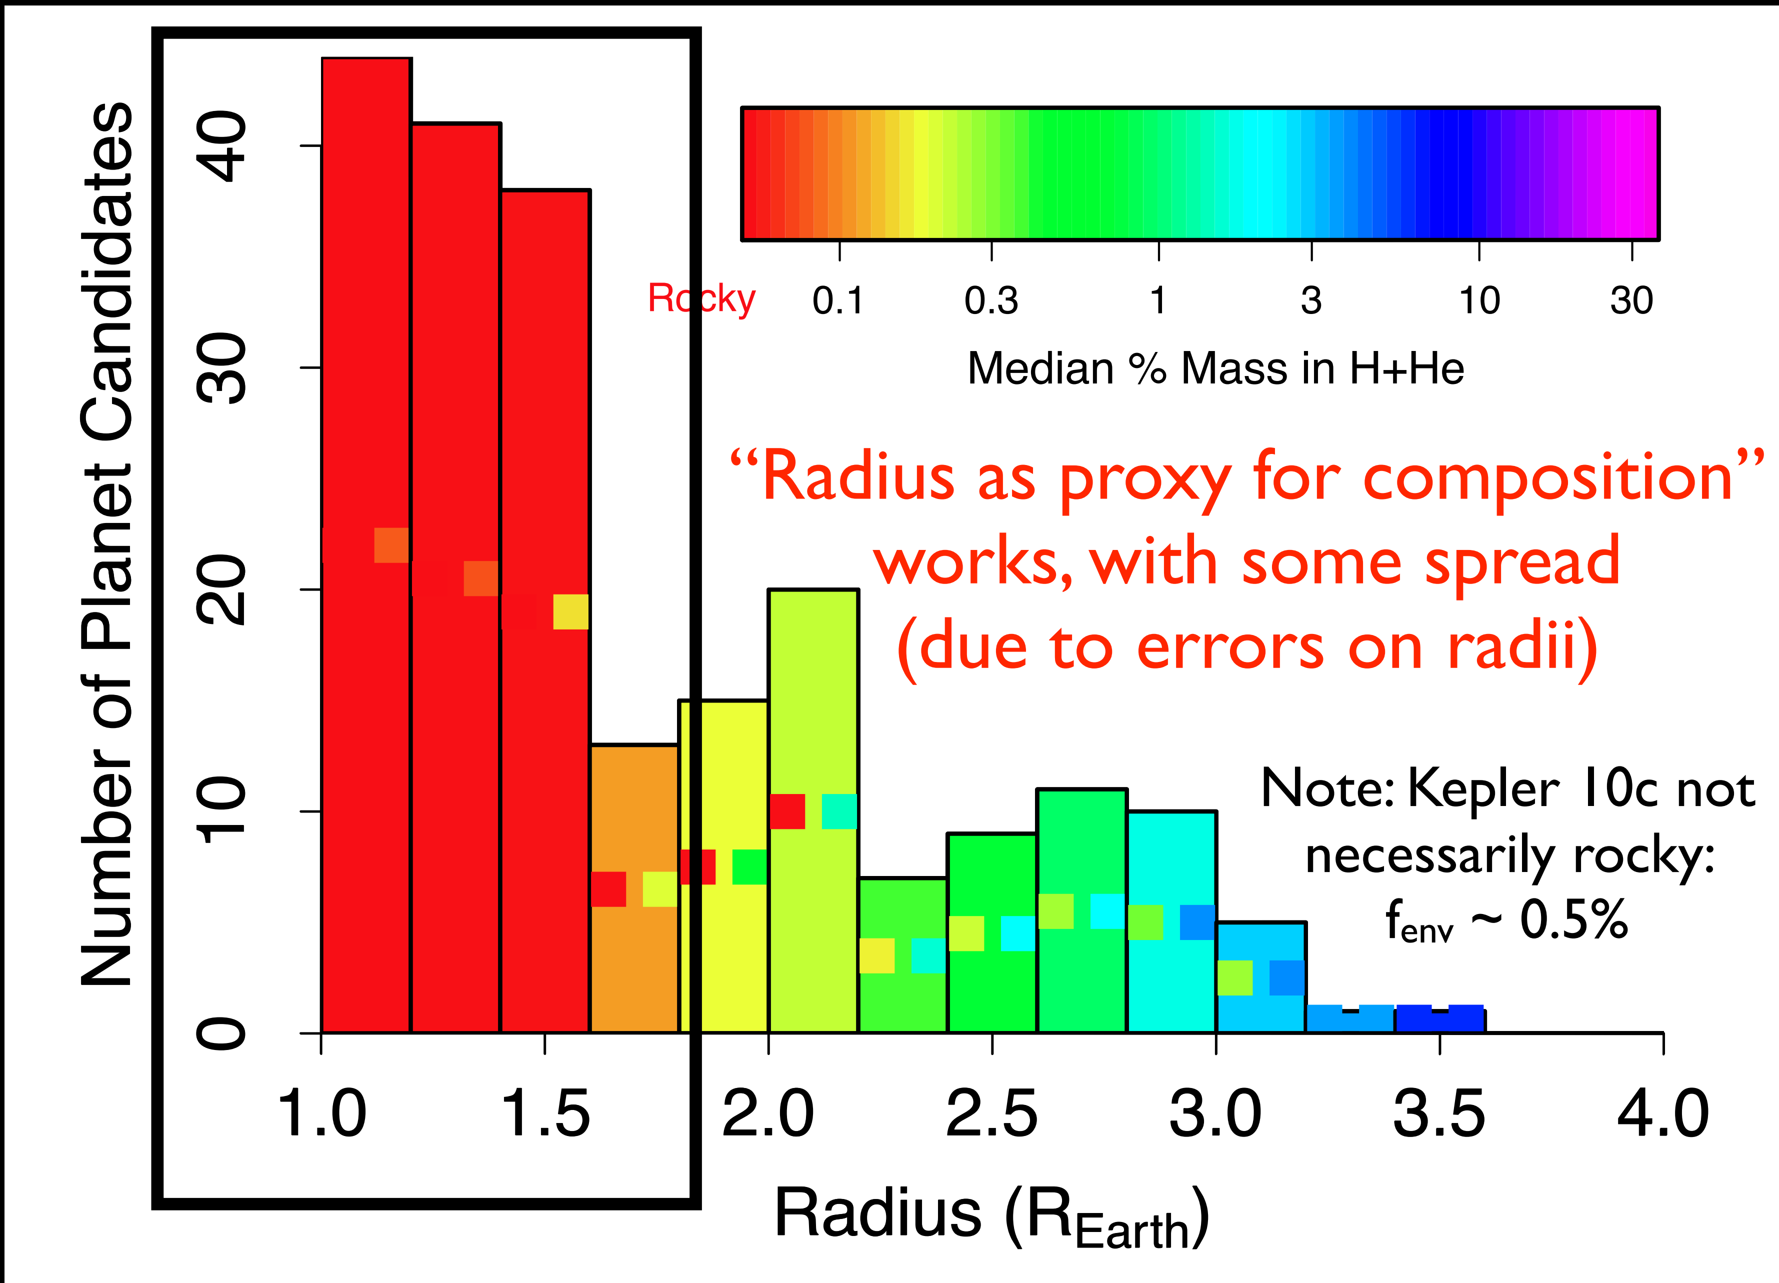

Rocky-Gaseous Transition

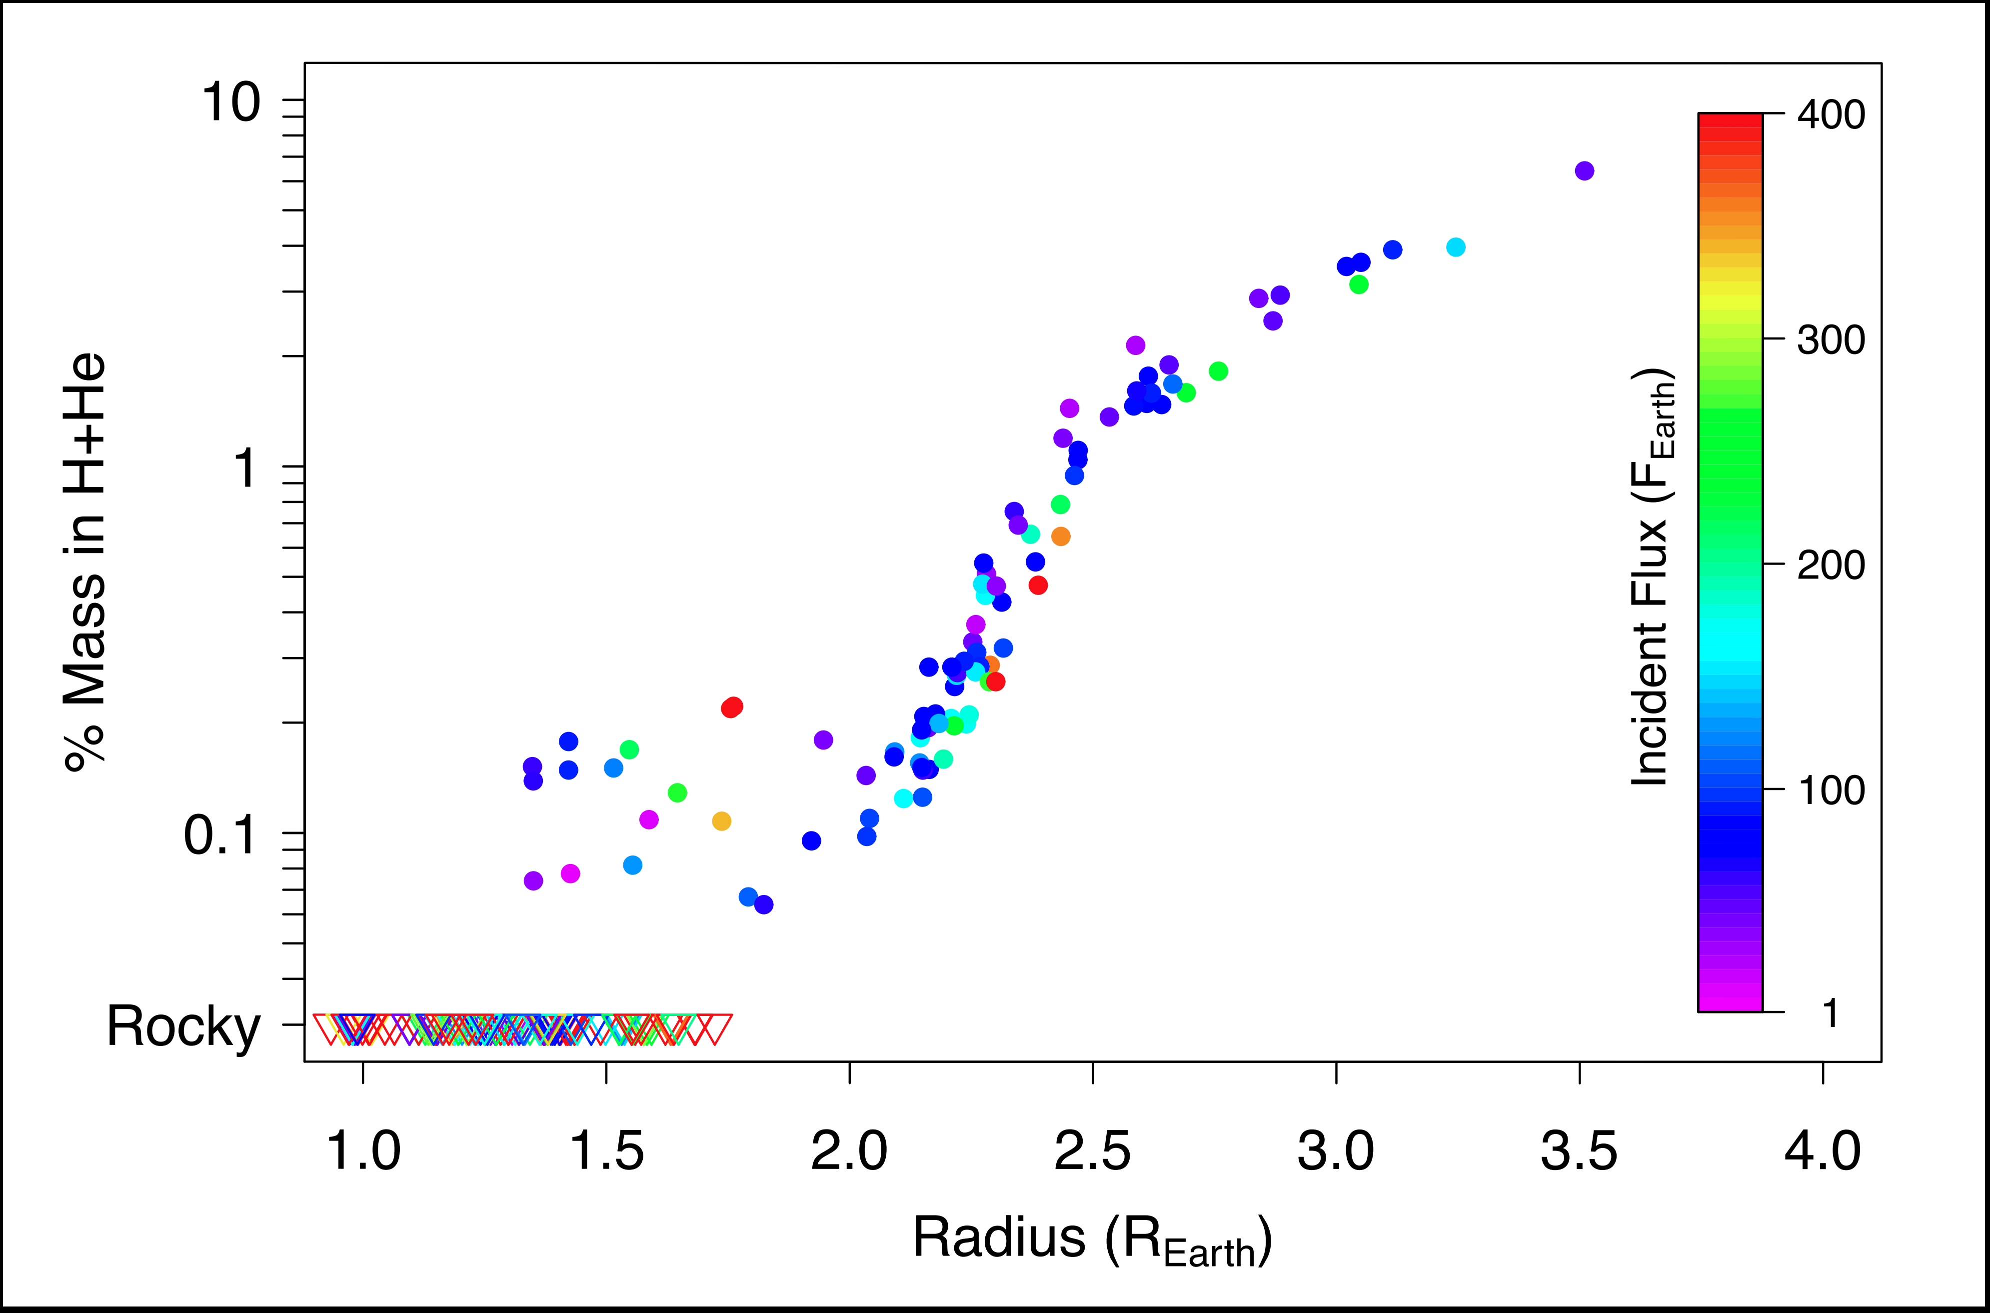

Rocky-Gaseous Transition

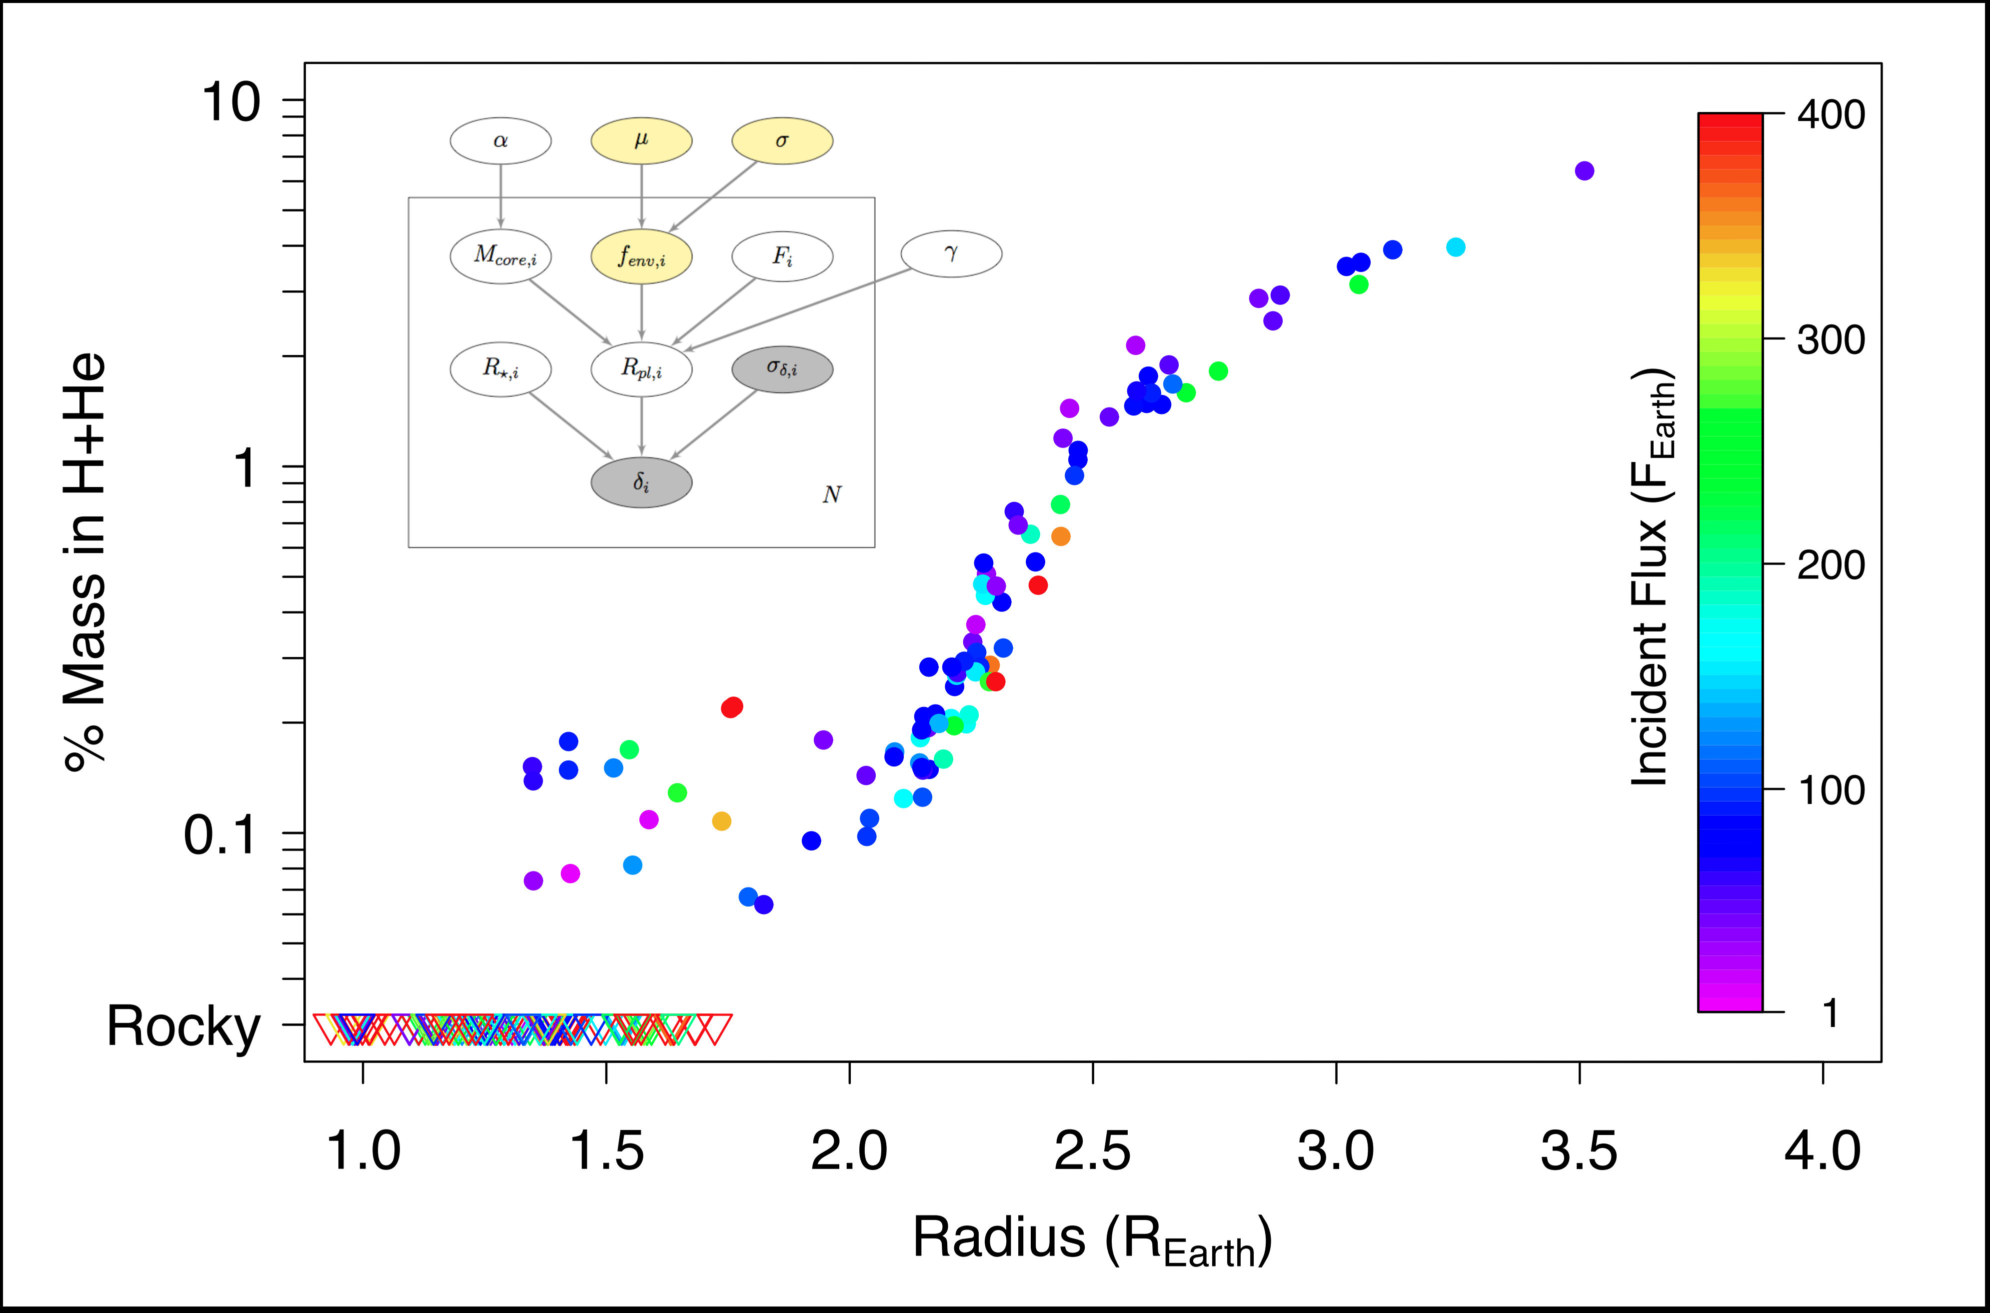

Rocky-Gaseous Transition

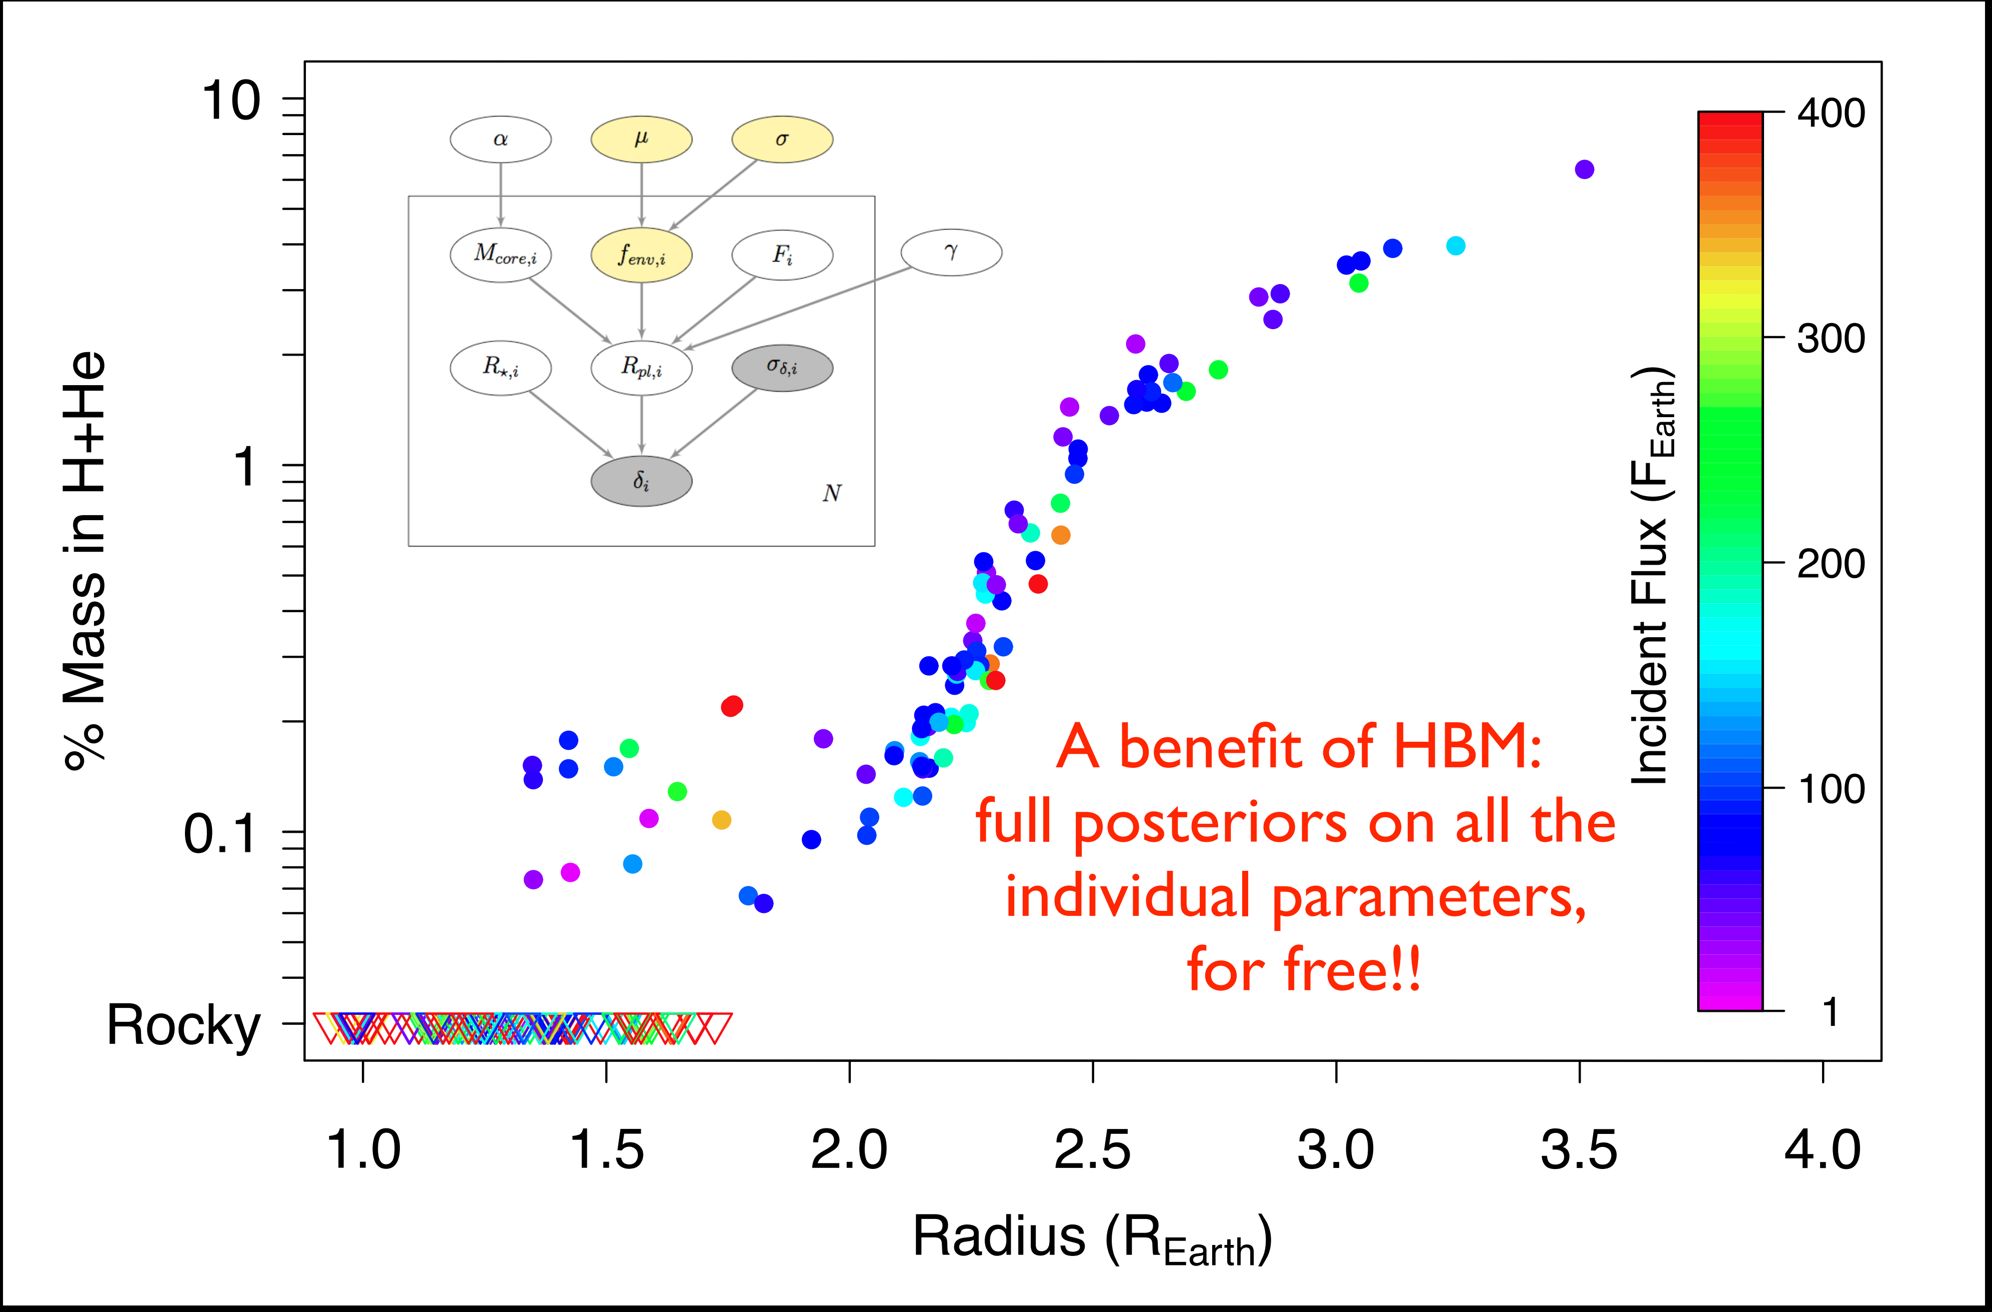

Rocky-Gaseous Transition

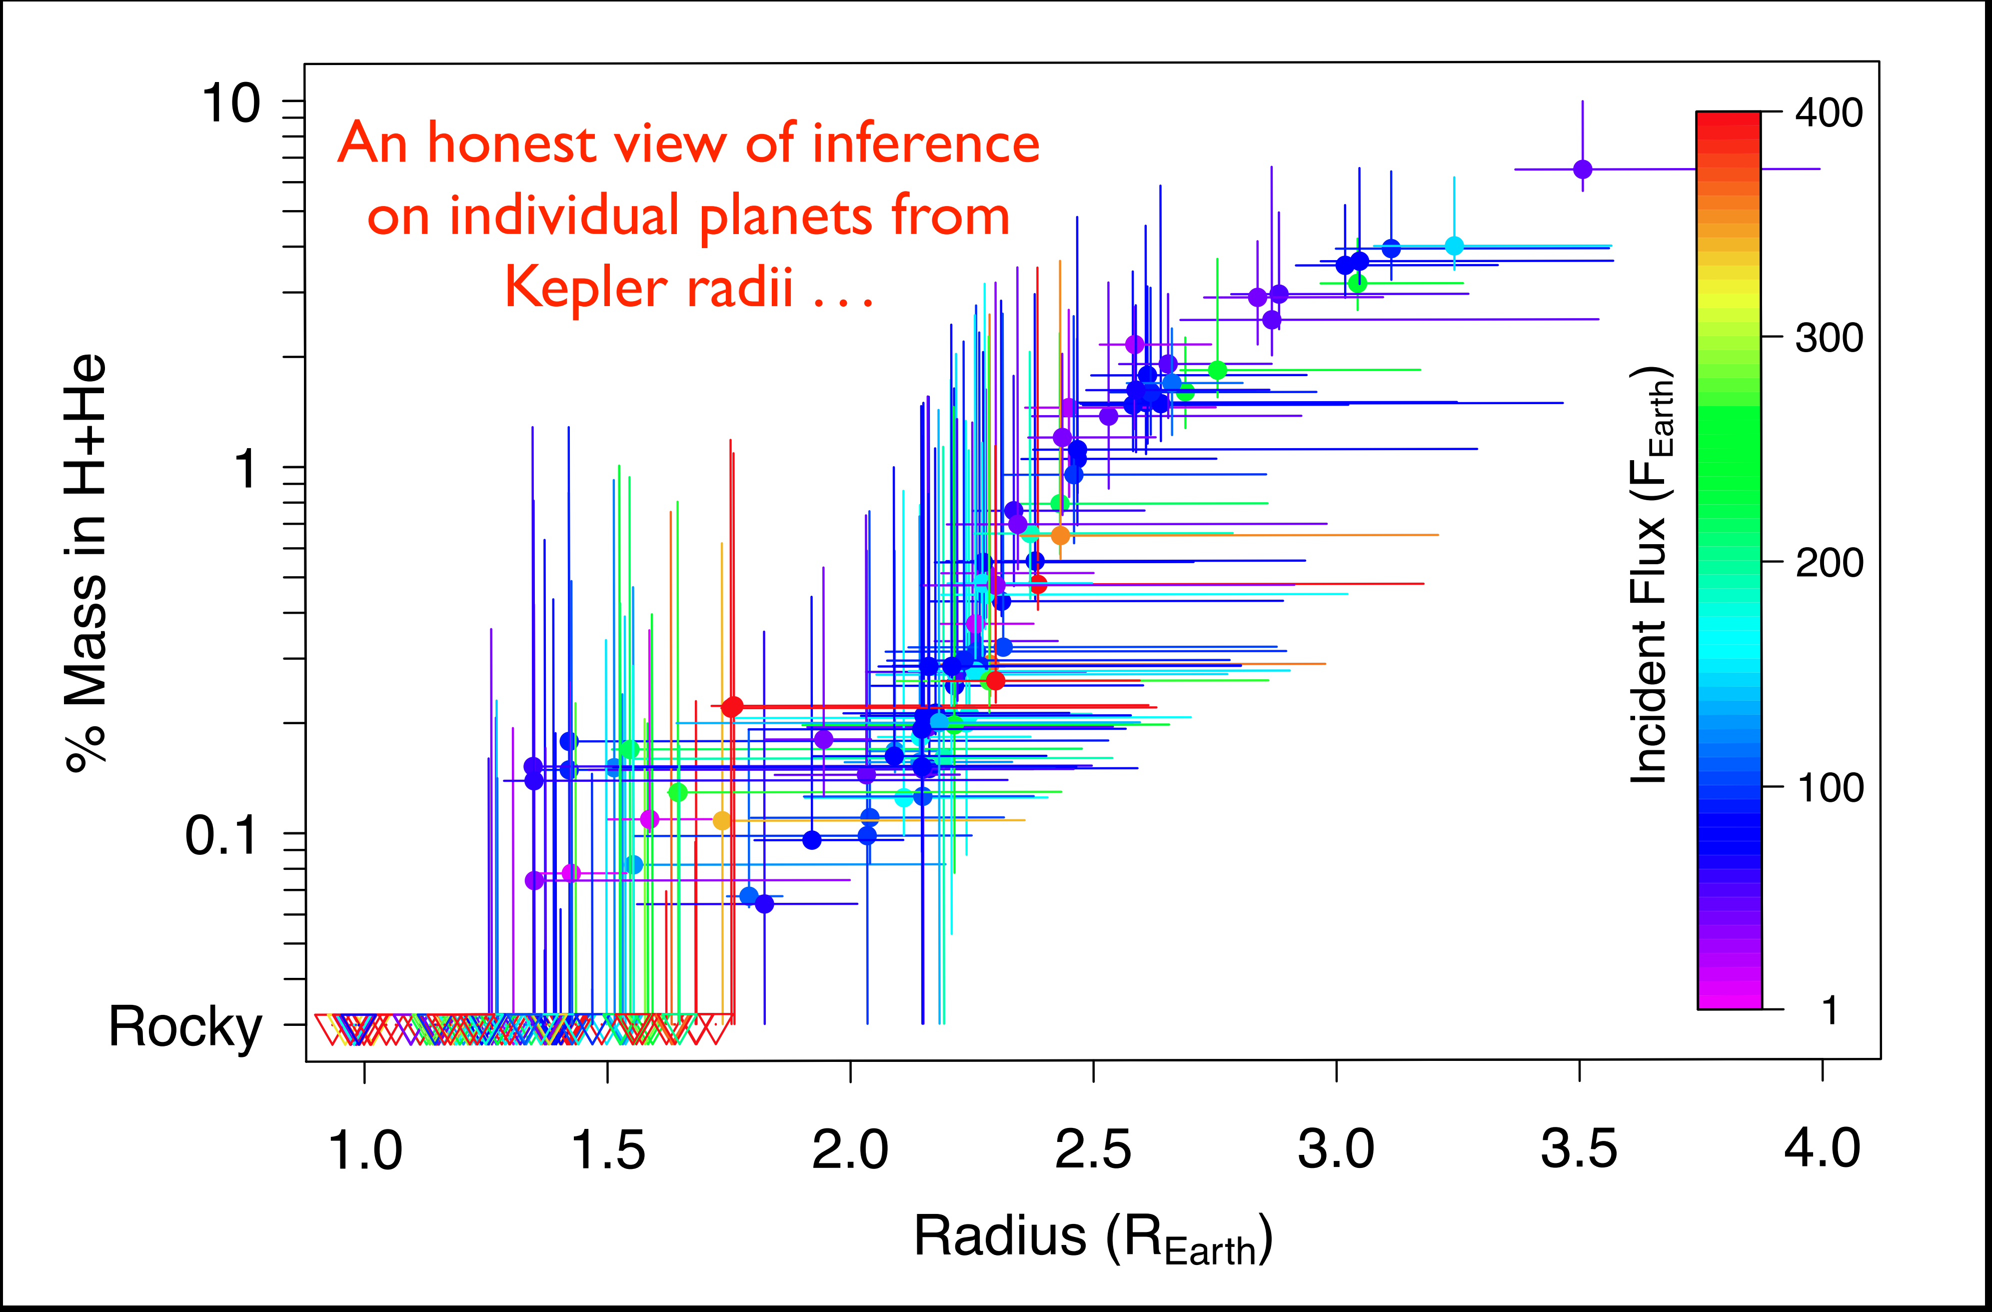

Rocky-Gaseous Transition

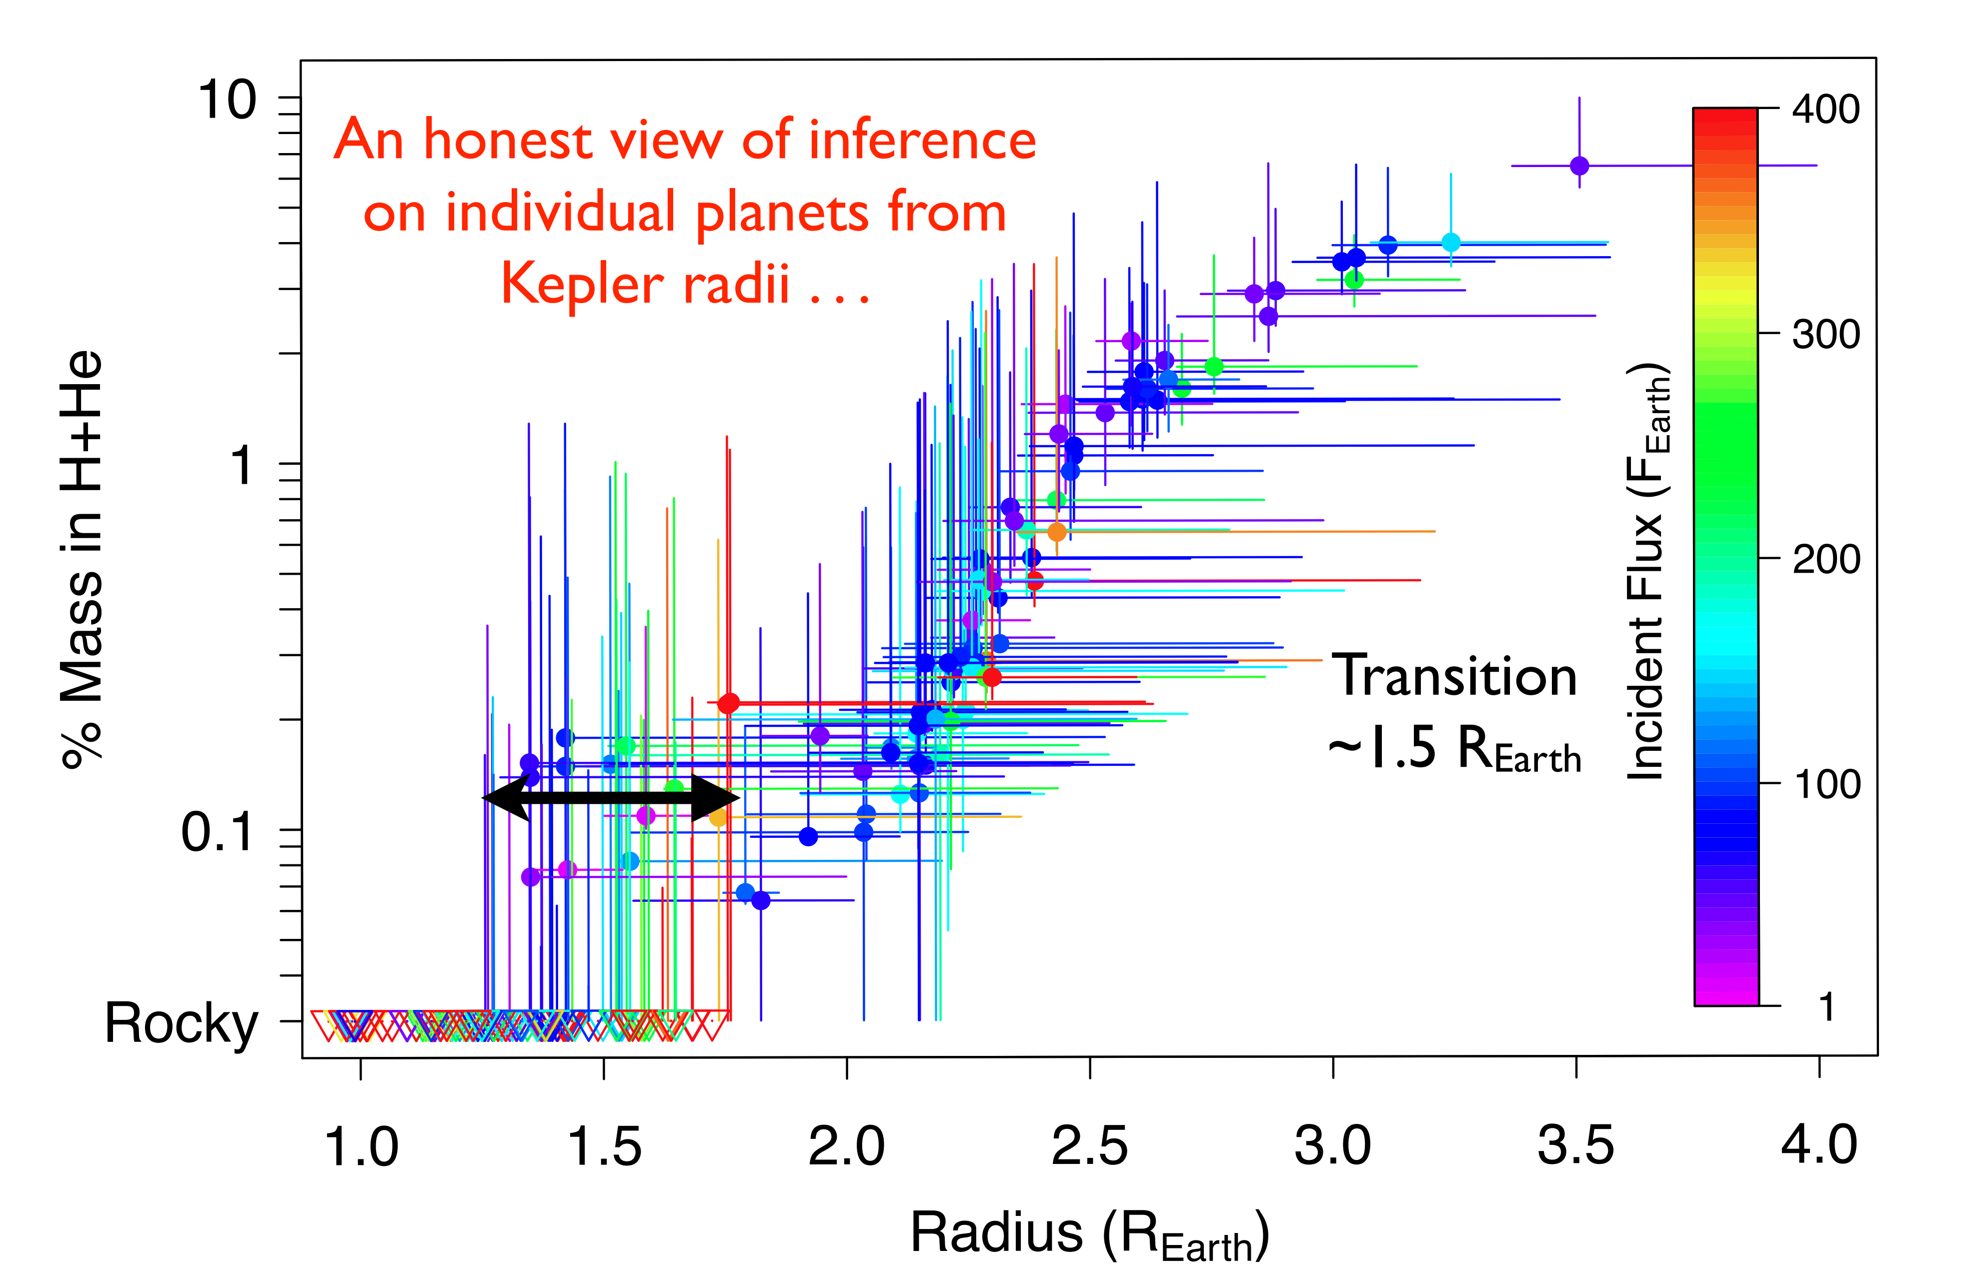

Rocky-Gaseous Transition

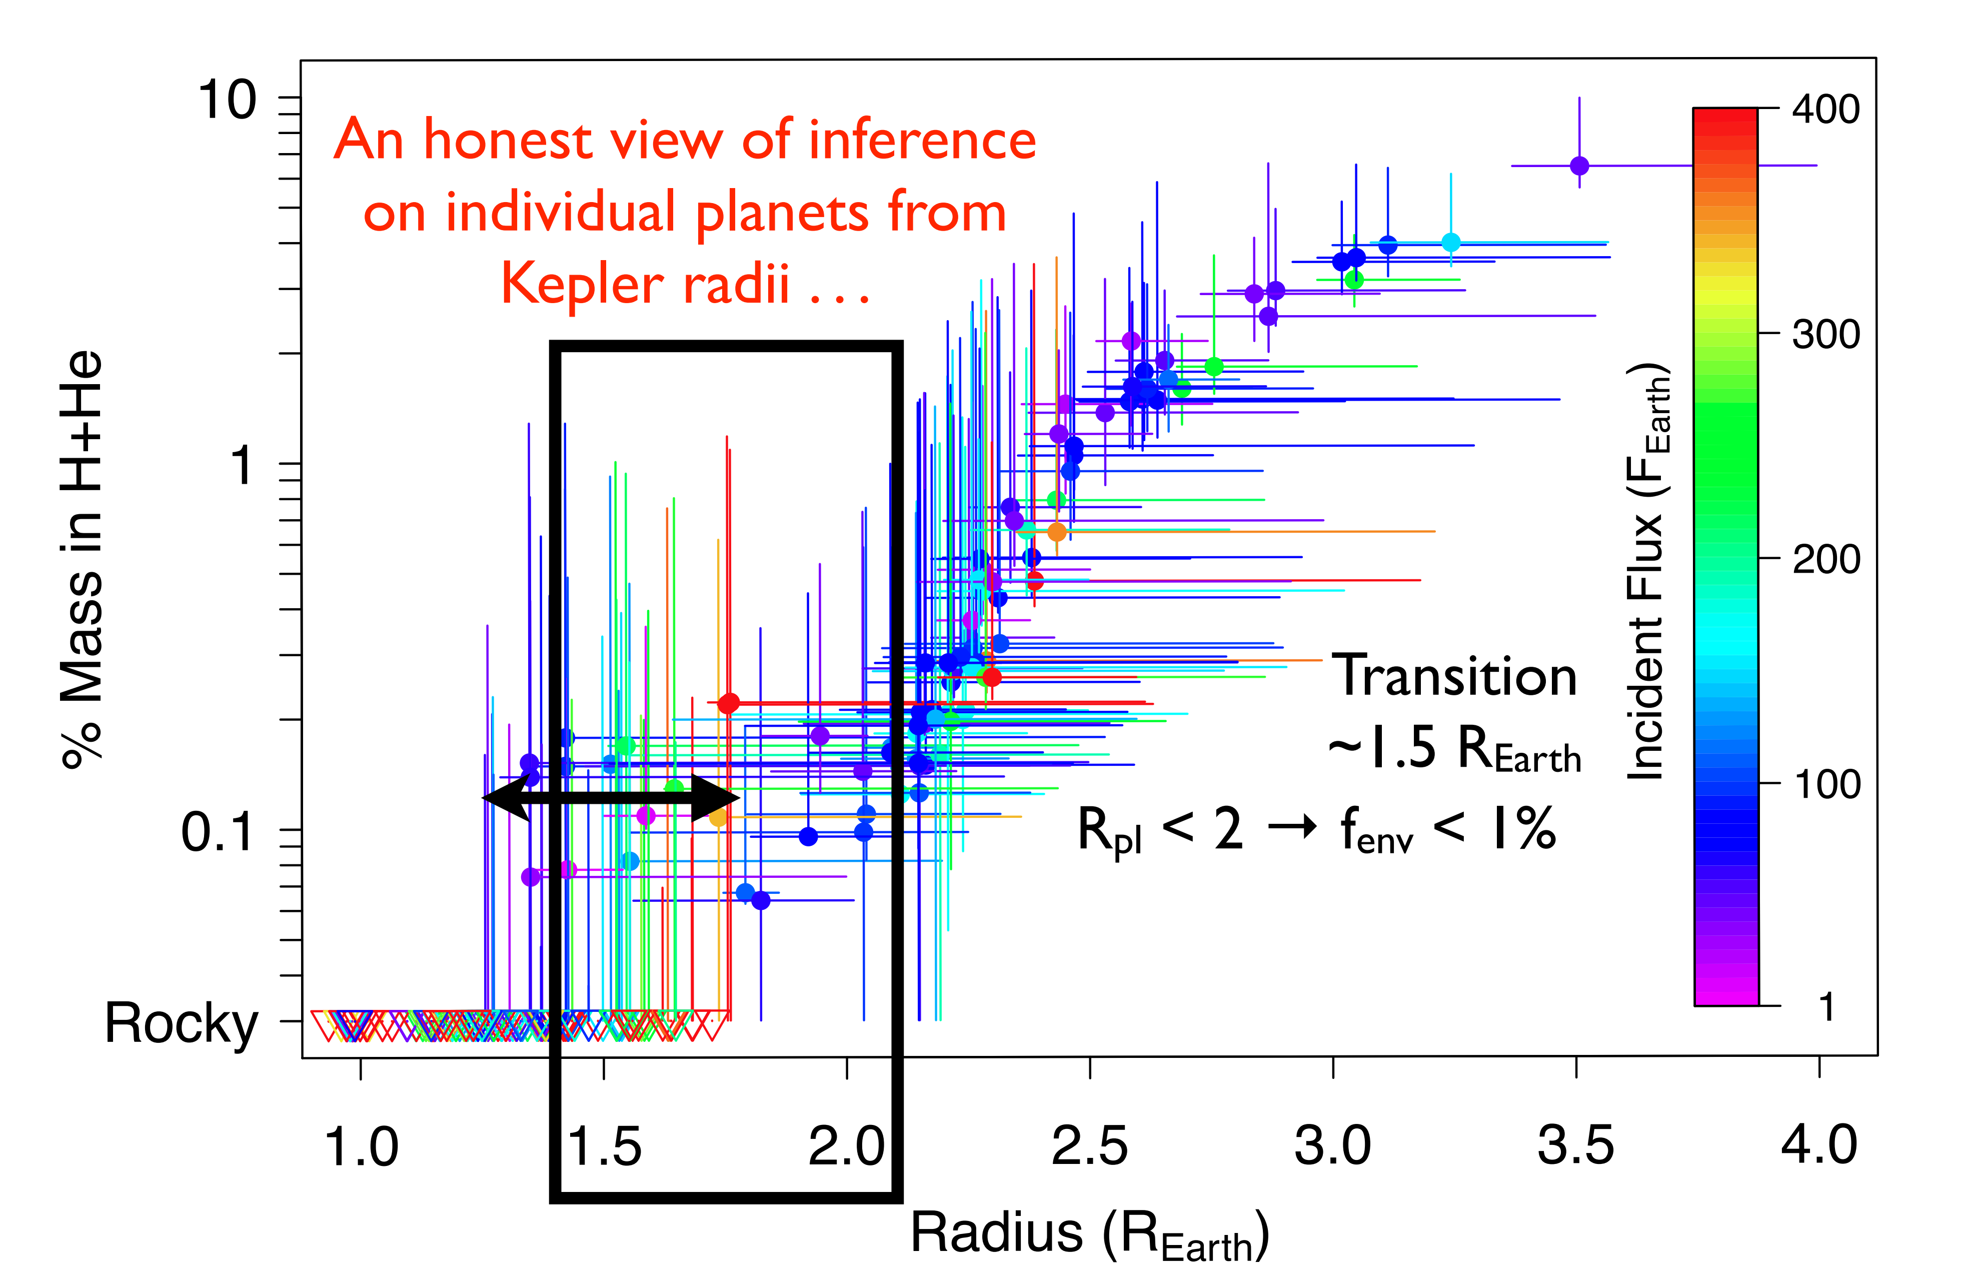

Rocky-Gaseous Transition

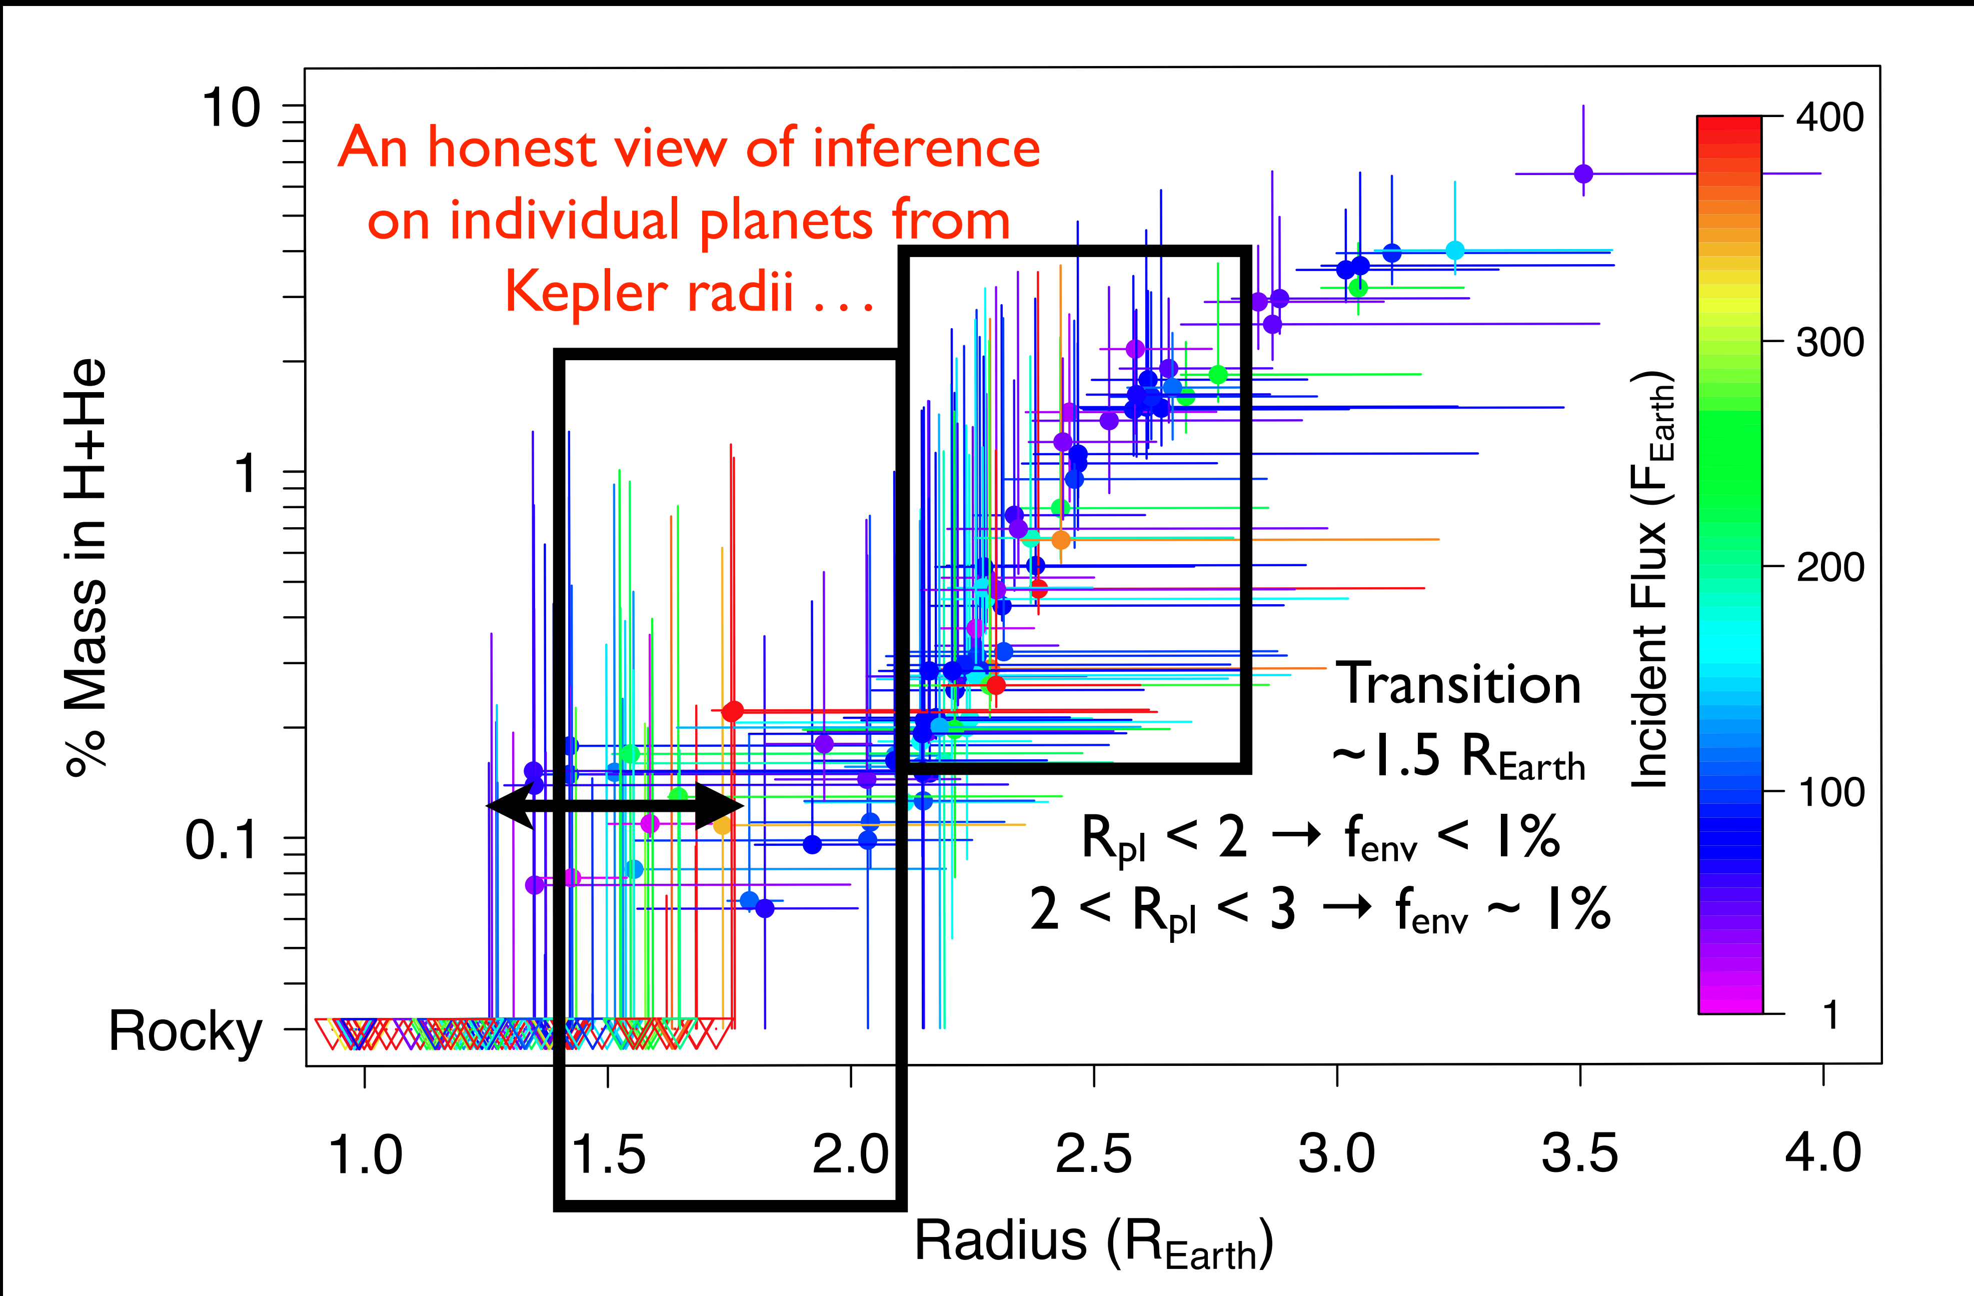

Rocky-Gaseous Transition

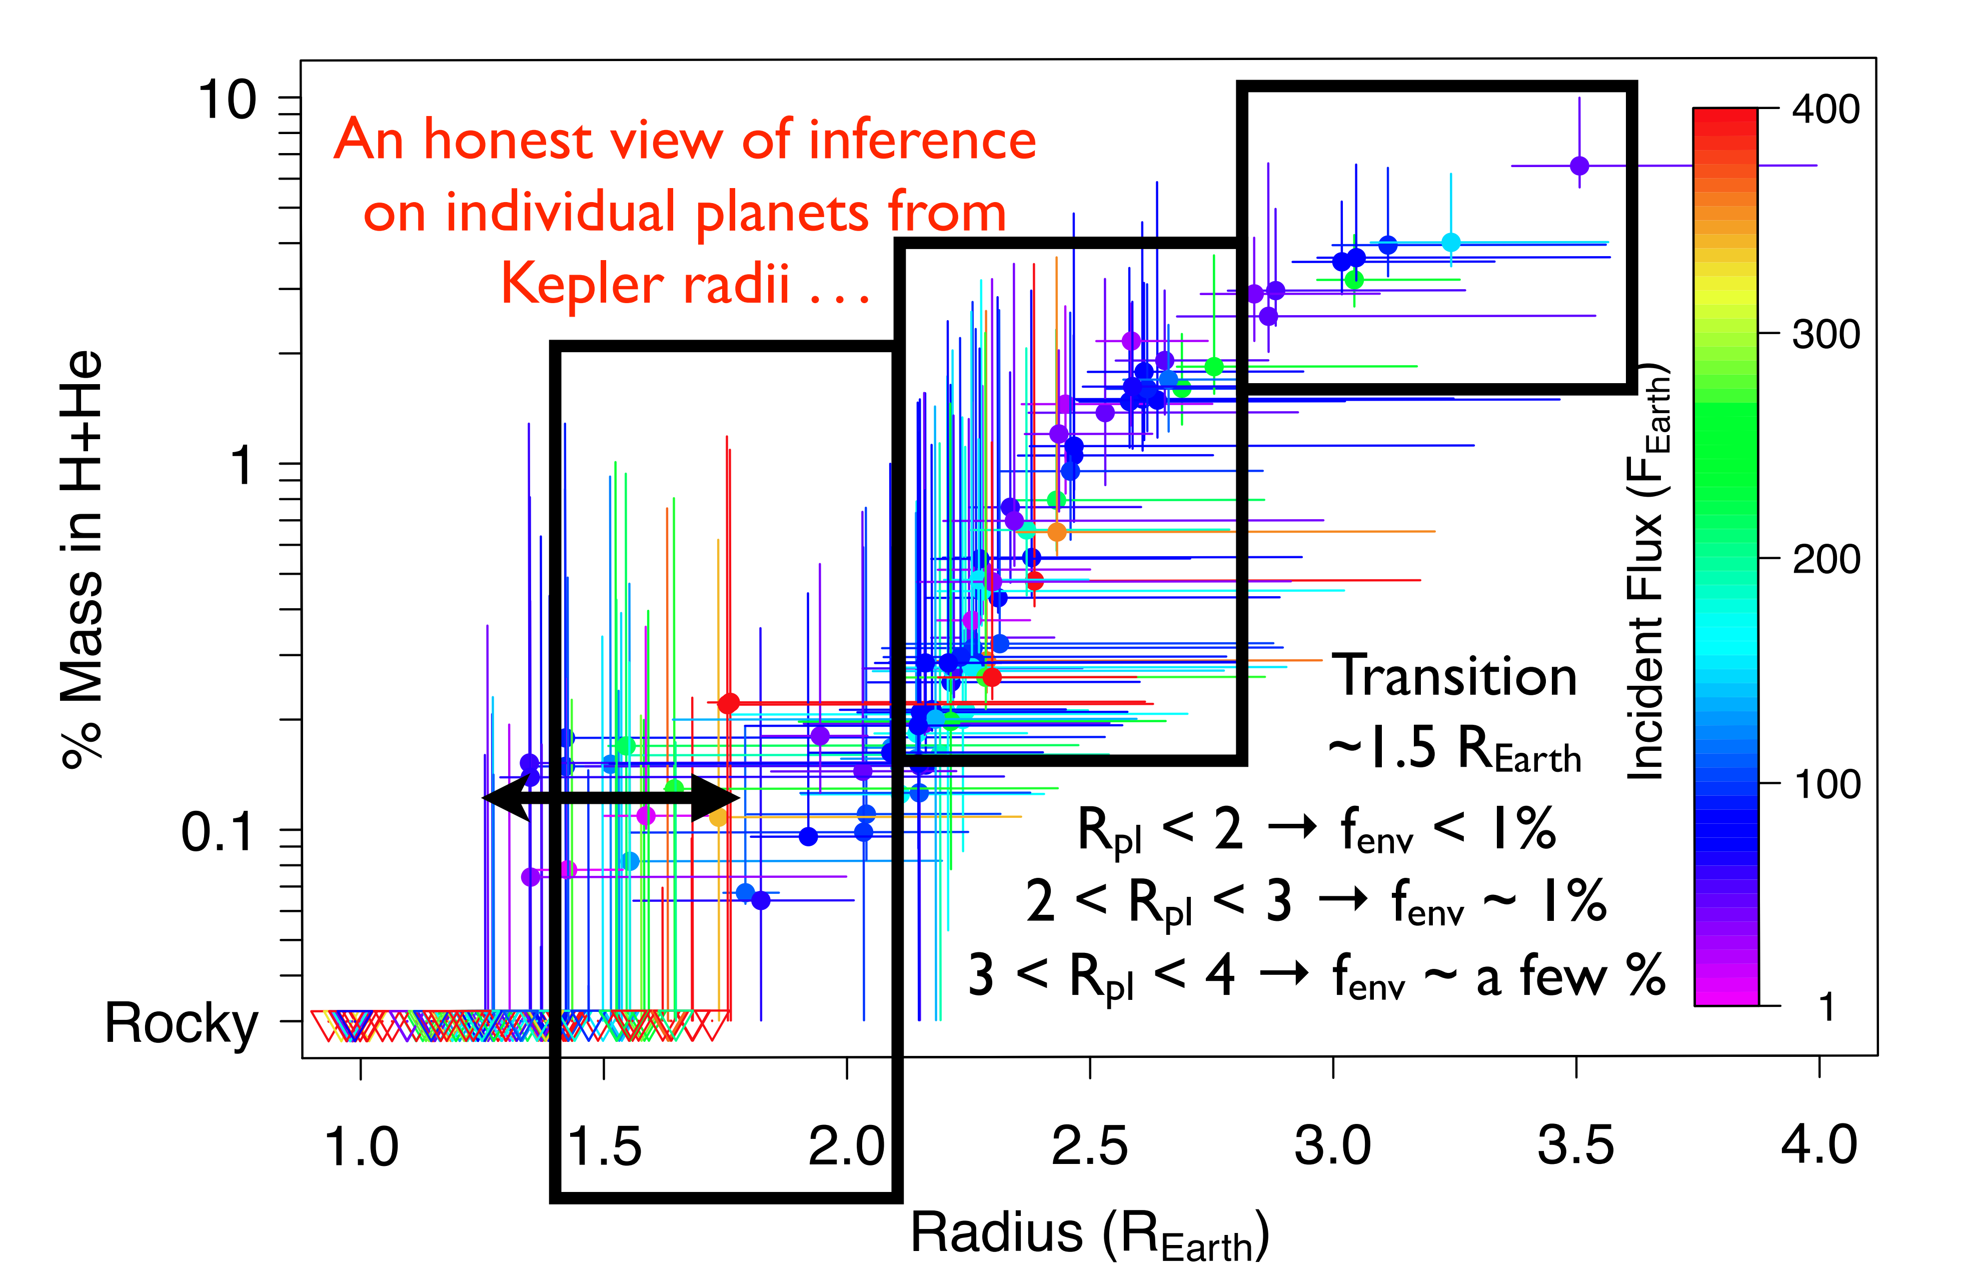

Mass-Radius Relationship

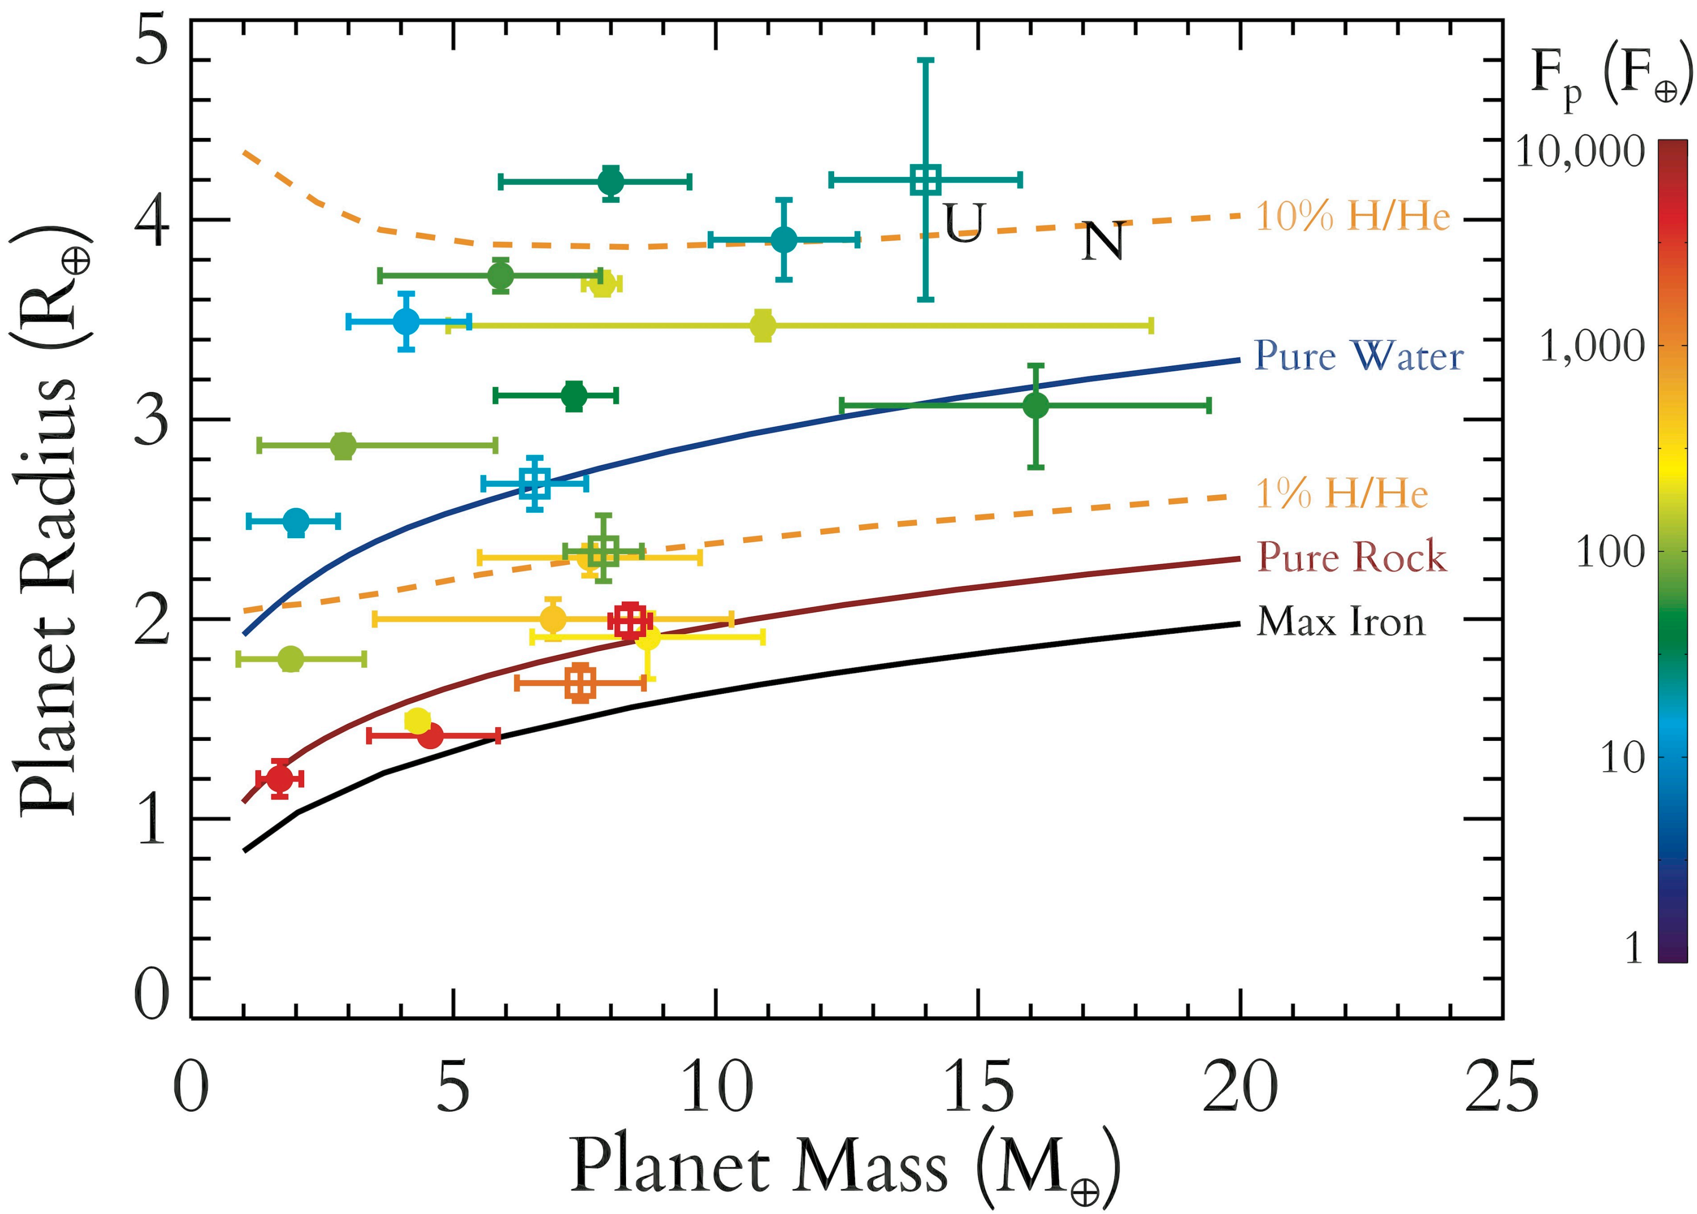

Mass-Radius Relationship

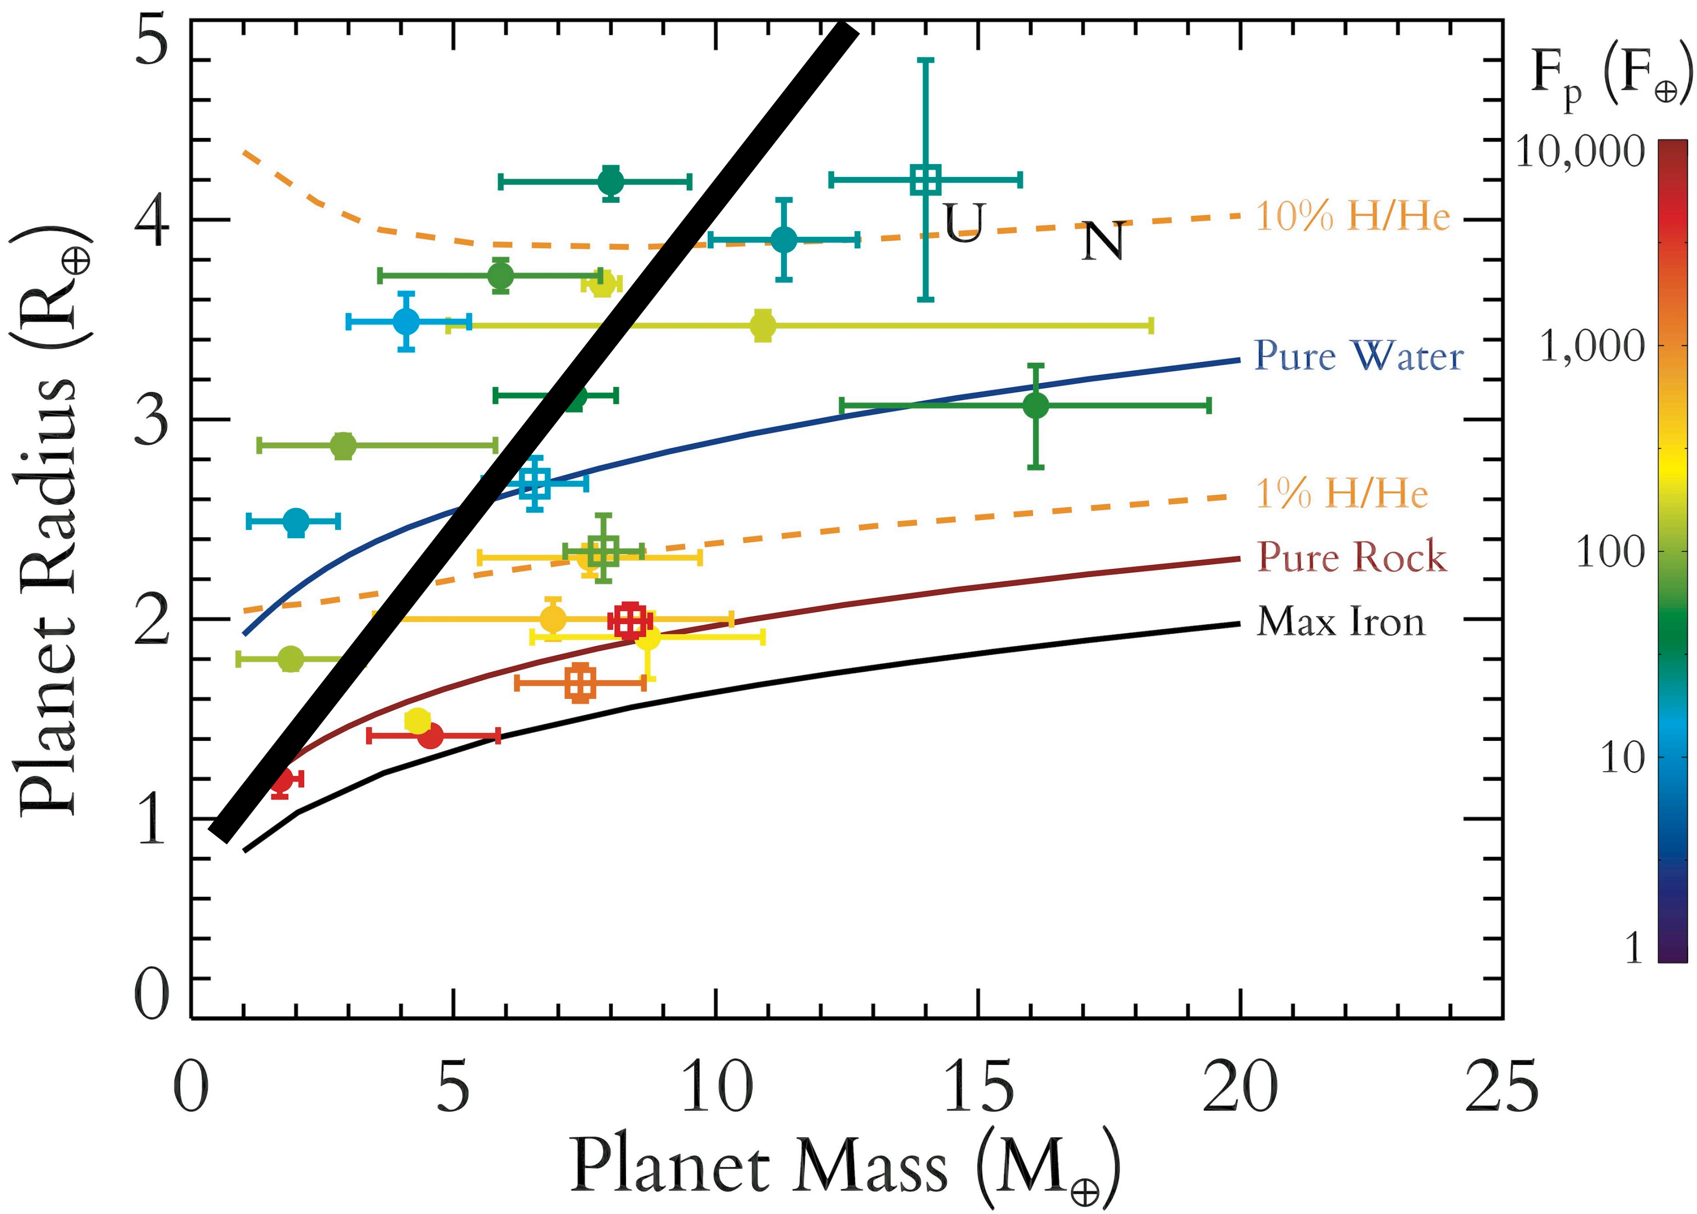

Mass-Radius Relationship

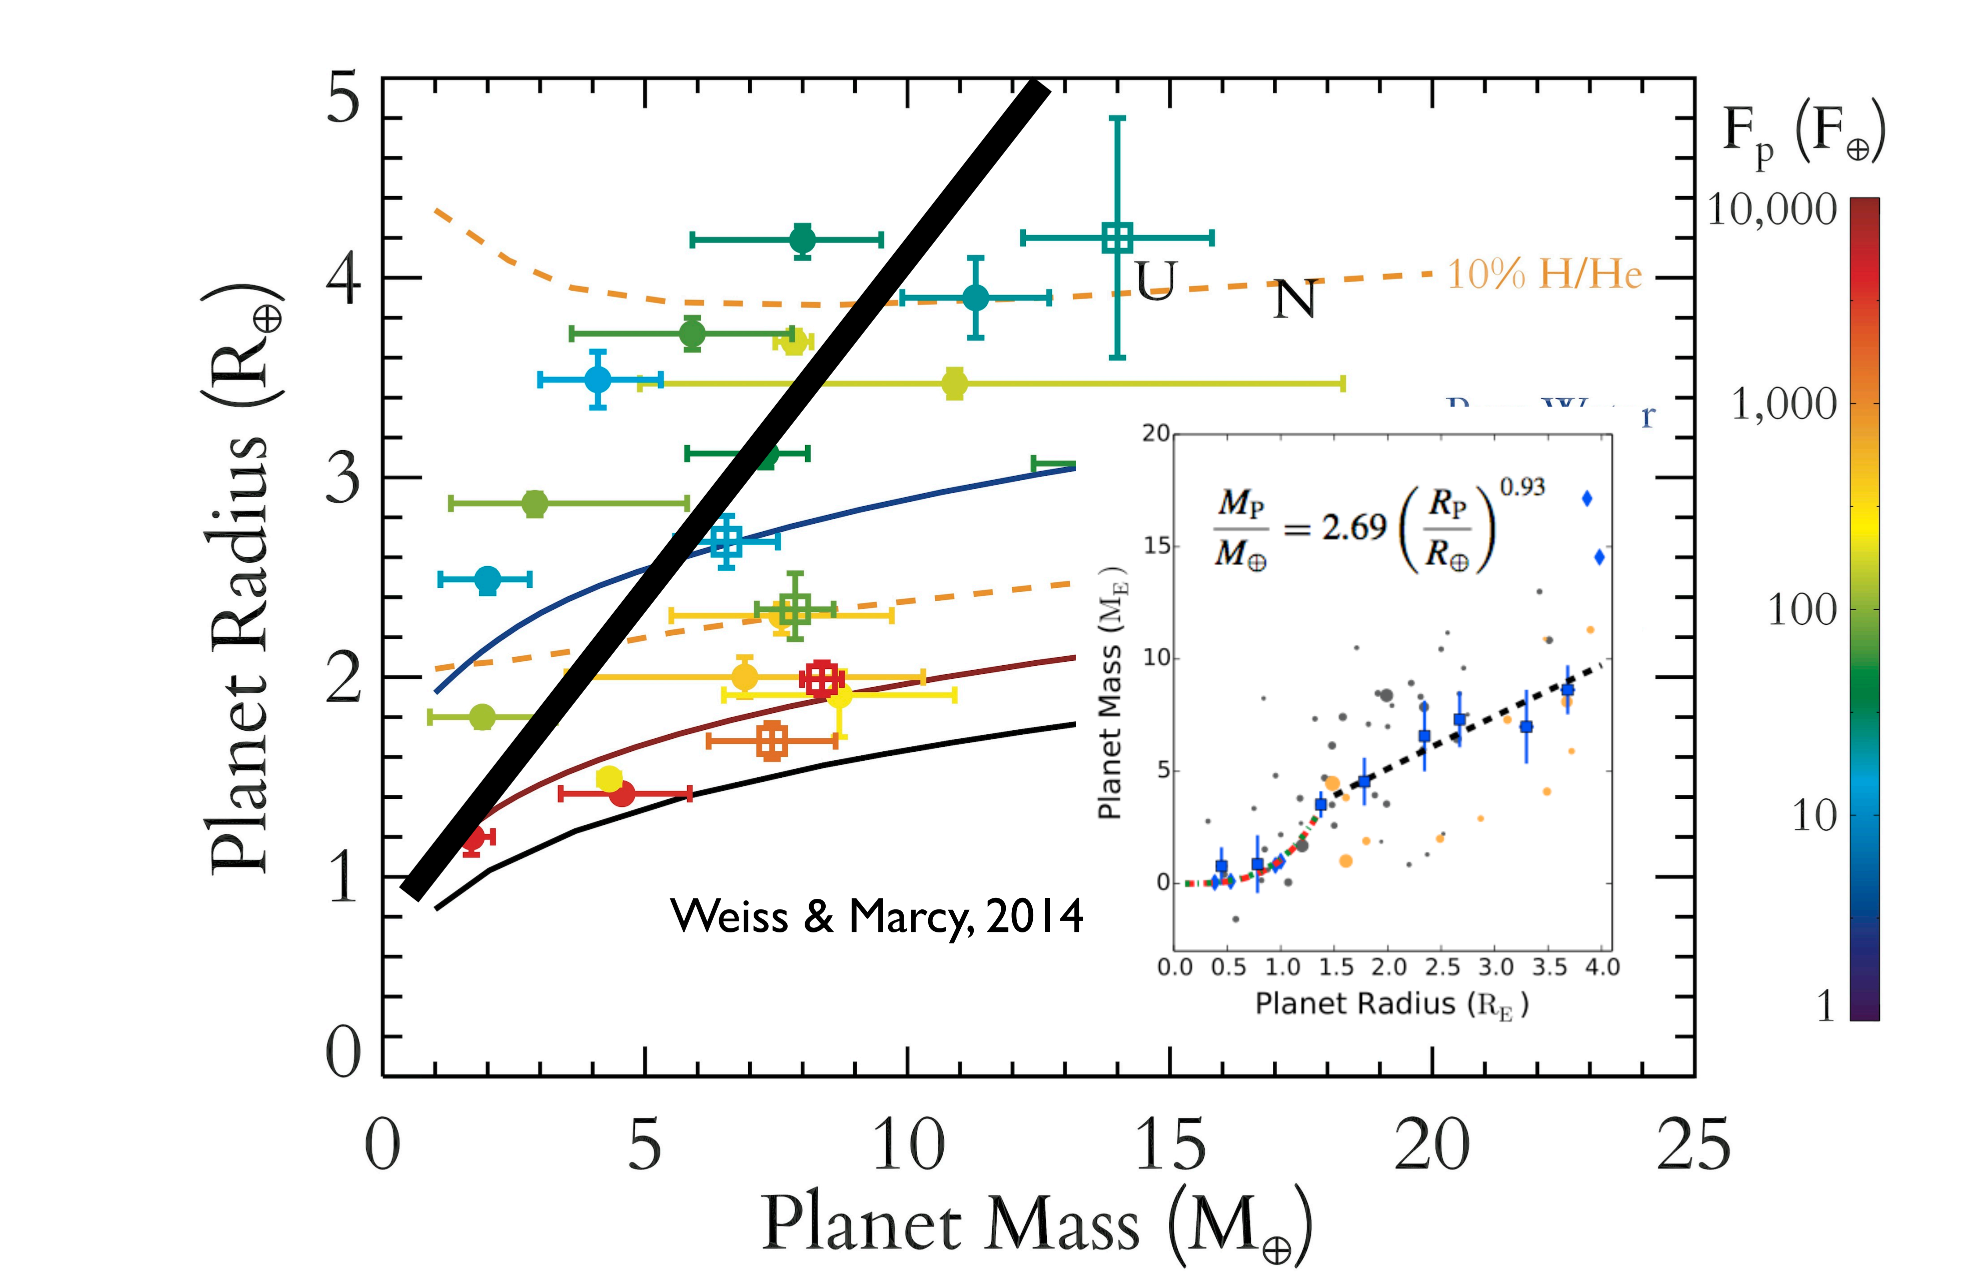

Mass-Radius Relationship

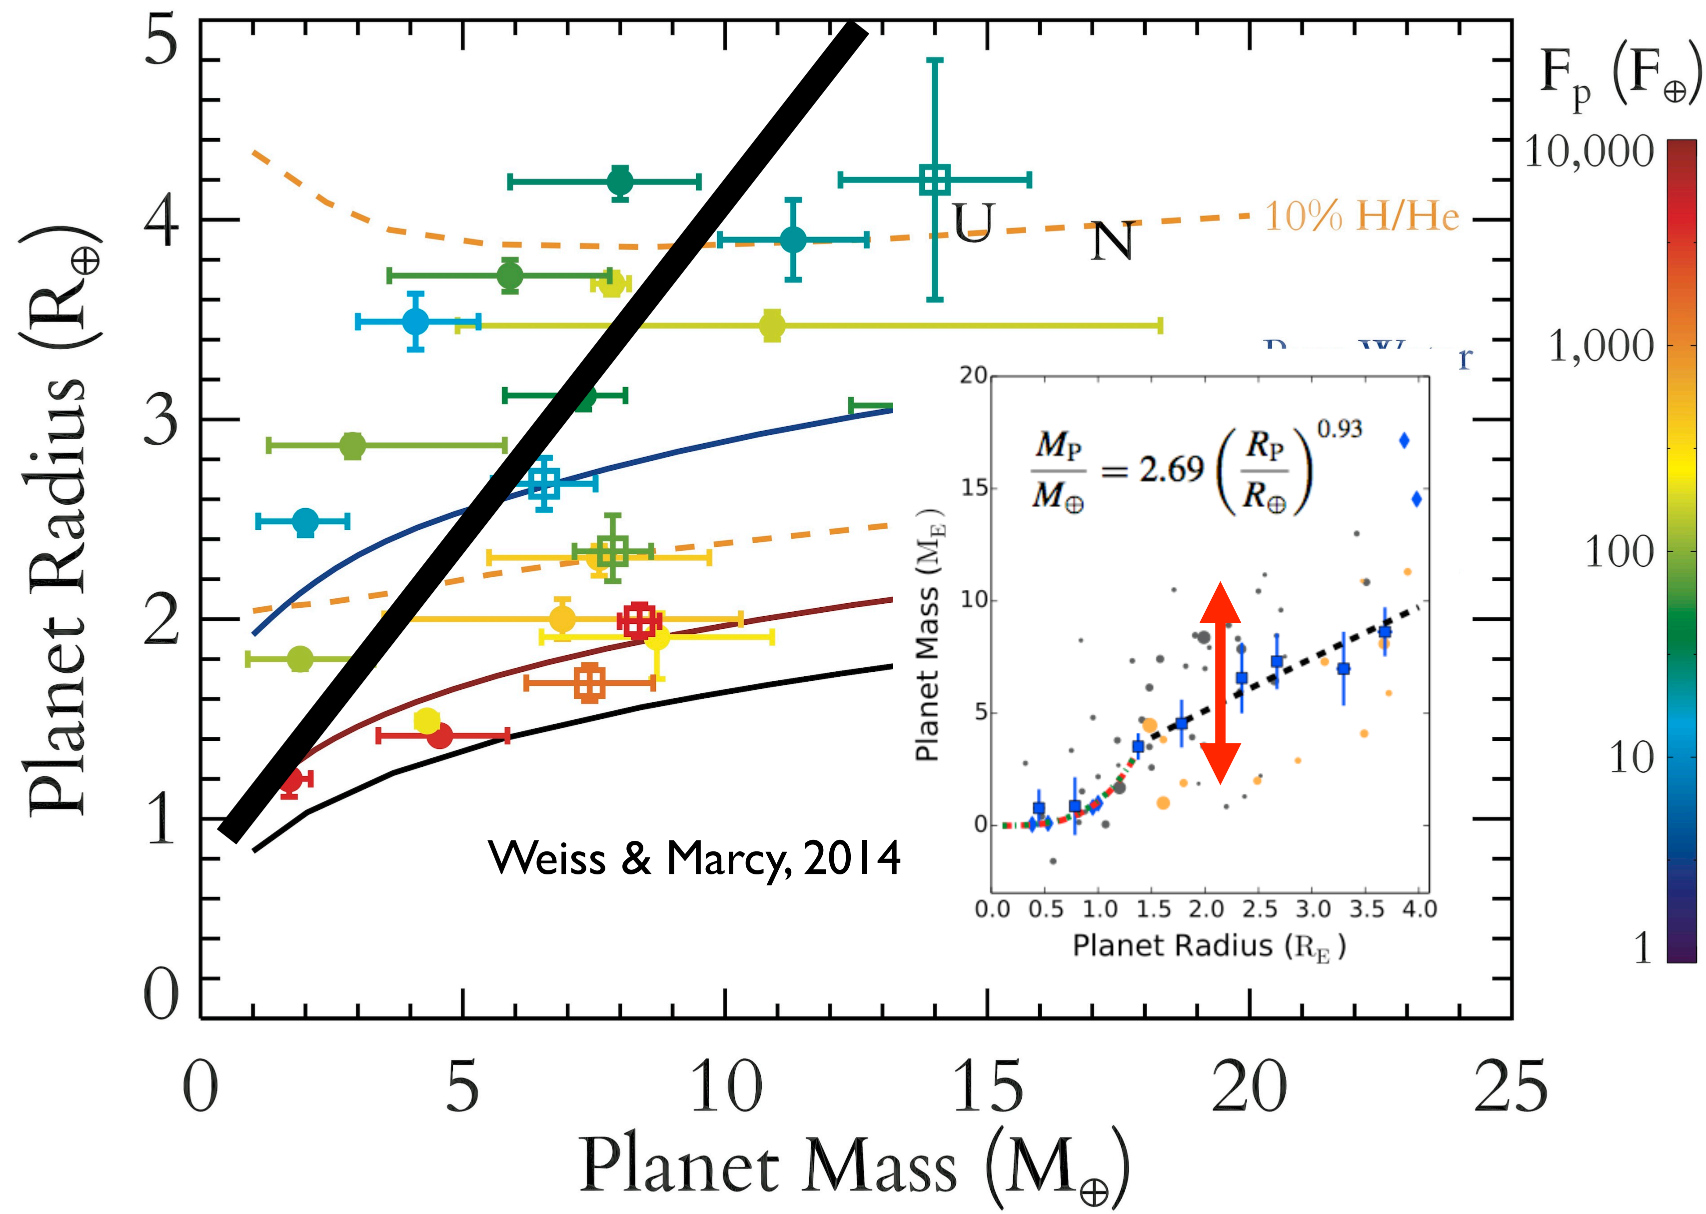

Mass-Radius ~~Relationship~~

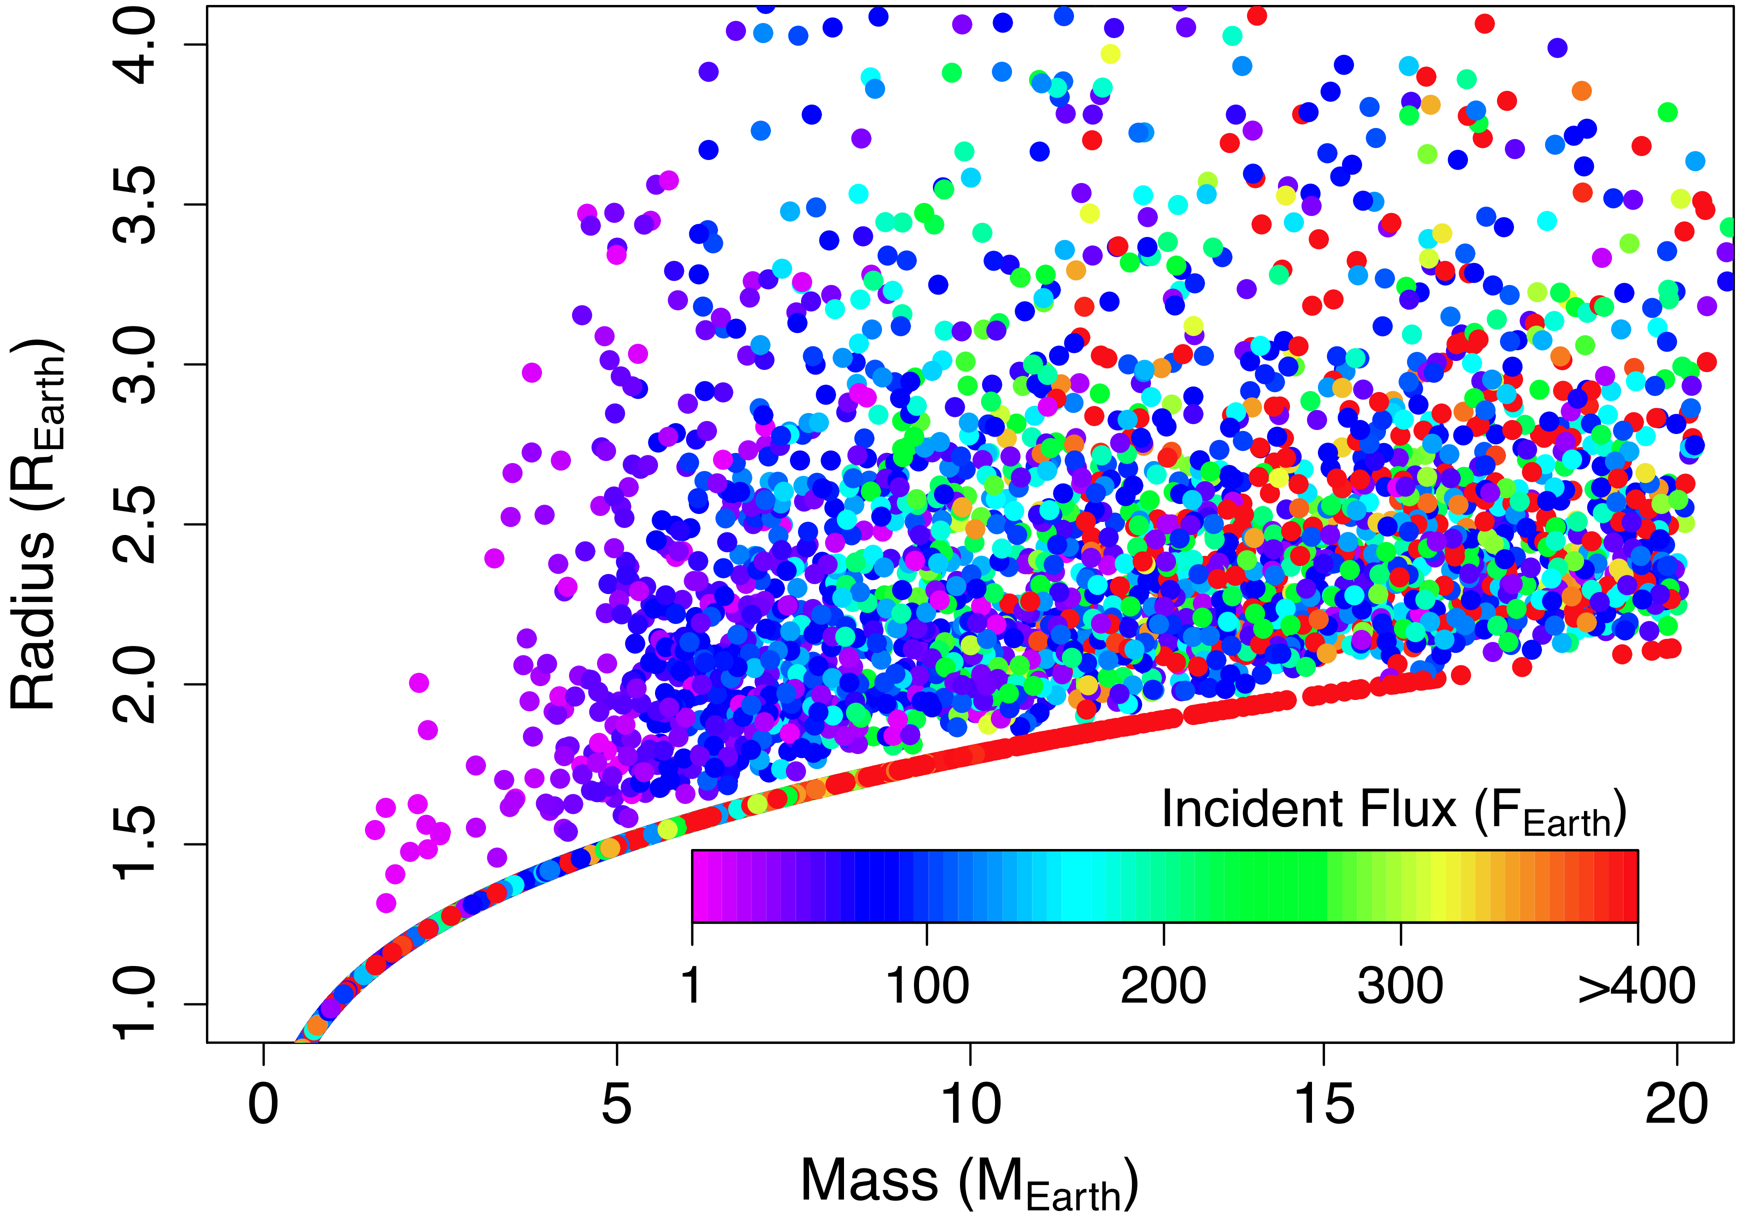

Mass-Radius PDF (probability density function)

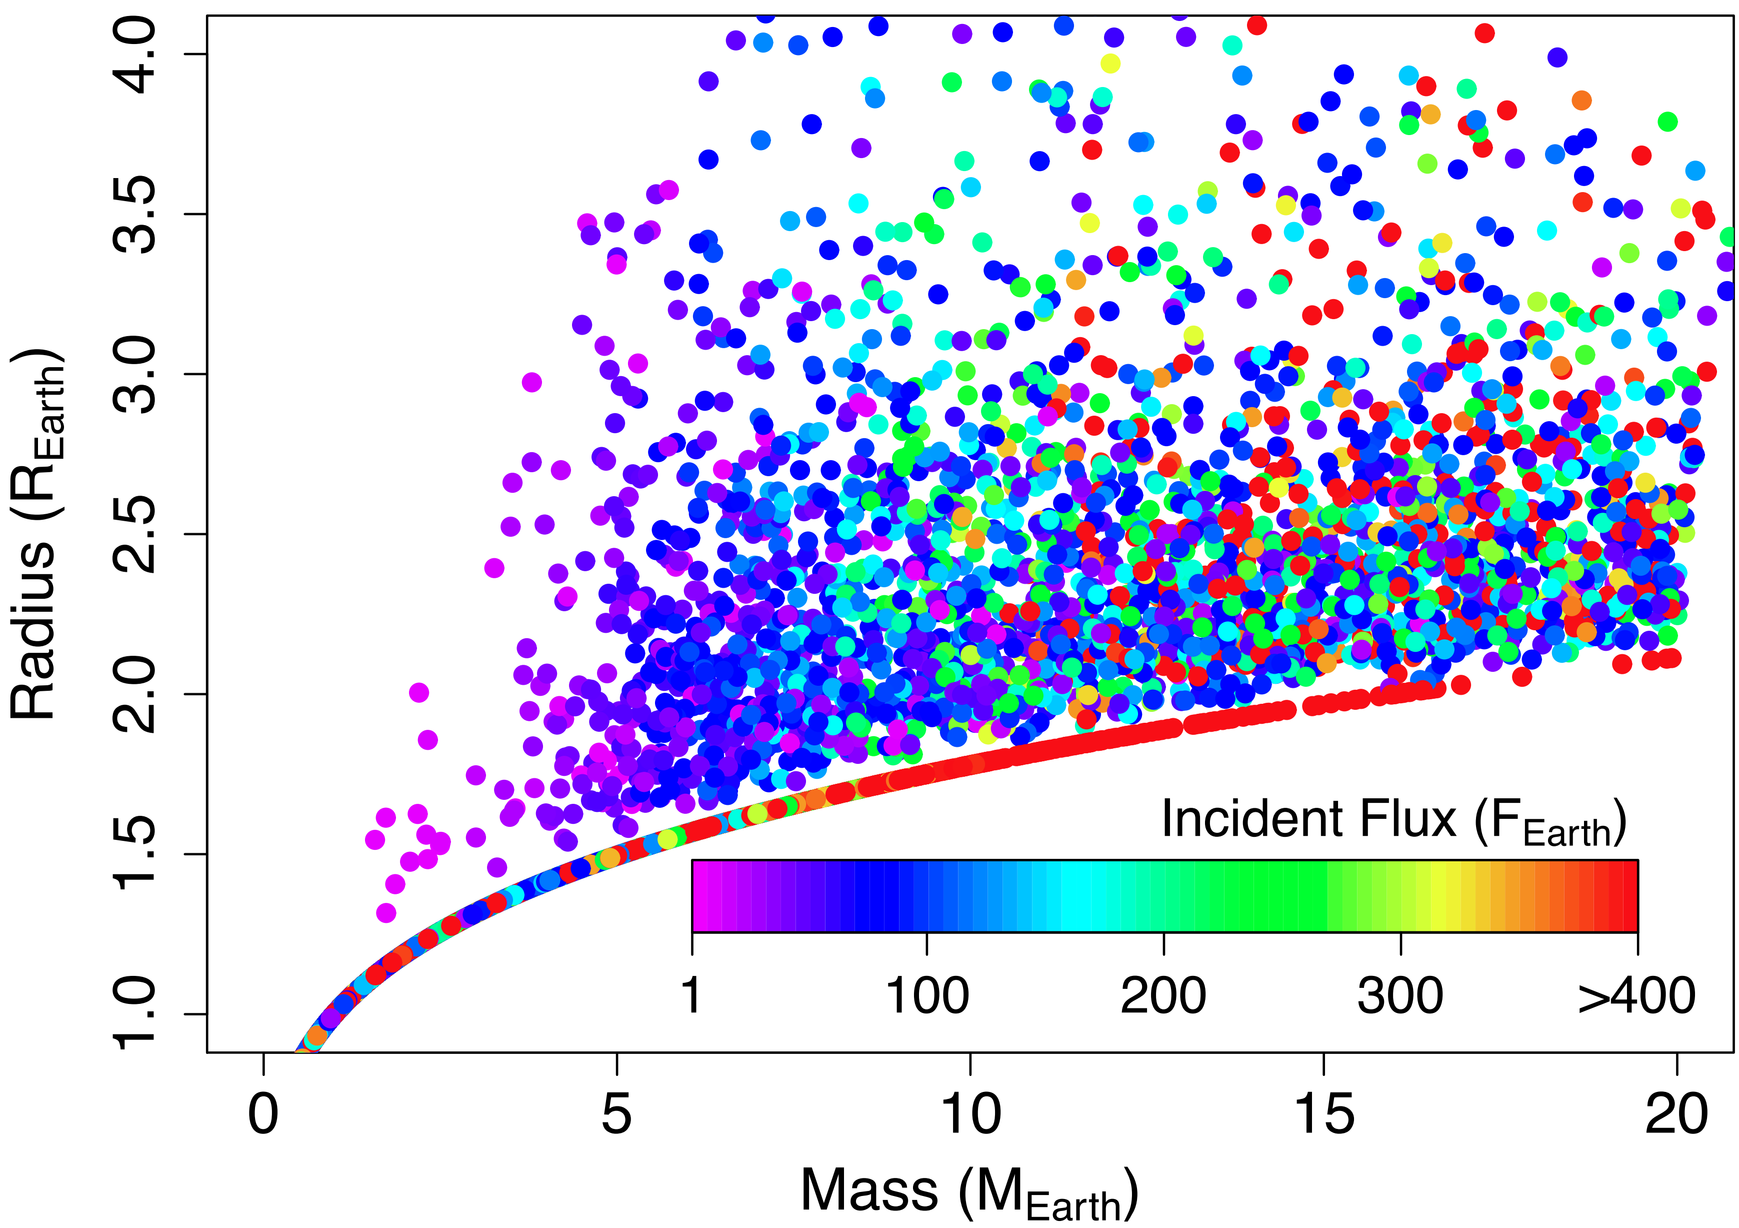

Mass-Radius PDF (probability density function)

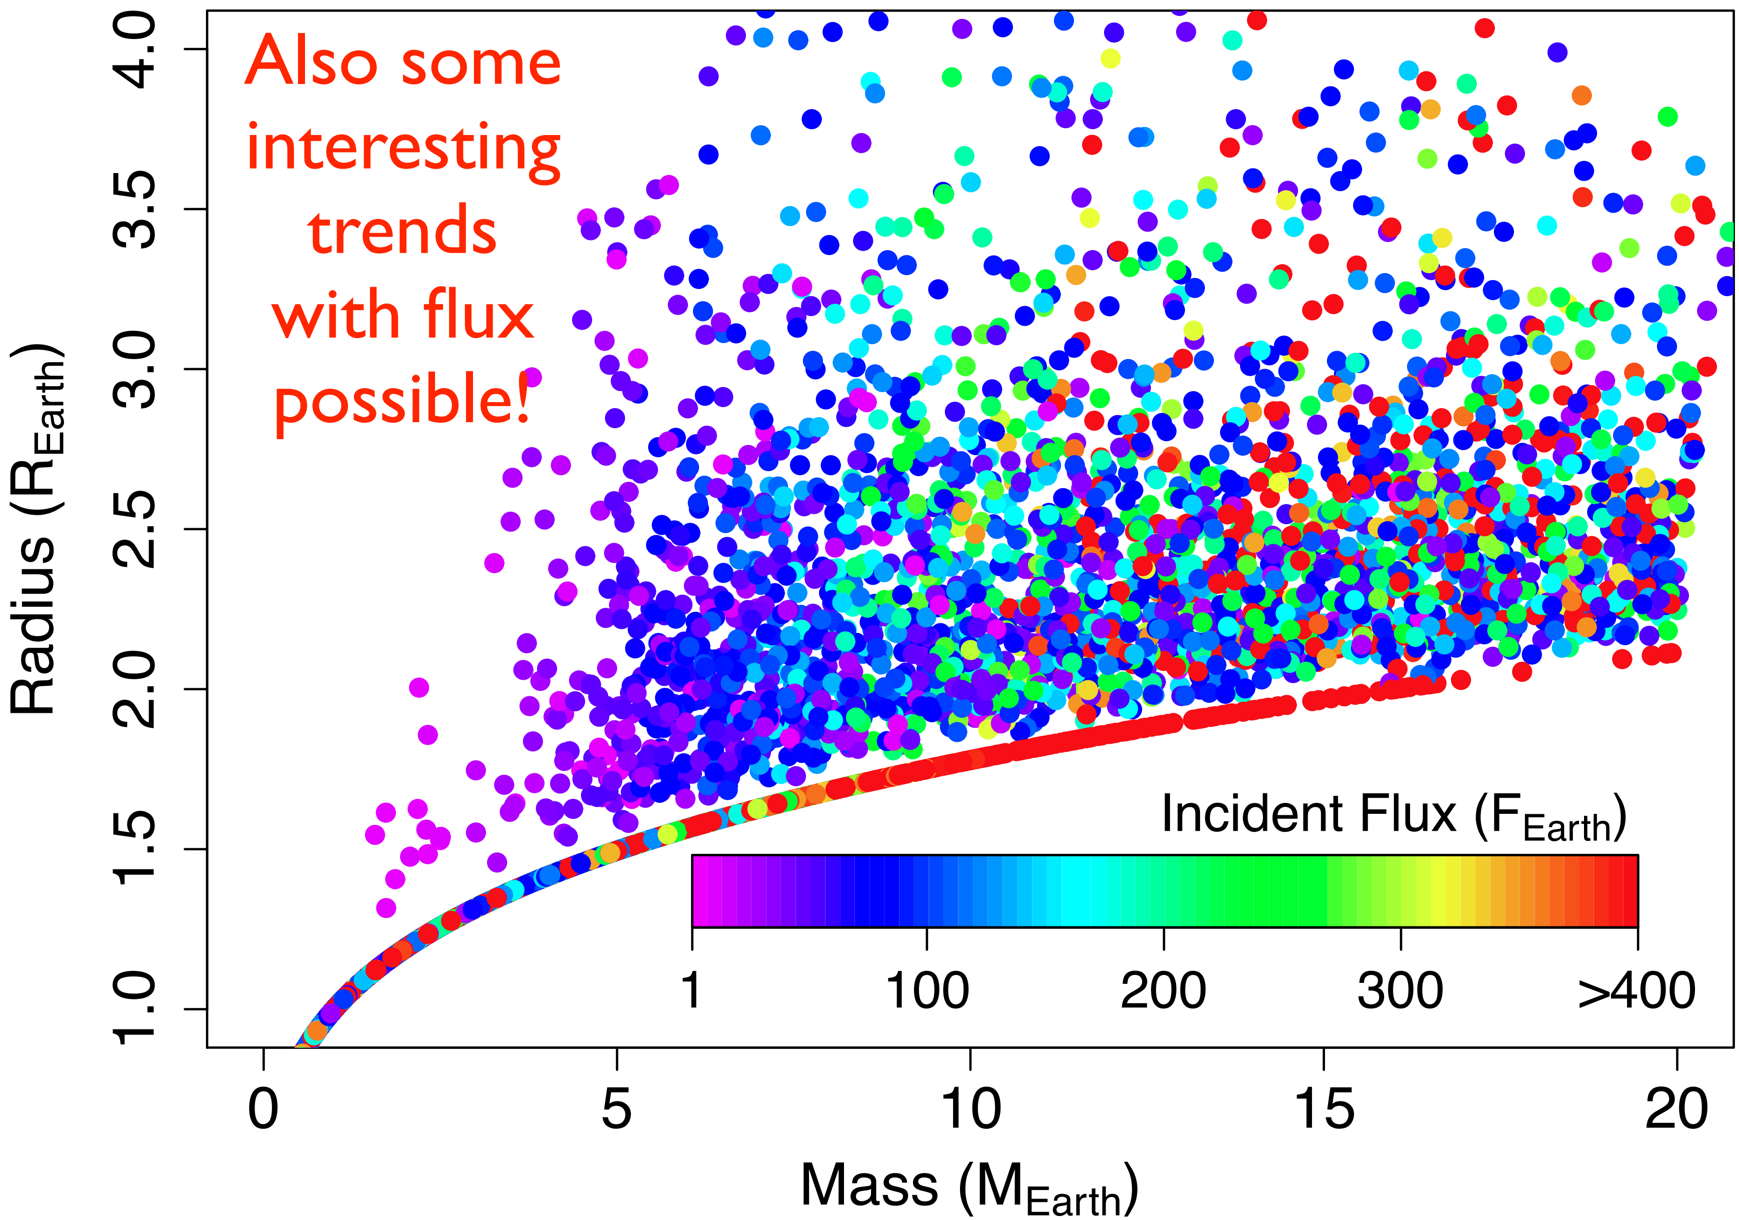

Mass-Radius PDF

$$\frac{M}{M_\oplus} = C\left(\frac{R}{R_\oplus}\right)^\gamma$$

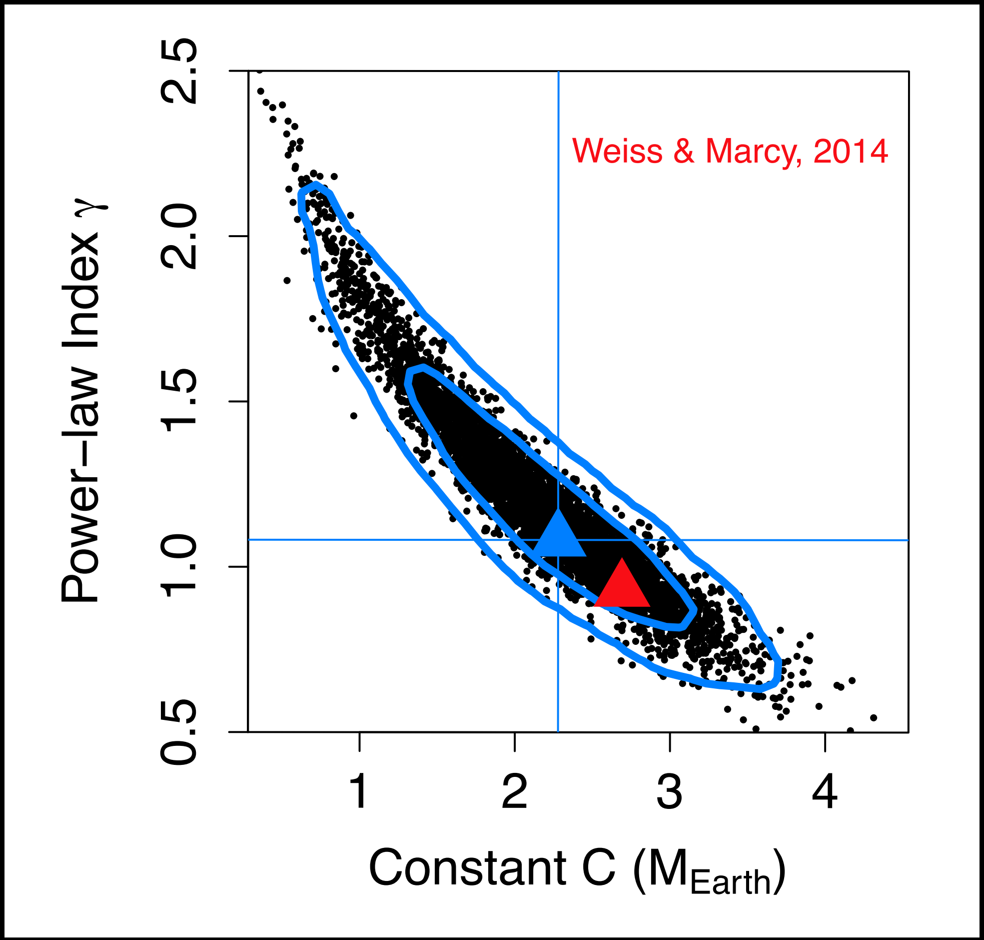

Mass-Radius PDF

$$\frac{M}{M_\oplus} = C\left(\frac{R}{R_\oplus}\right)^\gamma \longrightarrow \frac{M}{M_\oplus} \sim \mathrm{Normal}\left(\mu = C\left(\frac{R}{R_\oplus}\right)^\gamma, \sigma = \sigma_M\right)$$

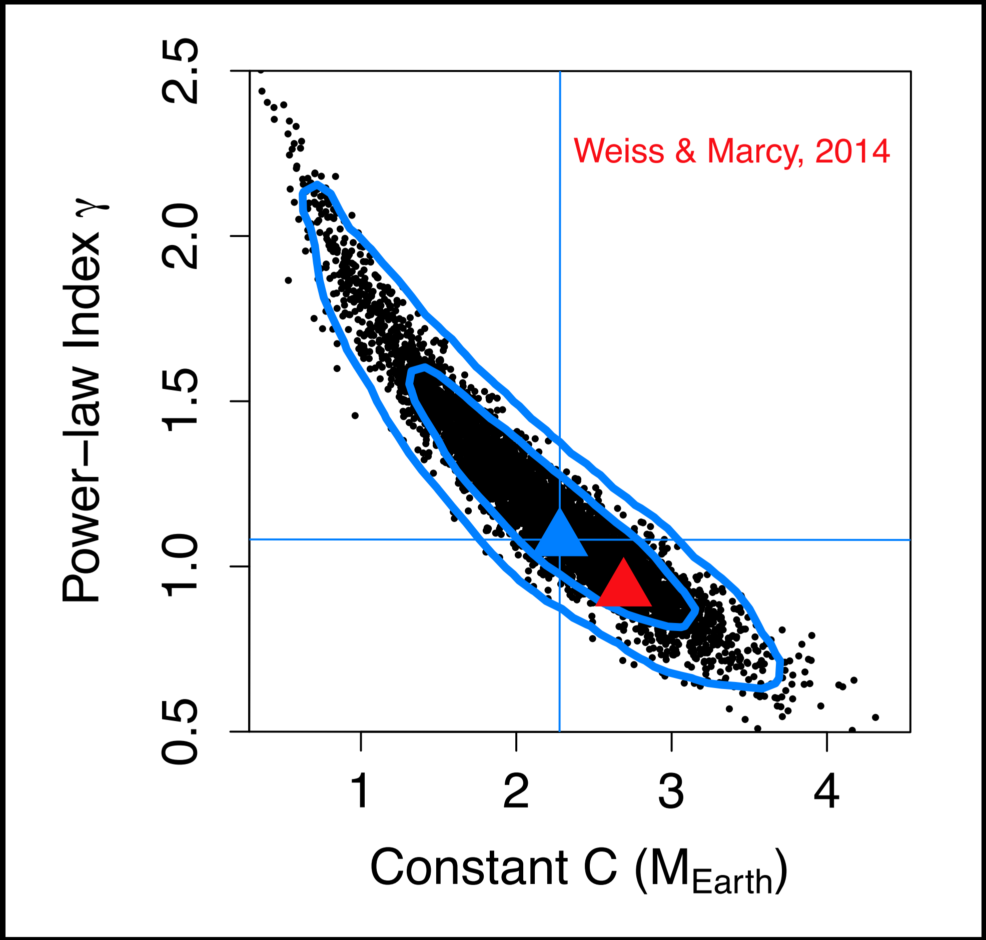

Mass-Radius PDF

$$\frac{M}{M_\oplus} = C\left(\frac{R}{R_\oplus}\right)^\gamma \longrightarrow \frac{M}{M_\oplus} \sim \mathrm{Normal}\left(\mu = C\left(\frac{R}{R_\oplus}\right)^\gamma, \sigma = \sigma_M\right)$$

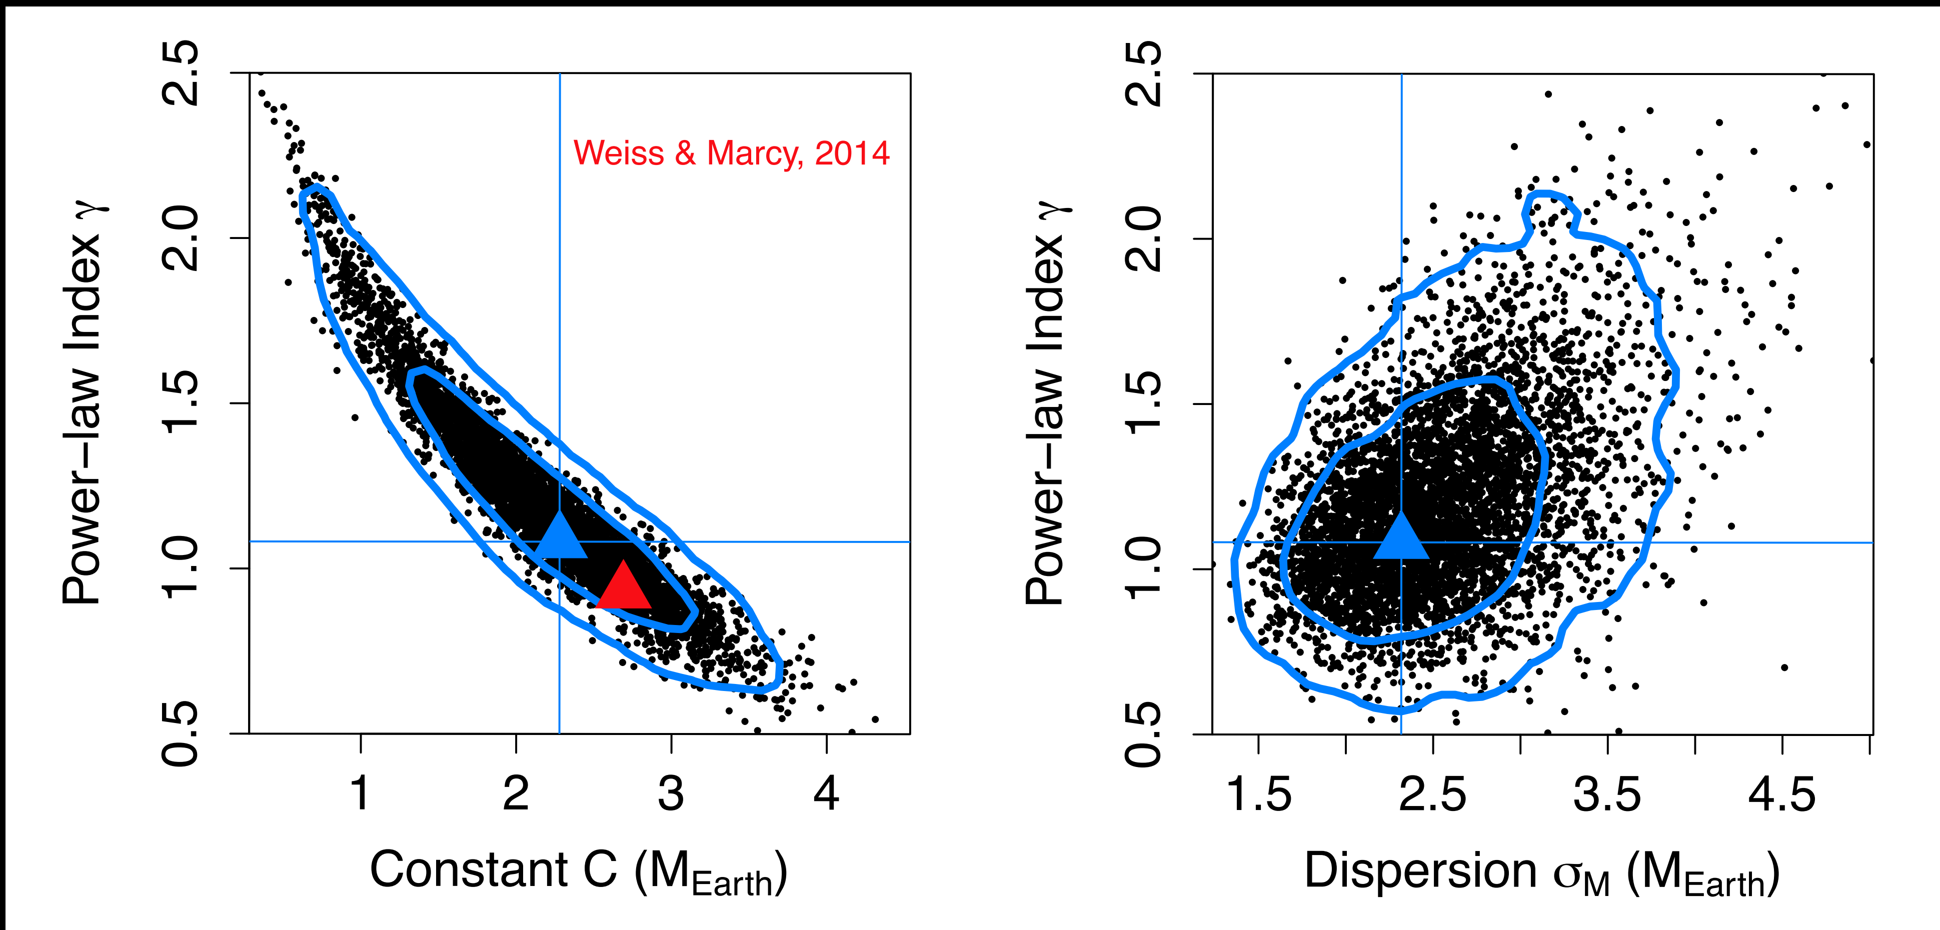

Mass-Radius PDF

$$\frac{M}{M_\oplus} \sim \text{Normal}\left(\mu = 2.28\left(\frac{R}{R_\oplus}\right)^{1.05}, \sigma = 2.32\right)$$

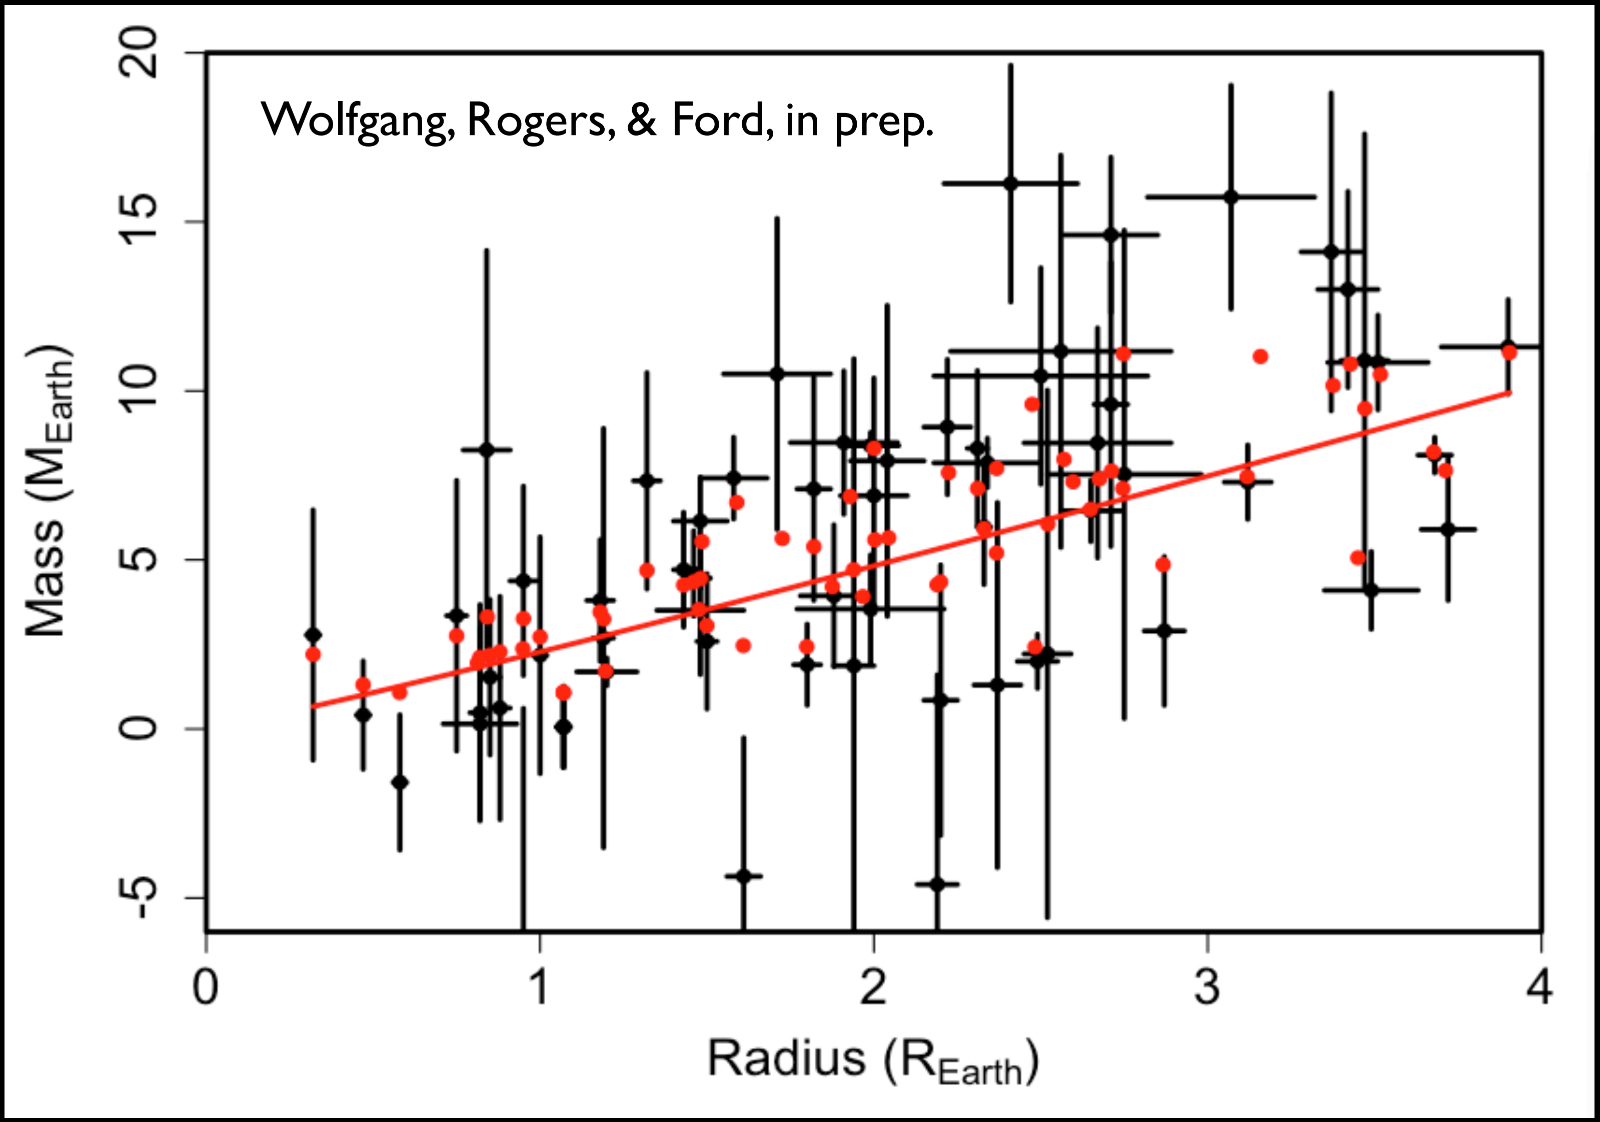

Mass-Radius PDF

$$\frac{M}{M_\oplus} \sim \mathrm{Normal}\left(\mu = 2.28\left(\frac{R}{R_\oplus}\right)^{1.05}, \sigma = 2.32\right)$$

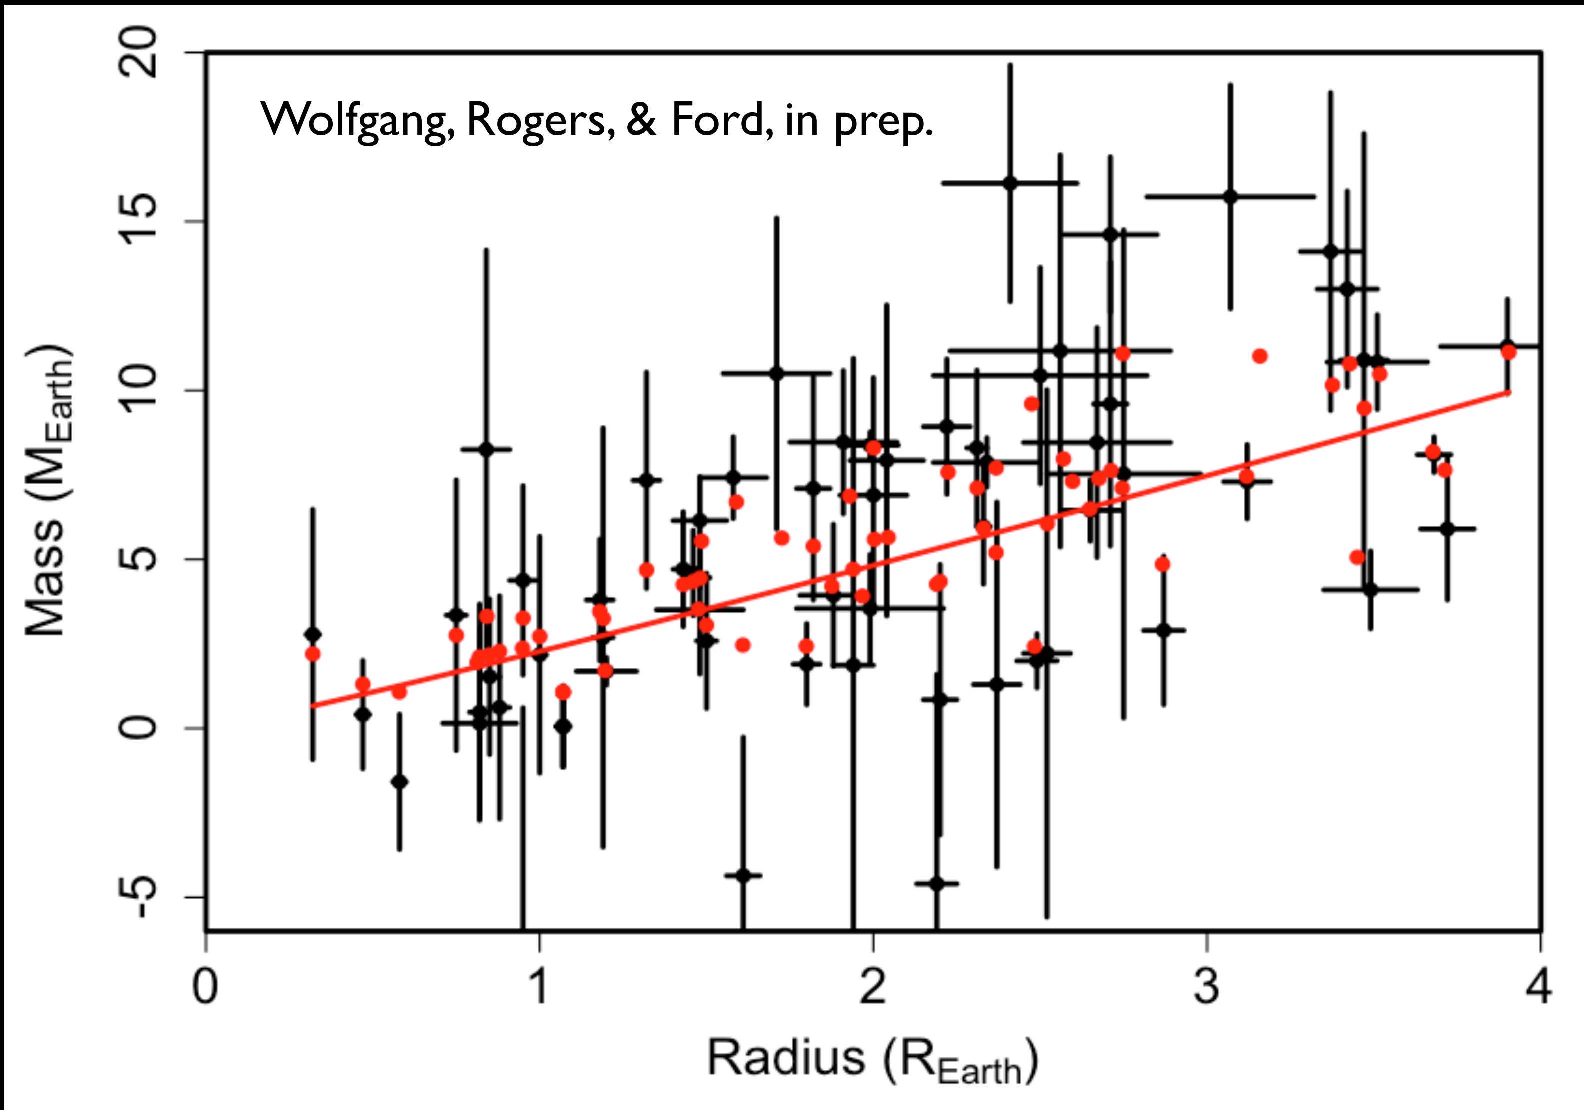

For dynamical studies, can now accurately represent how much we know about a planet's mass based on radius!

Summary of Implications

Summary of Implications

1) First composition *distribution:*

~1% envelope mass fractions are the most likely compositions for sub-Neptunes.

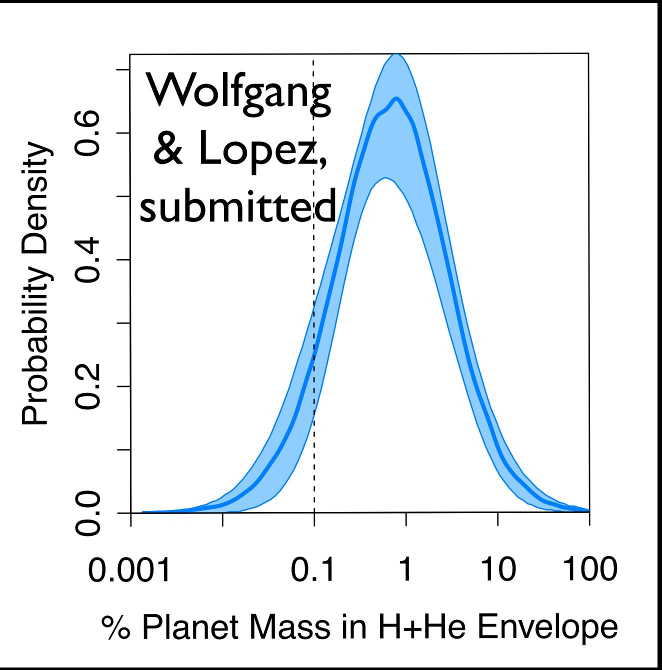

Summary of Implications

1) First composition *distribution:*

~1% envelope mass fractions are the most likely compositions for sub-Neptunes.

2) With radius uncertainties:

$R_{pl} < 2 \rightarrow f_{env} < 1\%$

$2 < R_{pl} < 3 \rightarrow f_{env} \sim 1\%$

$3 < R_{pl} < 4 \rightarrow f_{env} \sim$ a few %

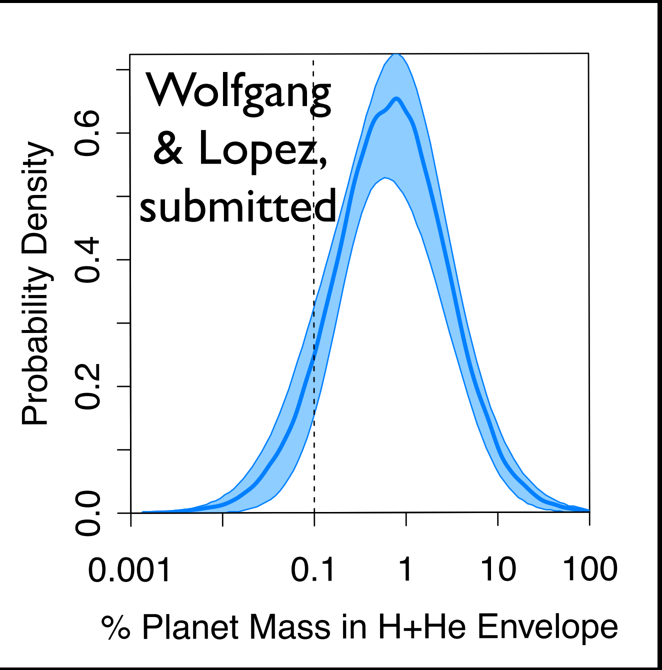

Summary of Implications

1) First composition *distribution:*

~1% envelope mass fractions are the most likely compositions for sub-Neptunes.

2) With radius uncertainties:

$R_{pl} < 2 \rightarrow f_{env} < 1\%$

$2 < R_{pl} < 3 \rightarrow f_{env} \sim 1\%$

$3 < R_{pl} < 4 \rightarrow f_{env} \sim$ a few %

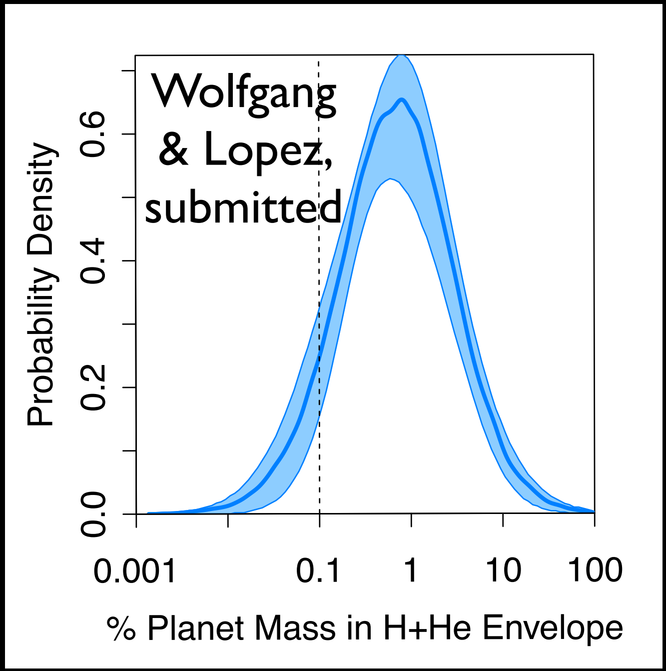

3) Rocky-gas transition around ~ 1.5 R_{Earth} (agrees with Rogers et. al. 2014), and could be fuzzy

Summary of Implications

1) First composition *distribution:*

~1% envelope mass fractions are the most likely compositions for sub-Neptunes.

2) With radius uncertainties:

$R_{pl} < 2 \rightarrow f_{env} < 1\%$

$2 < R_{pl} < 3 \rightarrow f_{env} \sim 1\%$

$3 < R_{pl} < 4 \rightarrow f_{env} \sim$ a few %

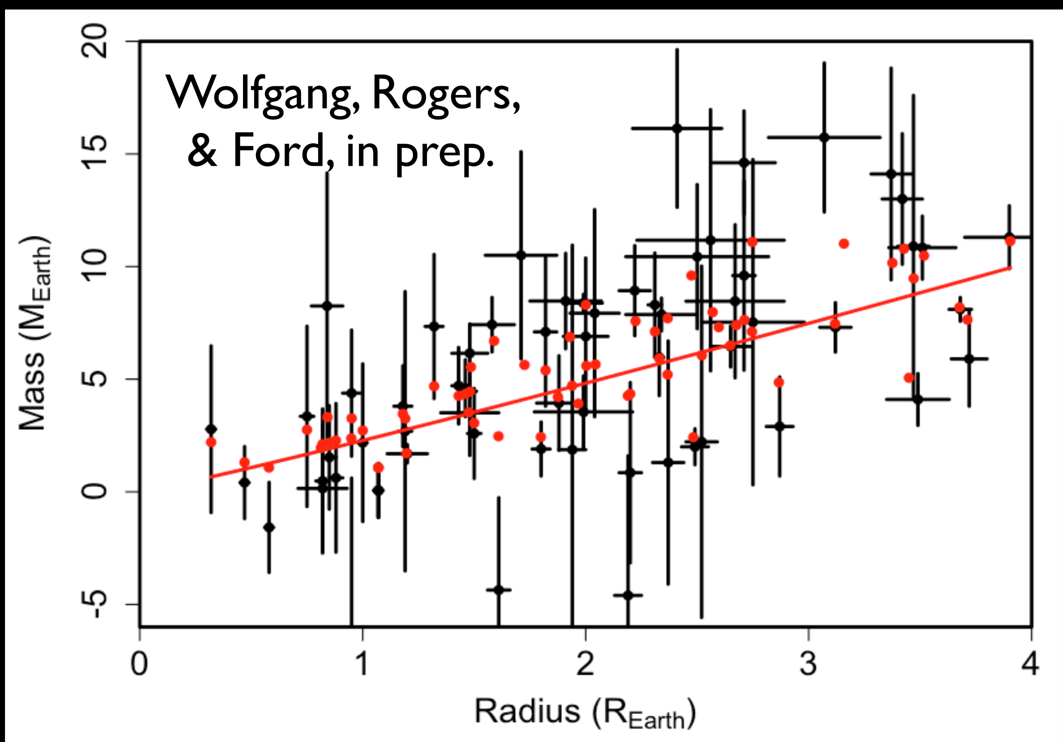

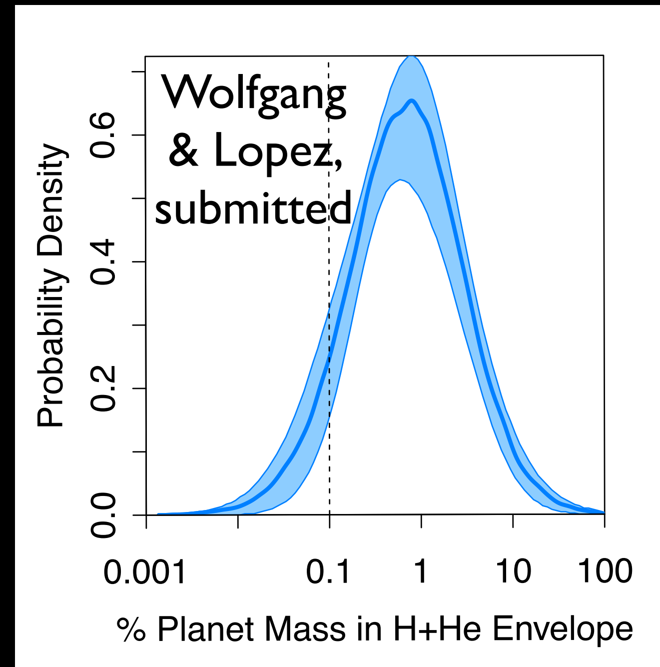

3) Rocky-gas transition around ~ 1.5 R_{Earth} (agrees with Rogers et. al. 2014), and could be fuzzy

4) Need probabilistic treatment to "convert" radii into masses.

Backup slides

Interpreting (M,R)

Interpreting (M,R)

Inferring a composition requires modeling these planets' internal structures:

Interpreting (M,R)

Inferring a composition requires modeling these planets' internal structures:

Continuity:

$$\frac{dm(r)}{dr} = 4\pi r^2 \rho(r)$$

Hydrostatic Equilibrium:

$$\frac{dP(r)}{dr} = \frac{-Gm(r)\rho(r)}{r^2}$$

Equation of State:

$$P(r) = f(\rho(r),\ T(r))$$

Interpreting (M,R)

Inferring a composition requires modeling these planets' internal structures:

Continuity:

$$\frac{dm(r)}{dr} = 4\pi r^2 \rho(r)$$

Hydrostatic Equilibrium:

$$\frac{dP(r)}{dr} = \frac{-Gm(r)\rho(r)}{r^2}$$

Equation of State:

$$P(r) = f(\rho(r), T(r))$$

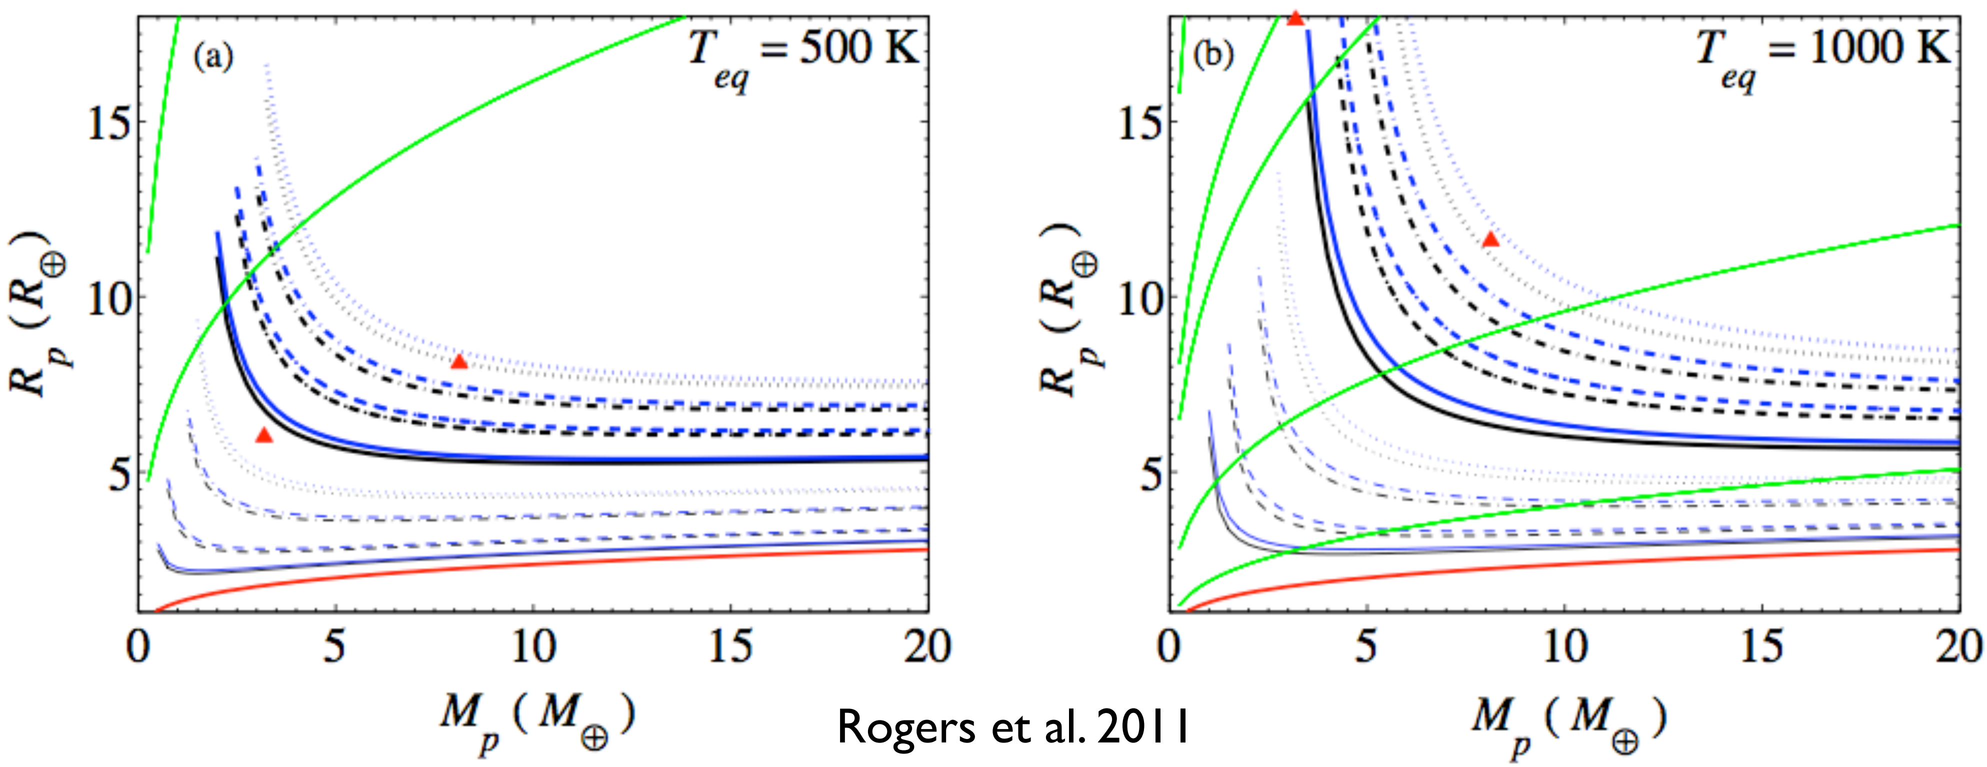

Rogers et al. 2011

Interpreting (M,R)

Inferring a composition requires modeling these planets’ internal structures:

Continuity:

$$\frac{dm(r)}{dr} = 4\pi r^2 \rho(r)$$

Hydrostatic Equilibrium:

$$\frac{dP(r)}{dr} = \frac{-Gm(r)\rho(r)}{r^2}$$

Equation of State:

$$P(r) = f(\rho(r),\, T(r))$$

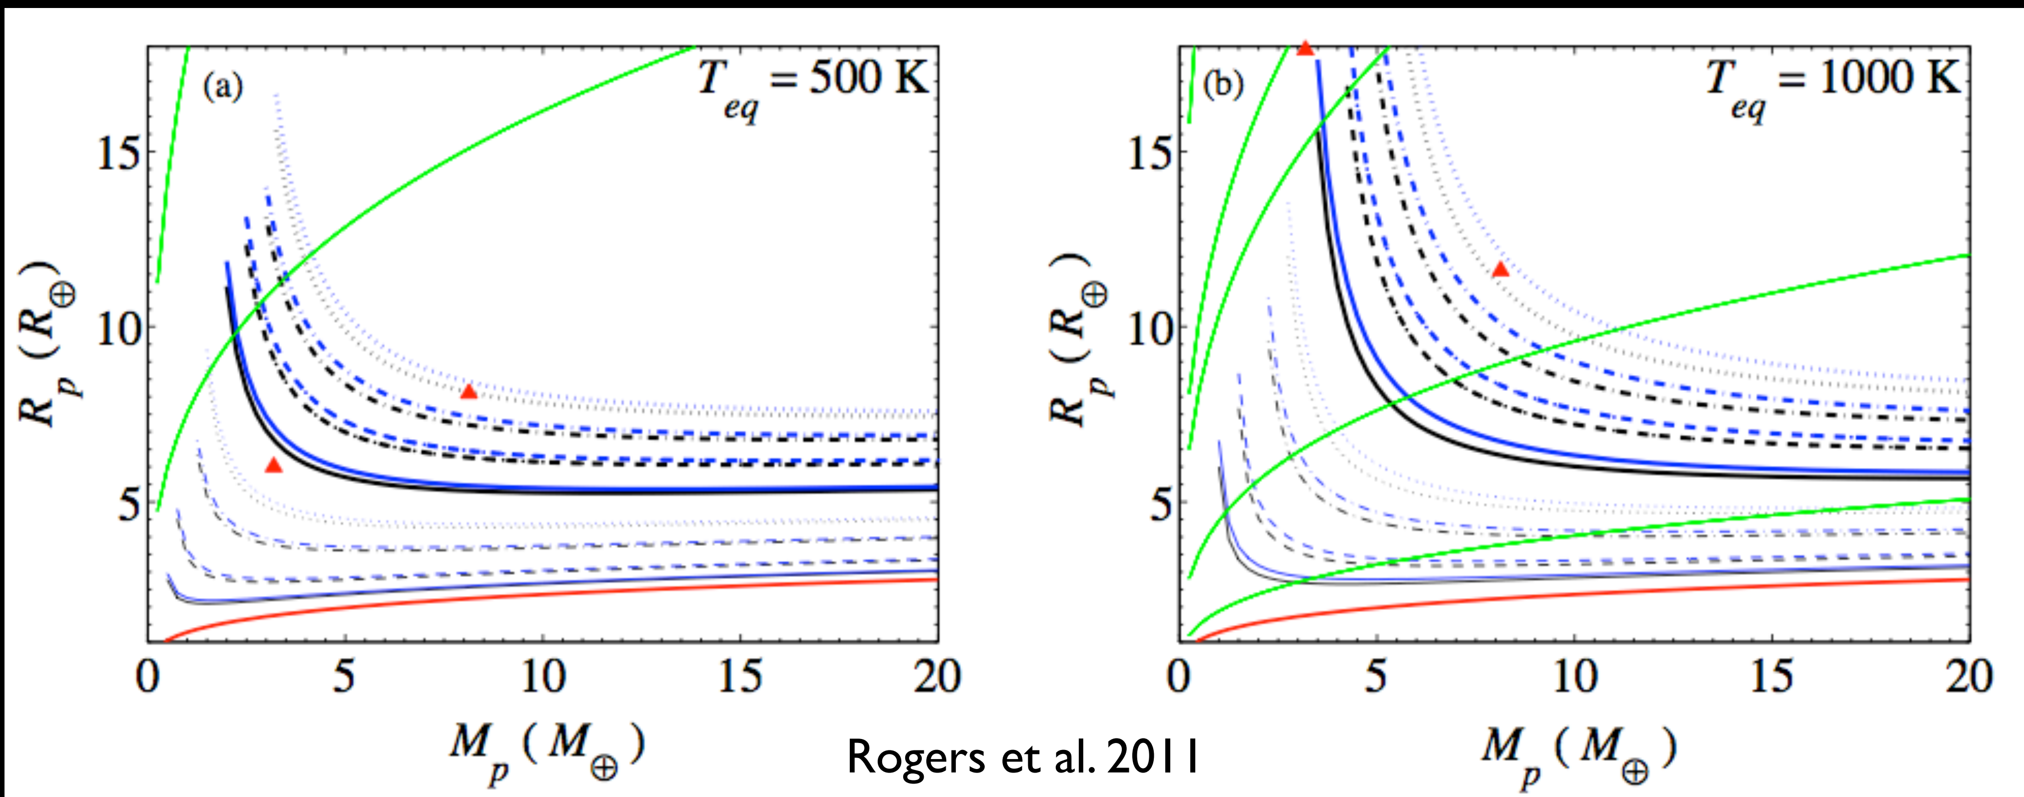

Rogers et al. 2011

Intrinsic luminosity is set to a constant value on a grid

Interpreting (M,R)

Inferring a composition requires modeling these planets' internal structures:

Continuity:

$$\frac{dm(r)}{dr} = 4\pi r^2 \rho(r)$$

Hydrostatic Equilibrium:

$$\frac{dP(r)}{dr} = \frac{-Gm(r)\rho(r)}{r^2}$$

Equation of State:

$$P(r) = f(\rho(r), T(r))$$

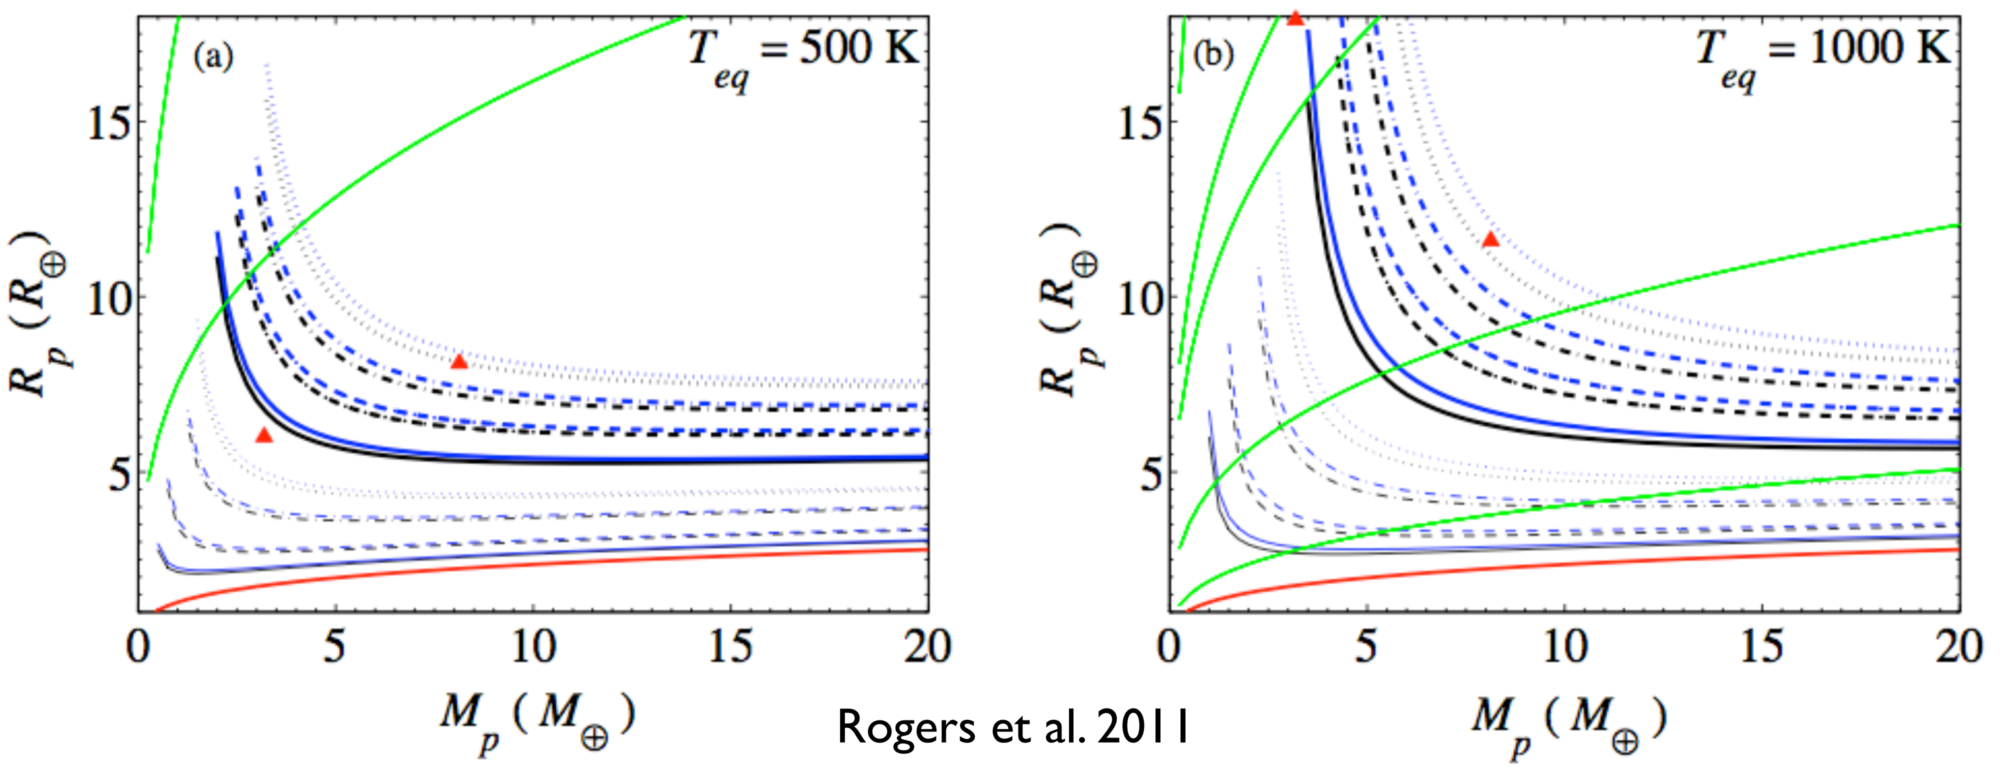

Rogers et al. 2011

Intrinsic luminosity is set to a constant value on a grid

. . . but how to determine an astrophysically appropriate value?

But a planet cools . . .

$$\int_{M_{\rm core}}^{M_{\rm p}} dm \frac{T\,dS}{dt} = -L_{\rm int} + L_{\rm radio} - c_{\rm v} M_{\rm core} \frac{dT_{\rm core}}{dt}$$

But a planet cools . . .

$$\int_{M_{\rm core}}^{M_{\rm p}} dm \frac{T\,dS}{dt} = -L_{\rm int} + L_{\rm radio} - c_{\rm v} M_{\rm core} \frac{dT_{\rm core}}{dt}$$

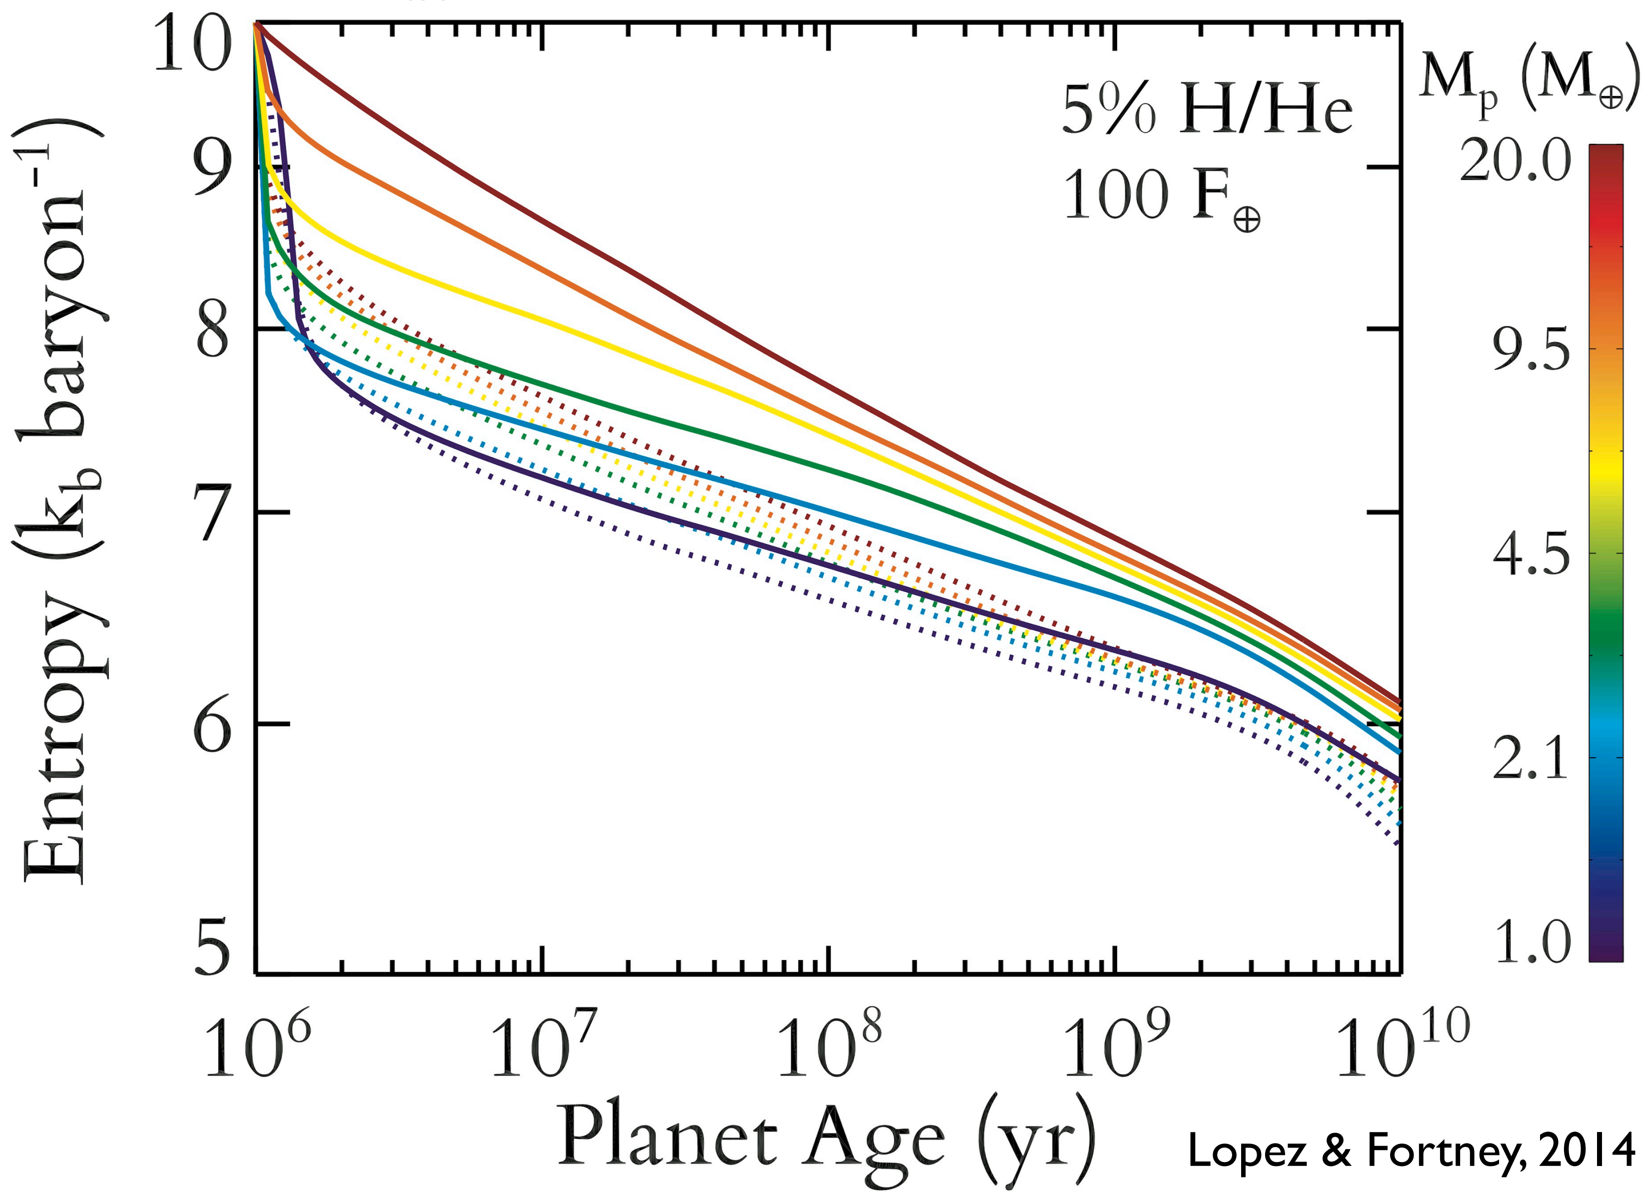

Lopez & Fortney, 2014

But a planet cools . . .

$$\int_{M_{\rm core}}^{M_{\rm p}} dm \frac{T\,dS}{dt} = -L_{\rm int} + L_{\rm radio} - c_{\rm v} M_{\rm core} \frac{dT_{\rm core}}{dt}$$

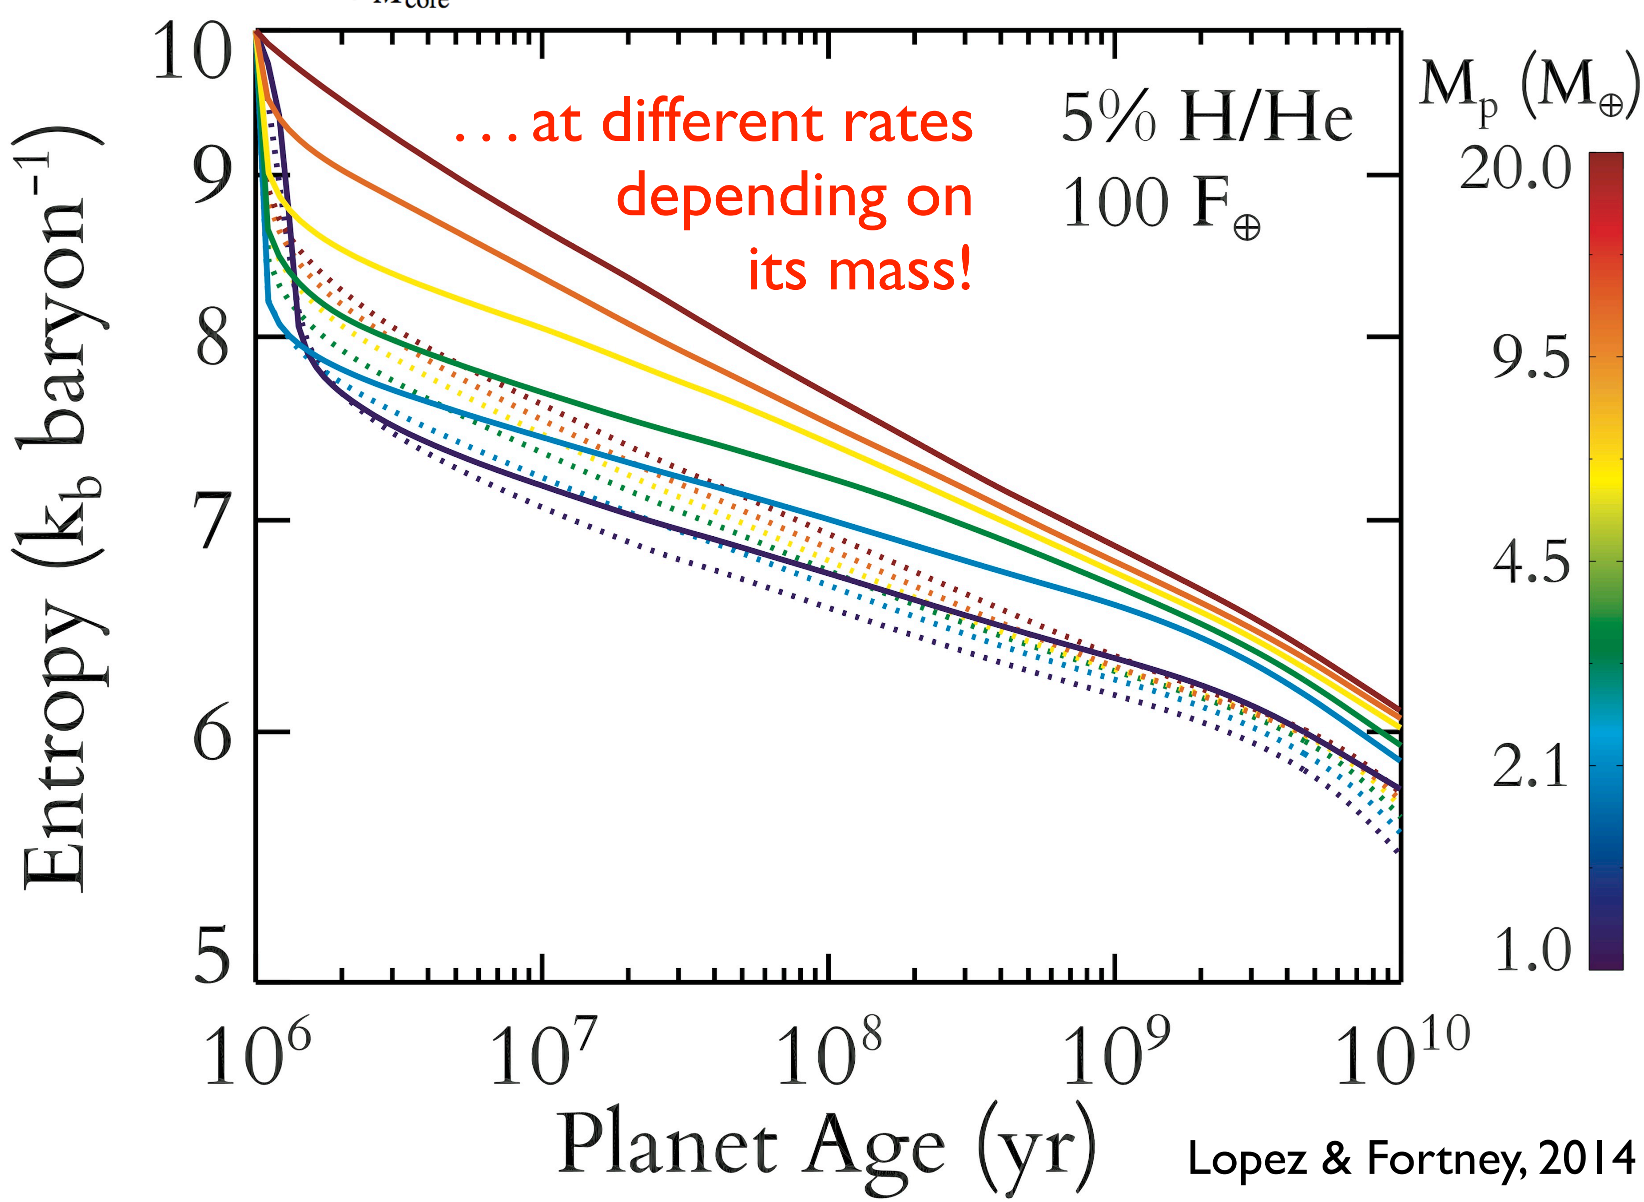

Lopez & Fortney, 2014

But a planet cools . . .

$$\int_{M_{\rm core}}^{M_{\rm p}} dm \frac{T\,dS}{dt} = -L_{\rm int} + L_{\rm radio} - c_{\rm v} M_{\rm core} \frac{dT_{\rm core}}{dt}$$

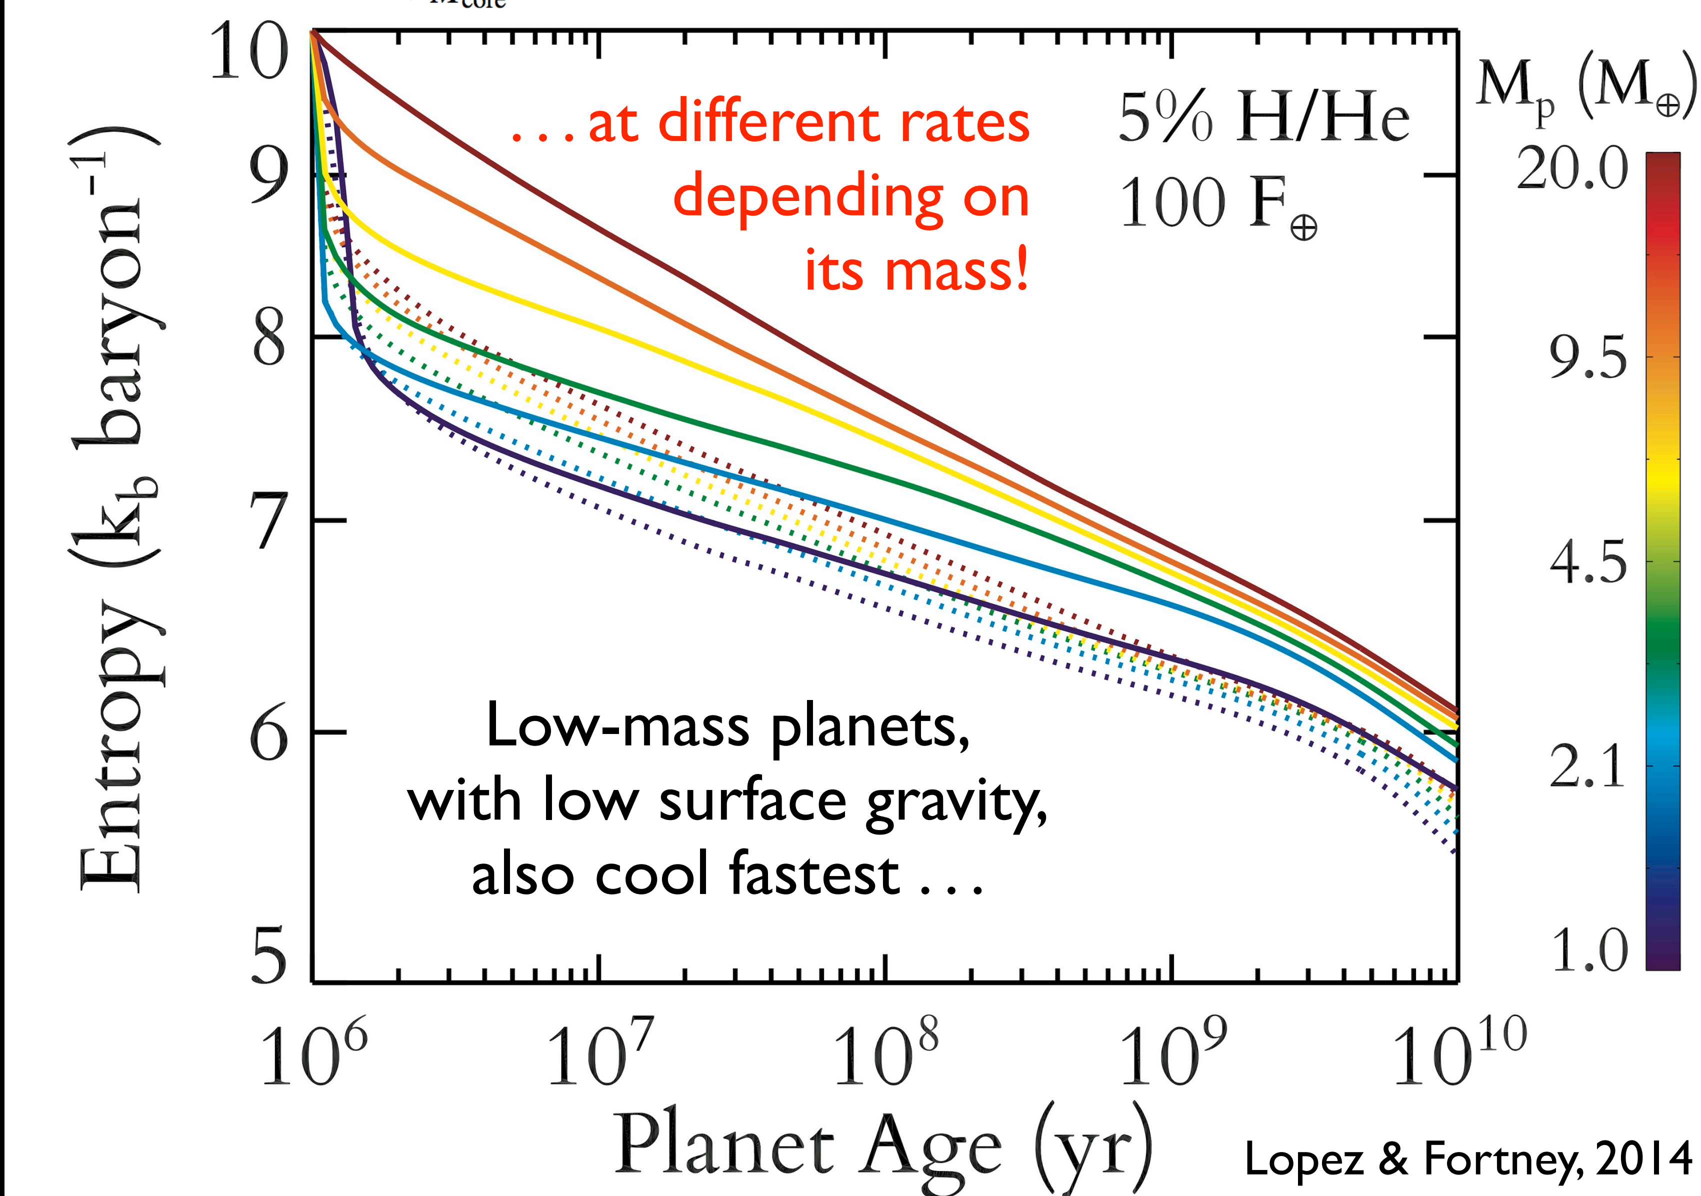

Why is this exciting?

Why is this exciting?

Can sidestep the need for expensive mass measurements!

(but need to assume rocky core with H+He envelope)

Why is this exciting?

Can sidestep the need for expensive mass measurements!

(but need to assume rocky core with H+He envelope)

Radial Velocities:

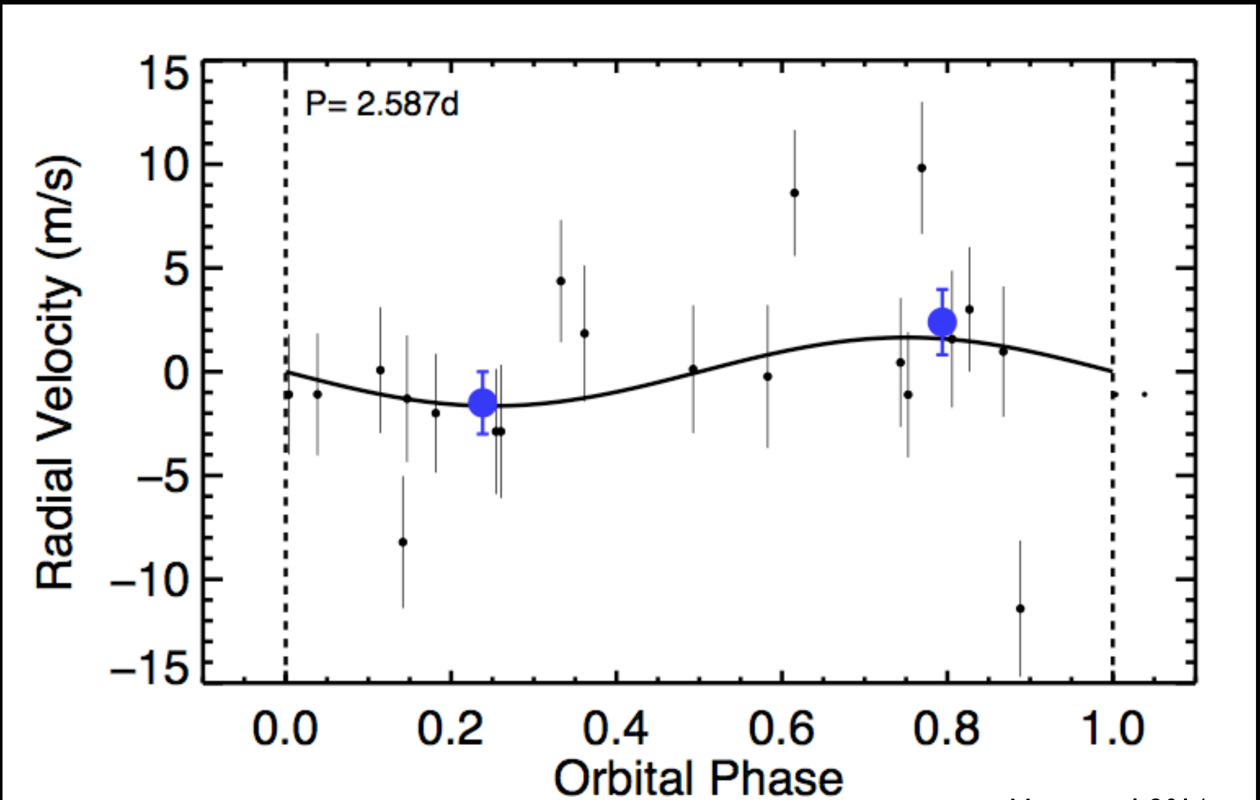

Marcy et al. 2014

Why is this exciting?

Can sidestep the need for expensive mass measurements!

(but need to assume rocky core with H+He envelope)

Radial Velocities:

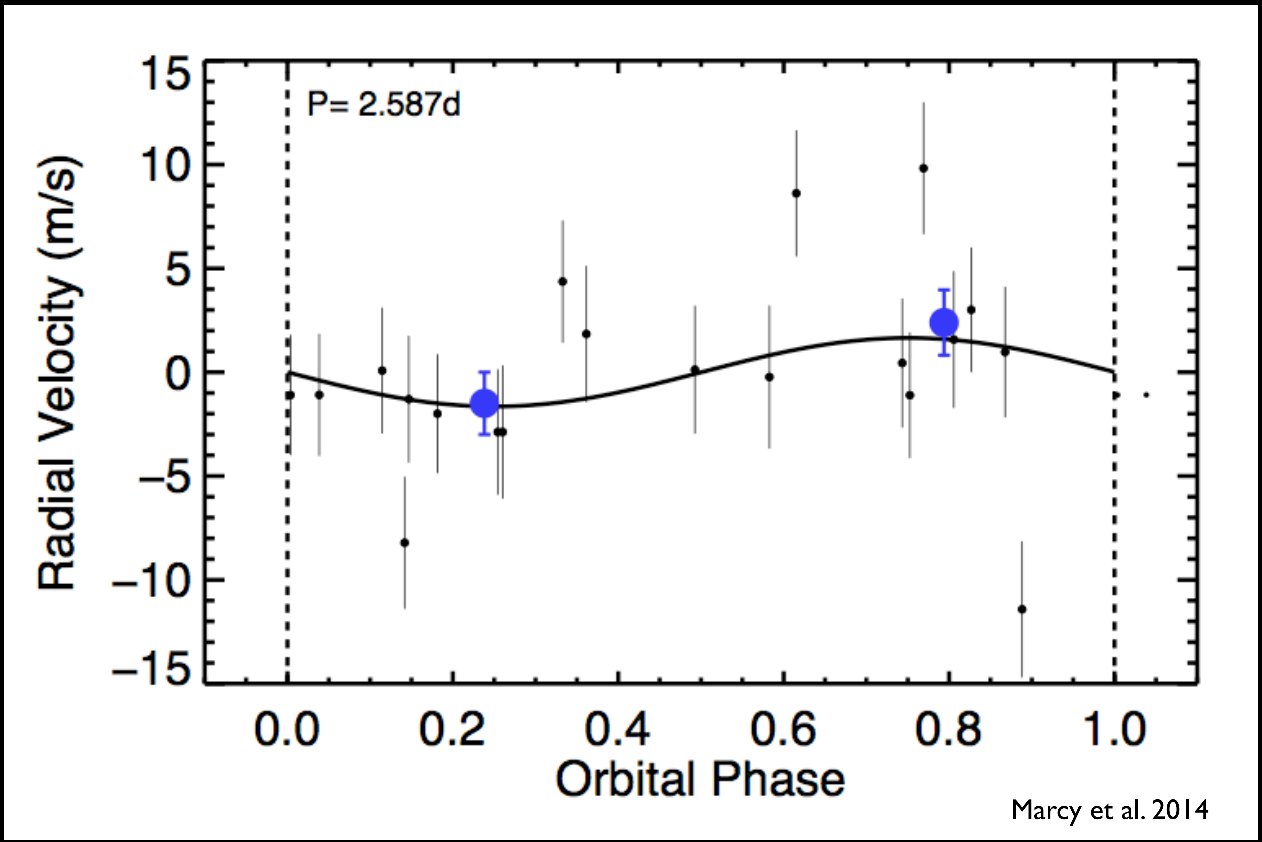

Requires years of observations for good phase coverage;

majority of Kepler targets too faint

Why is this exciting?

Can sidestep the need for expensive mass measurements!

(but need to assume rocky core with H+He envelope)

Radial Velocities:

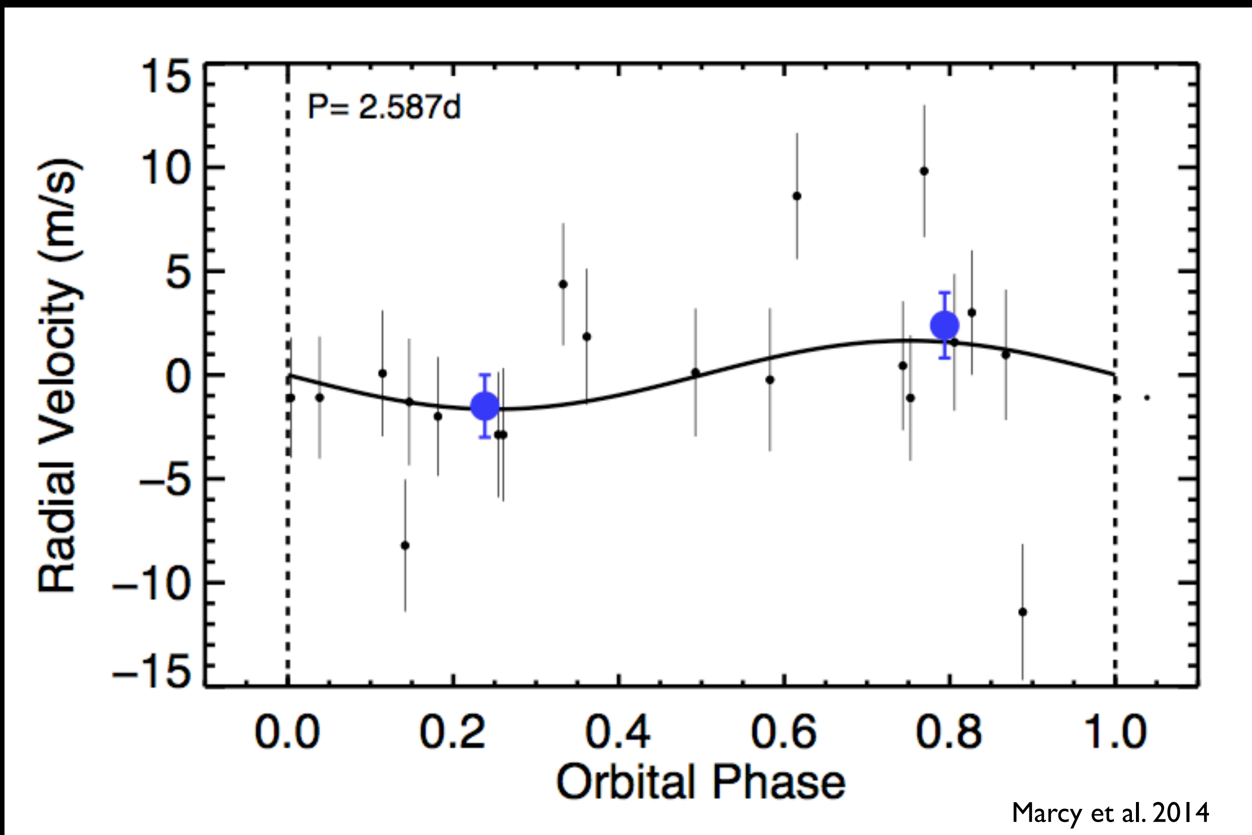

Transit Timing Variations:

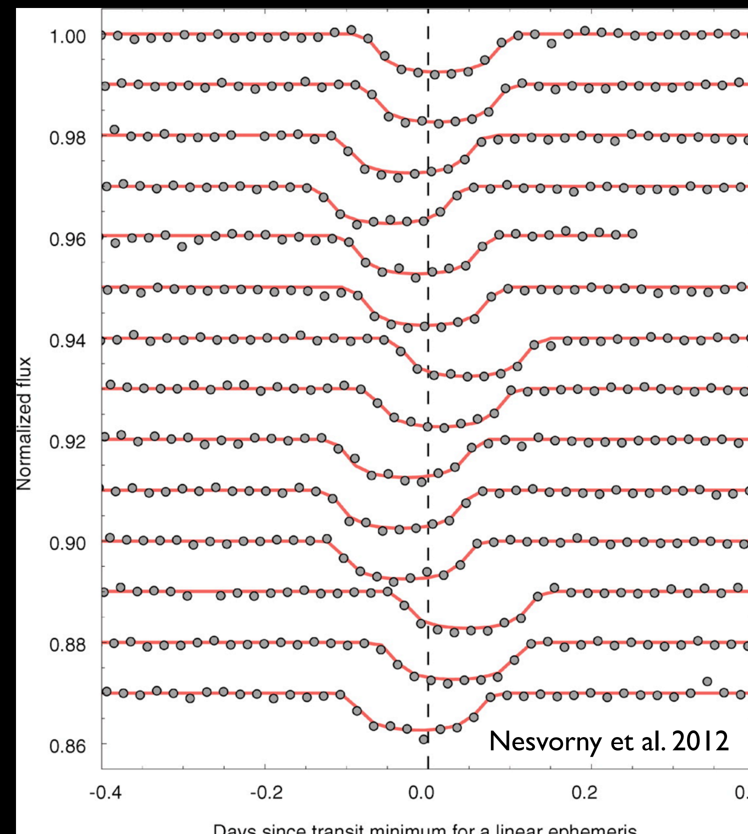

Requires years of observations for good phase coverage; majority of Kepler targets too faint

Why is this exciting?

Can sidestep the need for expensive mass measurements!

(but need to assume rocky core with H+He envelope)

Radial Velocities:

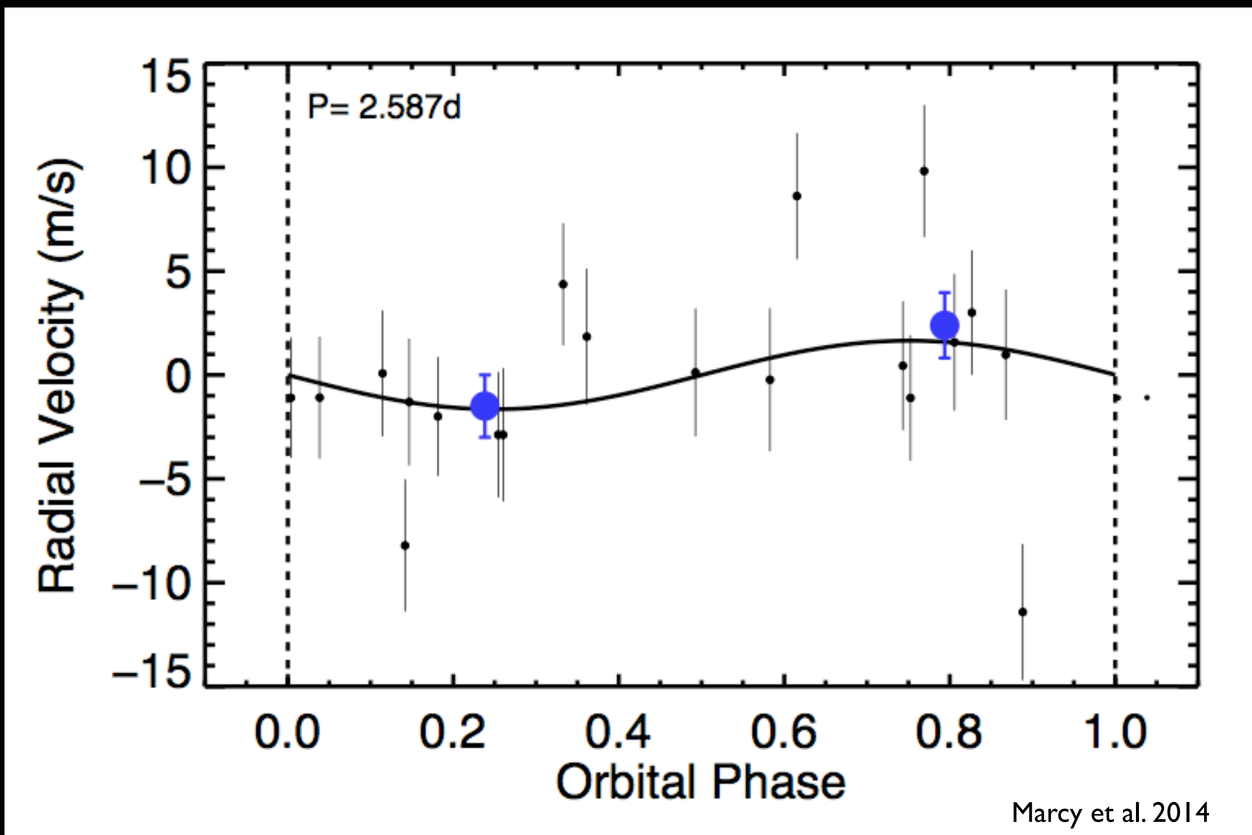

Requires years of observations for good phase coverage;
majority of Kepler targets too faint

Transit Timing Variations:

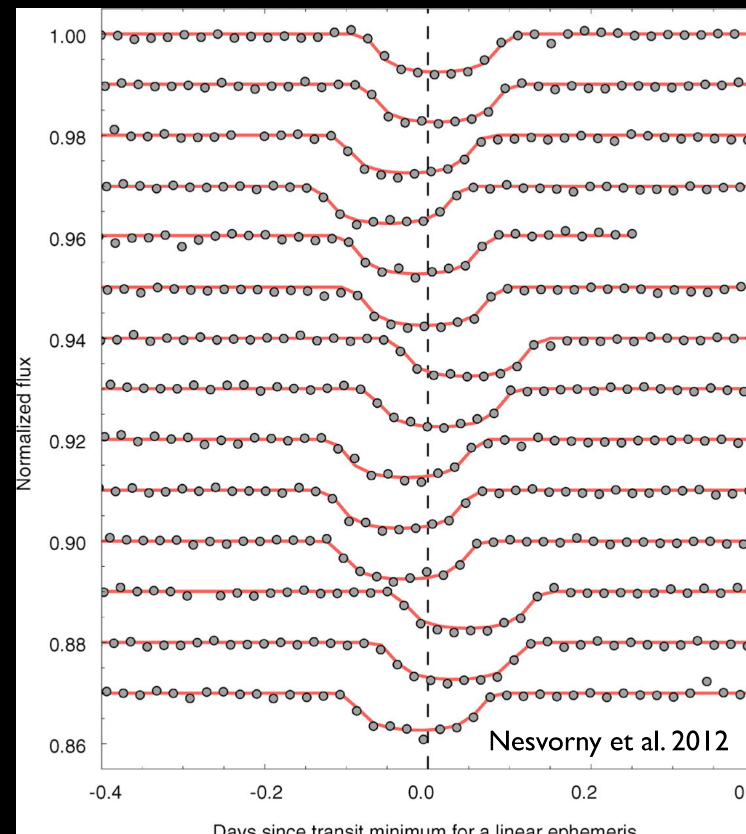

Planets must be in resonances, need high data cadence & long time baselines

Well Suited to HBM!

Well Suited to HBM!

“Regular” Bayes:

$$p(\theta|y) = \frac{p(y|\theta)p(\theta)}{\int_{\Theta} p(y|\tilde{\theta})p(\tilde{\theta})\ d\tilde{\theta}}$$

Well Suited to HBM!

“Regular” Bayes:

$$p(\theta|y) = \frac{p(y|\theta)p(\theta)}{\int_\Theta p(y|\tilde{\theta})p(\tilde{\theta})\ d\tilde{\theta}}$$

y = data

Well Suited to HBM!

“Regular” Bayes:

$$p(\theta|y) = \frac{p(y|\theta)p(\theta)}{\int_{\Theta} p(y|\tilde{\theta})p(\tilde{\theta})\ d\tilde{\theta}}$$

y = data

θ = the parameters of a model that can produce the data

Well Suited to HBM!

"Regular" Bayes:

$$p(\theta|y) = \frac{p(y|\theta)p(\theta)}{\int_{\Theta} p(y|\tilde{\theta})p(\tilde{\theta})\ d\tilde{\theta}}$$

y = data

θ = the parameters of a model that can produce the data

p() = probability density [distribution] of; | = "conditional on", or "given"

Well Suited to HBM!

“Regular” Bayes:

$$p(\theta|y) = \frac{p(y|\theta)p(\theta)}{\int_{\Theta} p(y|\tilde{\theta})p(\tilde{\theta})\ d\tilde{\theta}}$$

y = data
θ = the parameters of a model that can produce the data
p() = probability density [distribution] of; | = “conditional on”, or “given”

p(θ) = prior probability
(how probable are the possible values that you think θ could take?)

Well Suited to HBM!

“Regular” Bayes:

$$p(\theta|y) = \frac{p(y|\theta)p(\theta)}{\int_\Theta p(y|\tilde{\theta})p(\tilde{\theta})\ d\tilde{\theta}}$$

y = data
θ = the parameters of a model that can produce the data
p() = probability density [distribution] of; | = “conditional on”, or “given”

p(θ) = prior probability
(how probable are the possible values that you think θ could take?)

p(y|θ) = likelihood, or sampling distribution
(ties your model to the data probabilistically:
how likely is the data you observed given specific model parameters?)

Well Suited to HBM!

"Regular" Bayes:

$$p(\theta|y) = \frac{p(y|\theta)p(\theta)}{\int_\Theta p(y|\tilde{\theta})p(\tilde{\theta})\ d\tilde{\theta}}$$

y = data
θ = the parameters of a model that can produce the data
p() = probability density [distribution] of; | = "conditional on", or "given"

p(θ) = prior probability
(how probable are the possible values that you think θ could take?)

p(y|θ) = likelihood, or sampling distribution
(ties your model to the data probabilistically:
how likely is the data you observed given specific model parameters?)

(Often as χ^2)

Well Suited to HBM!

"Regular" Bayes:

$$p(\theta|y) = \frac{p(y|\theta)p(\theta)}{\int_\Theta p(y|\tilde{\theta})p(\tilde{\theta})\ d\tilde{\theta}}$$

y = data
θ = the parameters of a model that can produce the data
p() = probability density [distribution] of; | = "conditional on", or "given"

p(θ) = prior probability
(how probable are the possible values that you think θ could take?)

p(y|θ) = likelihood, or sampling distribution
(ties your model to the data probabilistically:
how likely is the data you observed given specific model parameters?)

(Often as χ^2)

p(θ|y) = posterior probability
(a "new prior" distribution, updated with information contained by the data:
what is the distribution of θ values given the data **and** your beliefs?)

Well Suited to HBM!

"Regular" Bayes:

$$p(\theta|y) = \frac{p(y|\theta)p(\theta)}{\int_{\Theta} p(y|\tilde{\theta})p(\tilde{\theta})\ d\tilde{\theta}}$$ = p(y) = constant

y = data
θ = the parameters of a model that can produce the data
p() = probability density [distribution] of; | = "conditional on", or "given"

p(θ) = prior probability
(how probable are the possible values that you think θ could take?)

p(y|θ) = likelihood, or sampling distribution
(ties your model to the data probabilistically:
how likely is the data you observed given specific model parameters?)

(Often as χ^2)

p(θ|y) = posterior probability
(a "new prior" distribution, updated with information contained by the data:
what is the distribution of θ values given the data **and** your beliefs?)

Well Suited to HBM!

"Regular" Bayes:

$$p(\theta|y) \propto p(y|\theta)\, p(\theta)$$

y = data
θ = the parameters of a model that can produce the data
p() = probability density [distribution] of; | = "conditional on", or "given"

$p(\theta)$ = prior probability
(how probable are the possible values that you think θ could take?)

$p(y|\theta)$ = likelihood, or sampling distribution
(ties your model to the data probabilistically:
how likely is the data you observed given specific model parameters?)

(Often as χ^2)

$p(\theta|y)$ = posterior probability
(a "new prior" distribution, updated with information contained by the data:
what is the distribution of θ values given the data **and** your beliefs?)

HBM: What is it?

Hierarchical Bayes:

HBM: What is it?

Hierarchical Bayes:

What if there isn’t just one “true” value of θ for all the data?

HBM: What is it?

Hierarchical Bayes:

What if there isn't just one "true" value of θ for all the data?
i.e. θ has its own intrinsic distribution?

HBM: What is it?

Hierarchical Bayes:

What if there isn't just one "true" value of θ for all the data?
i.e. θ has its own intrinsic distribution?

$$p(\theta,\alpha|y) \propto p(y|\theta,\alpha)\, p(\theta|\alpha)\, p(\alpha)$$

HBM: What is it?

Hierarchical Bayes:

What if there isn't just one "true" value of θ for all the data?
i.e. θ has its own intrinsic distribution?

posterior

$$p(\theta,\alpha|y) \propto p(y|\theta,\alpha)\, p(\theta|\alpha)\, p(\alpha)$$

HBM: What is it?

Hierarchical Bayes:

What if there isn't just one "true" value of θ for all the data?
i.e. θ has its own intrinsic distribution?

posterior likelihood

$$p(\theta,\alpha|y) \propto p(y|\theta,\alpha)\, p(\theta|\alpha)\, p(\alpha)$$

HBM: What is it?

Hierarchical Bayes:

What if there isn't just one "true" value of θ for all the data?
i.e. θ has its own intrinsic distribution?

posterior likelihood "prior"

$$\mathbf{p(\theta,\alpha|y) \propto p(y|\theta,\alpha)\, p(\theta|\alpha)\, p(\alpha)}$$

HBM: What is it?

Hierarchical Bayes:

What if there isn't just one "true" value of θ for all the data?
i.e. θ has its own intrinsic distribution?

$$\underbrace{p(\theta,\alpha|y)}_{\text{posterior}} \propto \underbrace{p(y|\theta,\alpha)}_{\text{likelihood}}\; p(\theta|\alpha)\; \underbrace{p(\alpha)}_{\text{“prior”}}$$

α = hyperparameters
(the parameters that describe the distribution of θ values)

HBM: What is it?

Hierarchical Bayes:

What if there isn't just one "true" value of θ for all the data?
i.e. θ has its own intrinsic distribution?

posterior likelihood "prior"

$$p(\theta,\alpha|y) \propto p(y|\theta,\alpha)\, p(\theta|\alpha)\, p(\alpha)$$

α = hyperparameters
(the parameters that describe the distribution of θ values)

$p(\theta|\alpha)$ = this "intrinsic distribution" for the parameters

HBM: What is it?

Hierarchical Bayes:

What if there isn't just one "true" value of θ for all the data?
i.e. θ has its own intrinsic distribution?

posterior | likelihood | "prior"

$$p(\theta,\alpha|y) \propto p(y|\theta,\alpha)\, p(\theta|\alpha)\, p(\alpha)$$

α = hyperparameters
(the parameters that describe the distribution of θ values)

$p(\theta|\alpha)$ = this "intrinsic distribution" for the parameters

$p(\alpha)$ = prior probability of the hyperparameters

HBM: What is it?

Hierarchical Bayes:

What if there isn't just one "true" value of θ for all the data?
i.e. θ has its own intrinsic distribution?

posterior likelihood "prior"

$$p(\theta,\alpha|y) \propto p(y|\theta,\alpha)\, p(\theta|\alpha)\, p(\alpha)$$

α = hyperparameters
(the parameters that describe the distribution of θ values)

p(θ|α) = this "intrinsic distribution" for the parameters

p(α) = prior probability of the hyperparameters

Happens often for population studies: data tends to be "grouped" (hierarchical),

HBM: What is it?

Hierarchical Bayes:

What if there isn't just one "true" value of θ for all the data?
i.e. θ has its own intrinsic distribution?

posterior — likelihood — "prior"

$$\underbrace{p(\theta,\alpha|y)}_{\text{posterior}} \propto \underbrace{p(y|\theta,\alpha)}_{\text{likelihood}}\ \underbrace{p(\theta|\alpha)}_{\text{“prior”}}\ p(\alpha)$$

α = hyperparameters
(the parameters that describe the distribution of θ values)

p(θ|α) = this "intrinsic distribution" for the parameters

p(α) = prior probability of the hyperparameters

Happens often for population studies: data tends to be "grouped" (hierarchical), and one value of θ may not be appropriate for all groups

HBM: What is it?

Hierarchical Bayes:

What if there isn't just one "true" value of θ for all the data?
i.e. θ has its own intrinsic distribution?

posterior likelihood "prior"

$$p(\theta,\alpha|y) \propto p(y|\theta,\alpha)\, p(\theta|\alpha)\, p(\alpha)$$

can still use MCMC!

α = hyperparameters
(the parameters that describe the distribution of θ values)

$p(\theta|\alpha)$ = this "intrinsic distribution" for the parameters

$p(\alpha)$ = prior probability of the hyperparameters

Happens often for population studies: data tends to be "grouped" (hierarchical), and one value of θ may not be appropriate for all groups

HBM: What is it?

Hierarchical Bayes:

What if there isn't just one "true" value of θ for all the data?
i.e. θ has its own intrinsic distribution?

posterior / likelihood / "prior"

$$\underbrace{p(\theta,\alpha|y)}_{\text{posterior}} \propto \underbrace{p(y|\theta,\alpha)}_{\text{likelihood}}\ \underbrace{p(\theta|\alpha)\ p(\alpha)}_{\text{"prior"}}$$

can still use MCMC!

α = hyperparameters
(the parameters that describe the distribution of θ values)

$p(\theta|\alpha)$ = this "intrinsic distribution" for the parameters

$p(\alpha)$ = prior probability of the hyperparameters

Happens often for population studies: data tends to be "grouped" (hierarchical),
and one value of θ may not be appropriate for all groups

Adding another layer of probabilistic structure

How much to believe this?

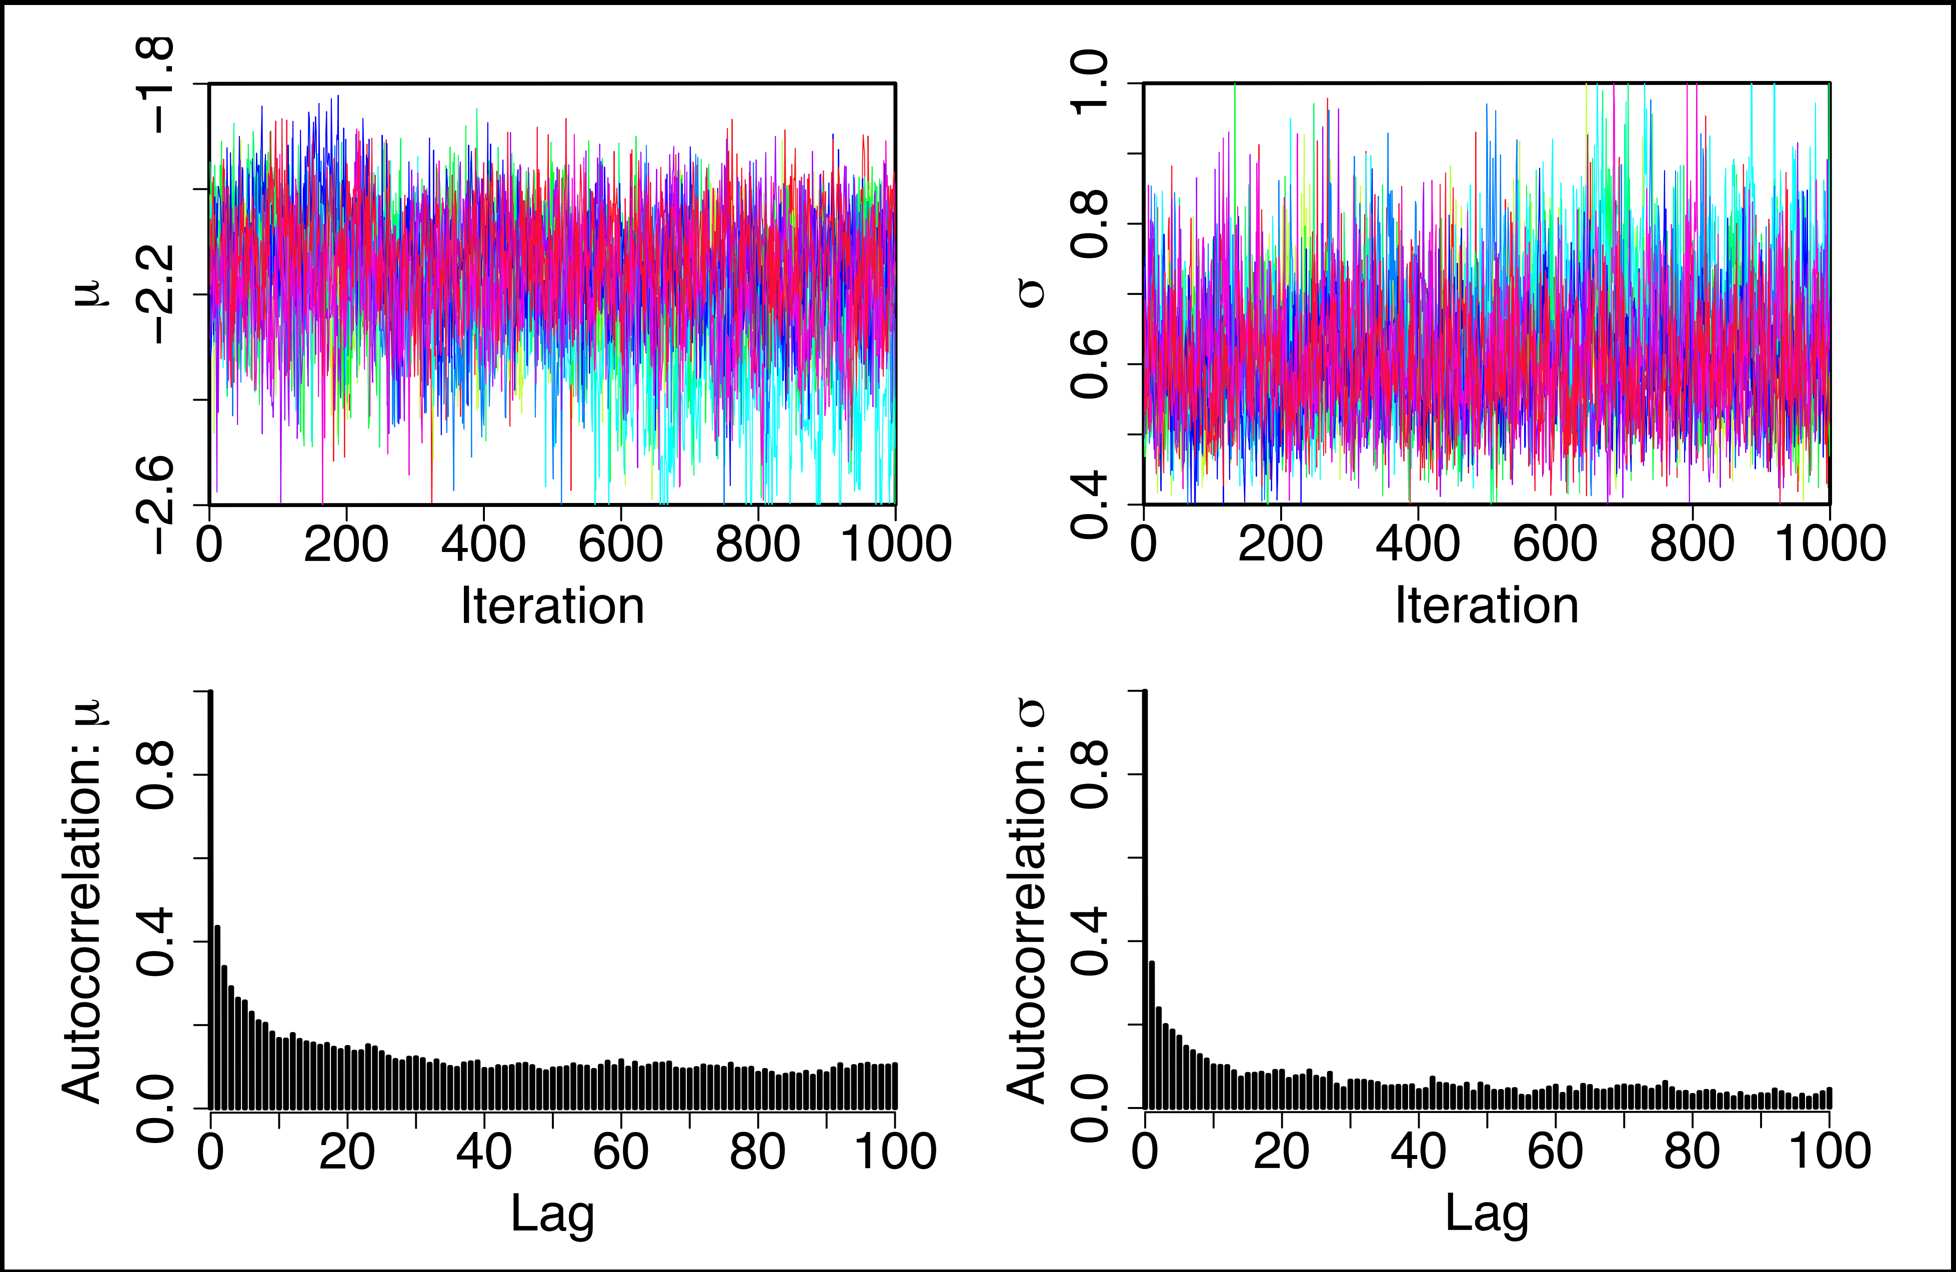

How much to believe this?

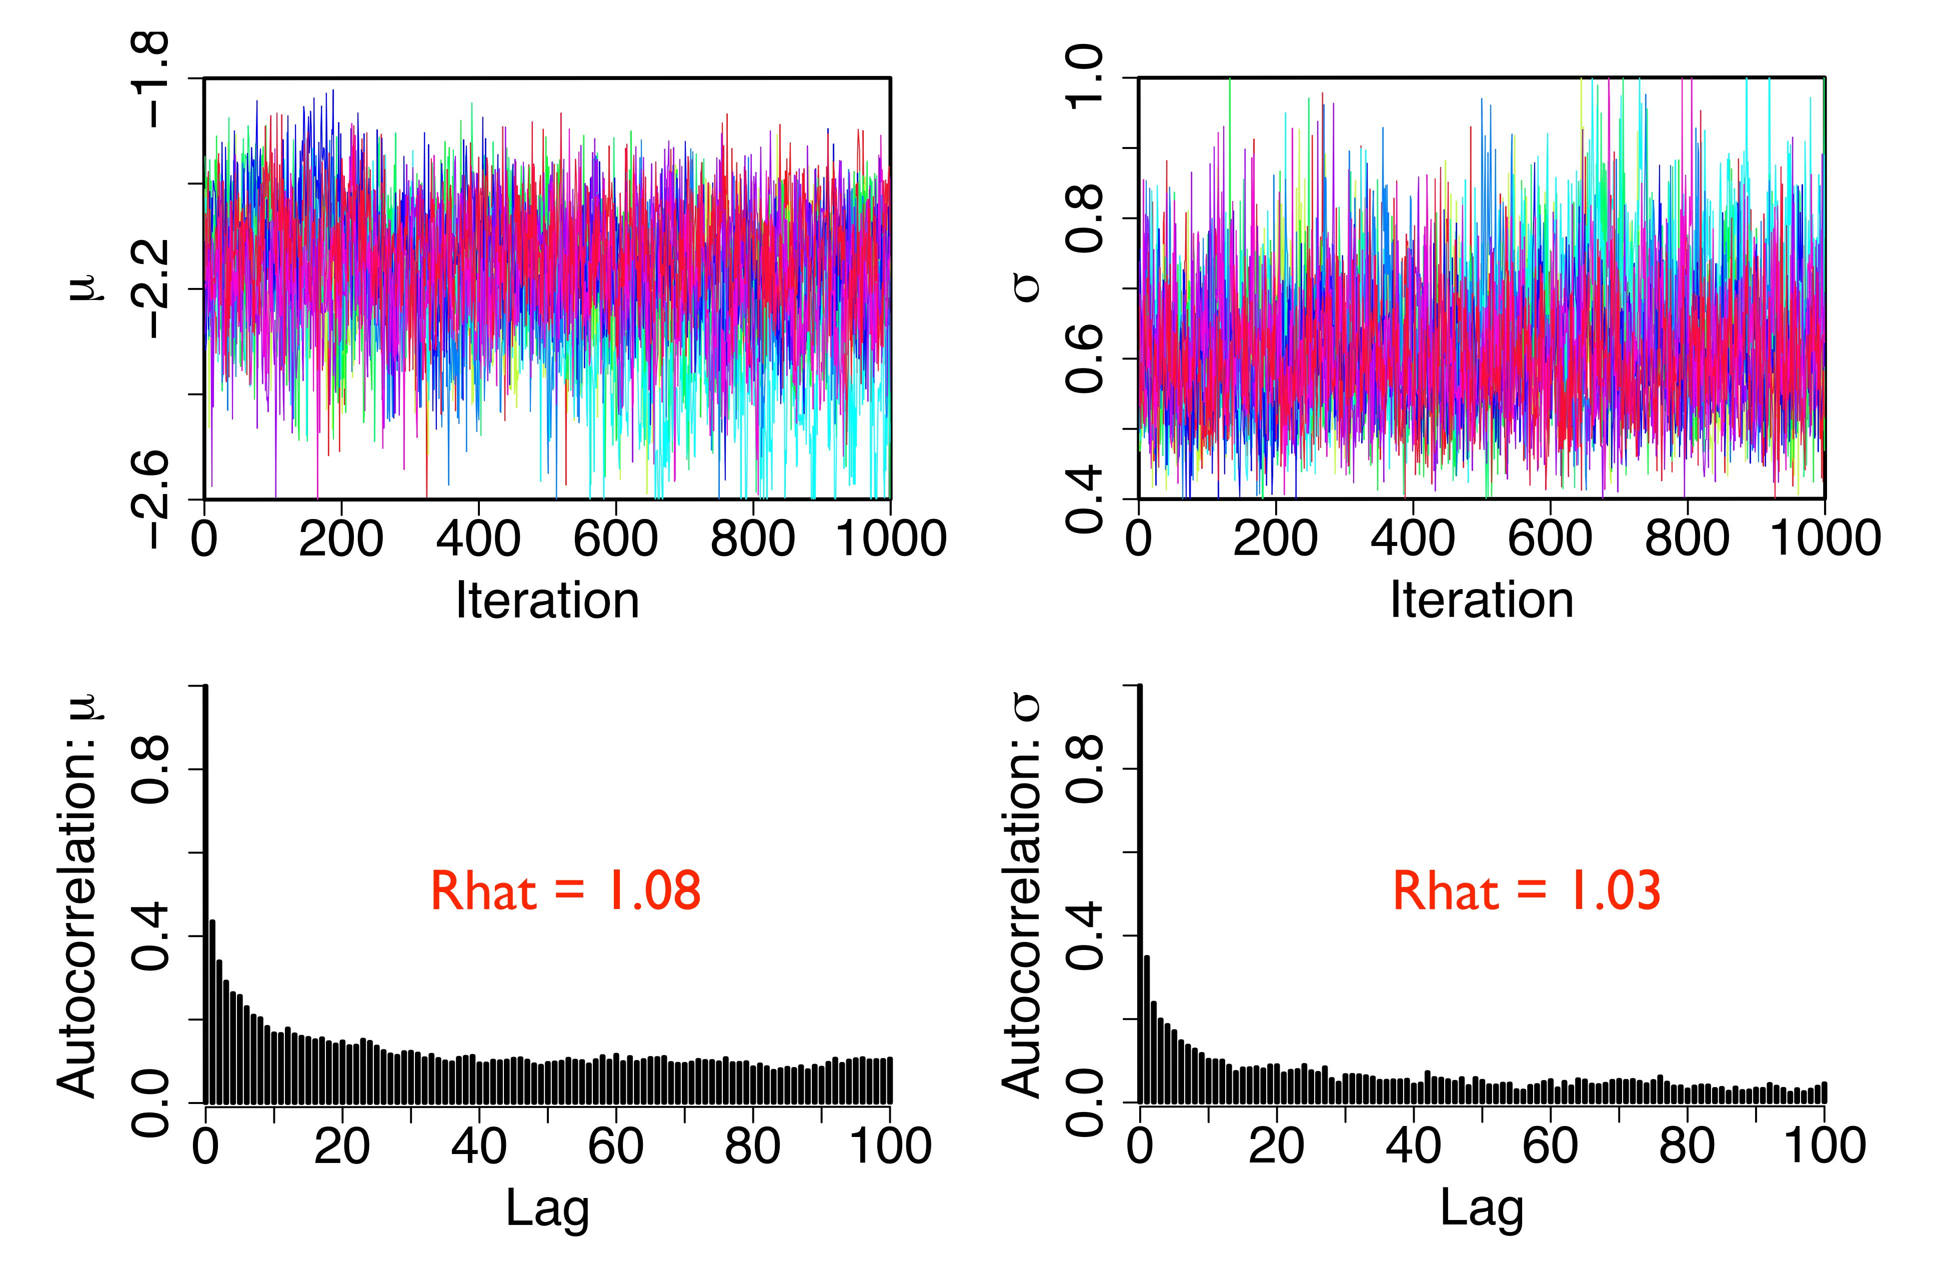

How much to believe this?

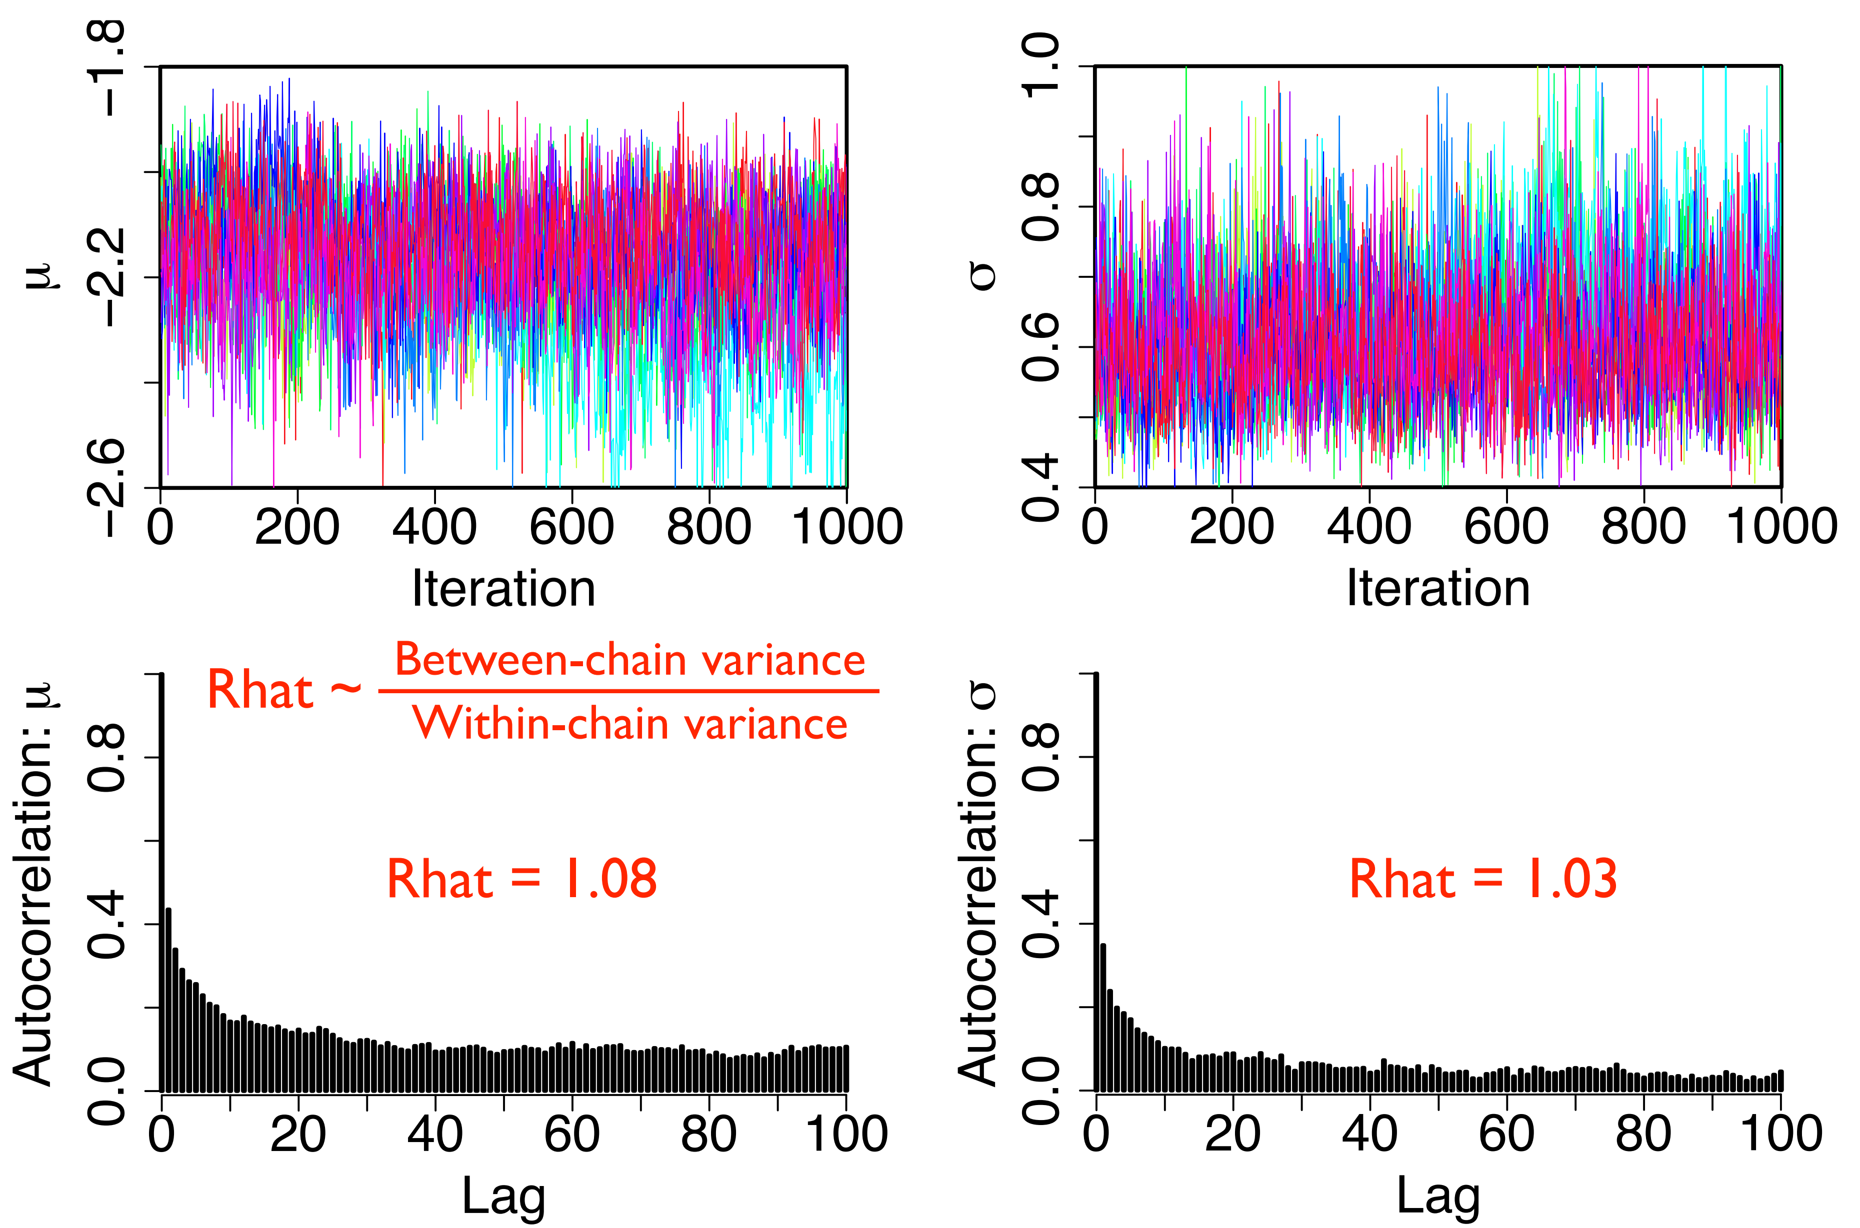

How γ affects composition?

How γ affects composition?

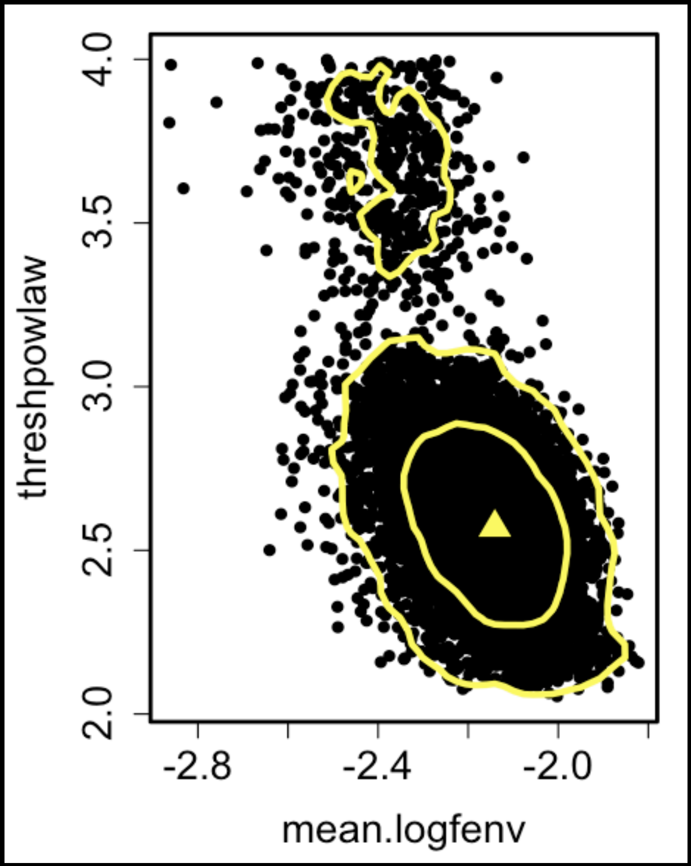

How γ affects composition?

Marginalized over γ
(threshpowlaw)

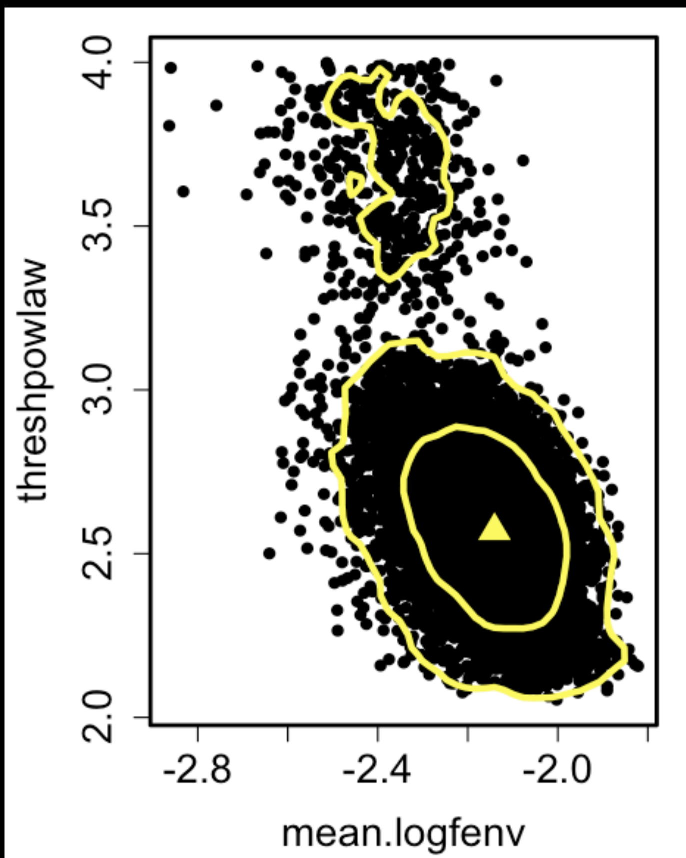

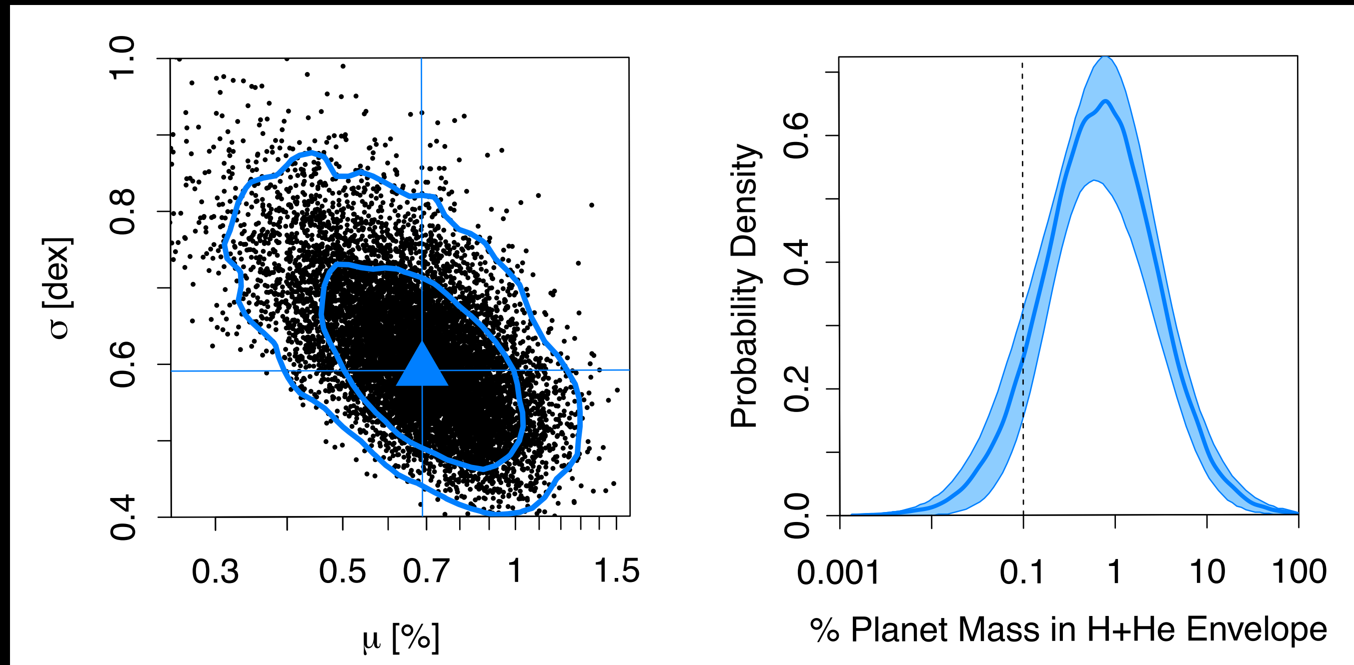

How γ affects composition?

Marginalized over γ
(threshpowlaw)

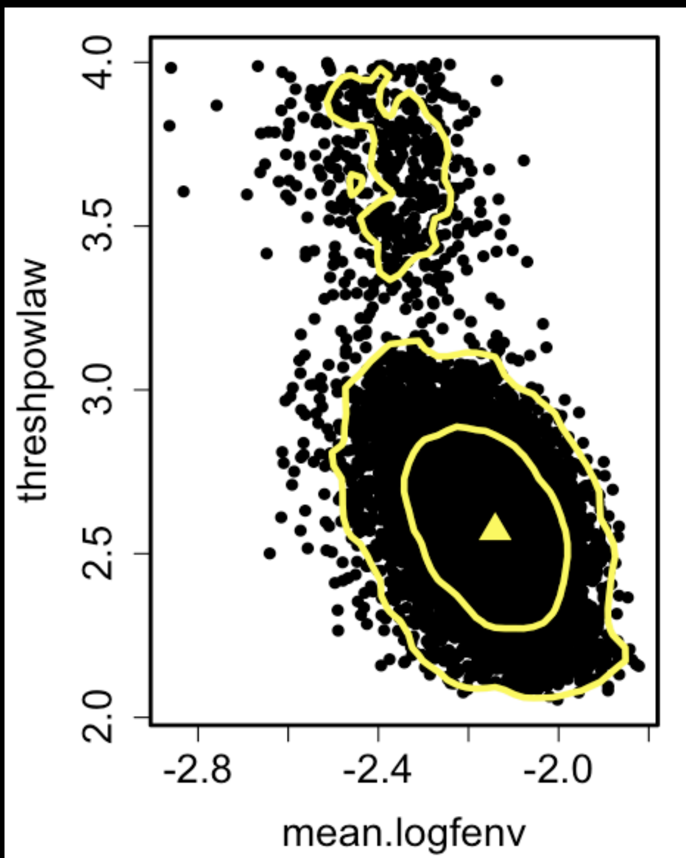

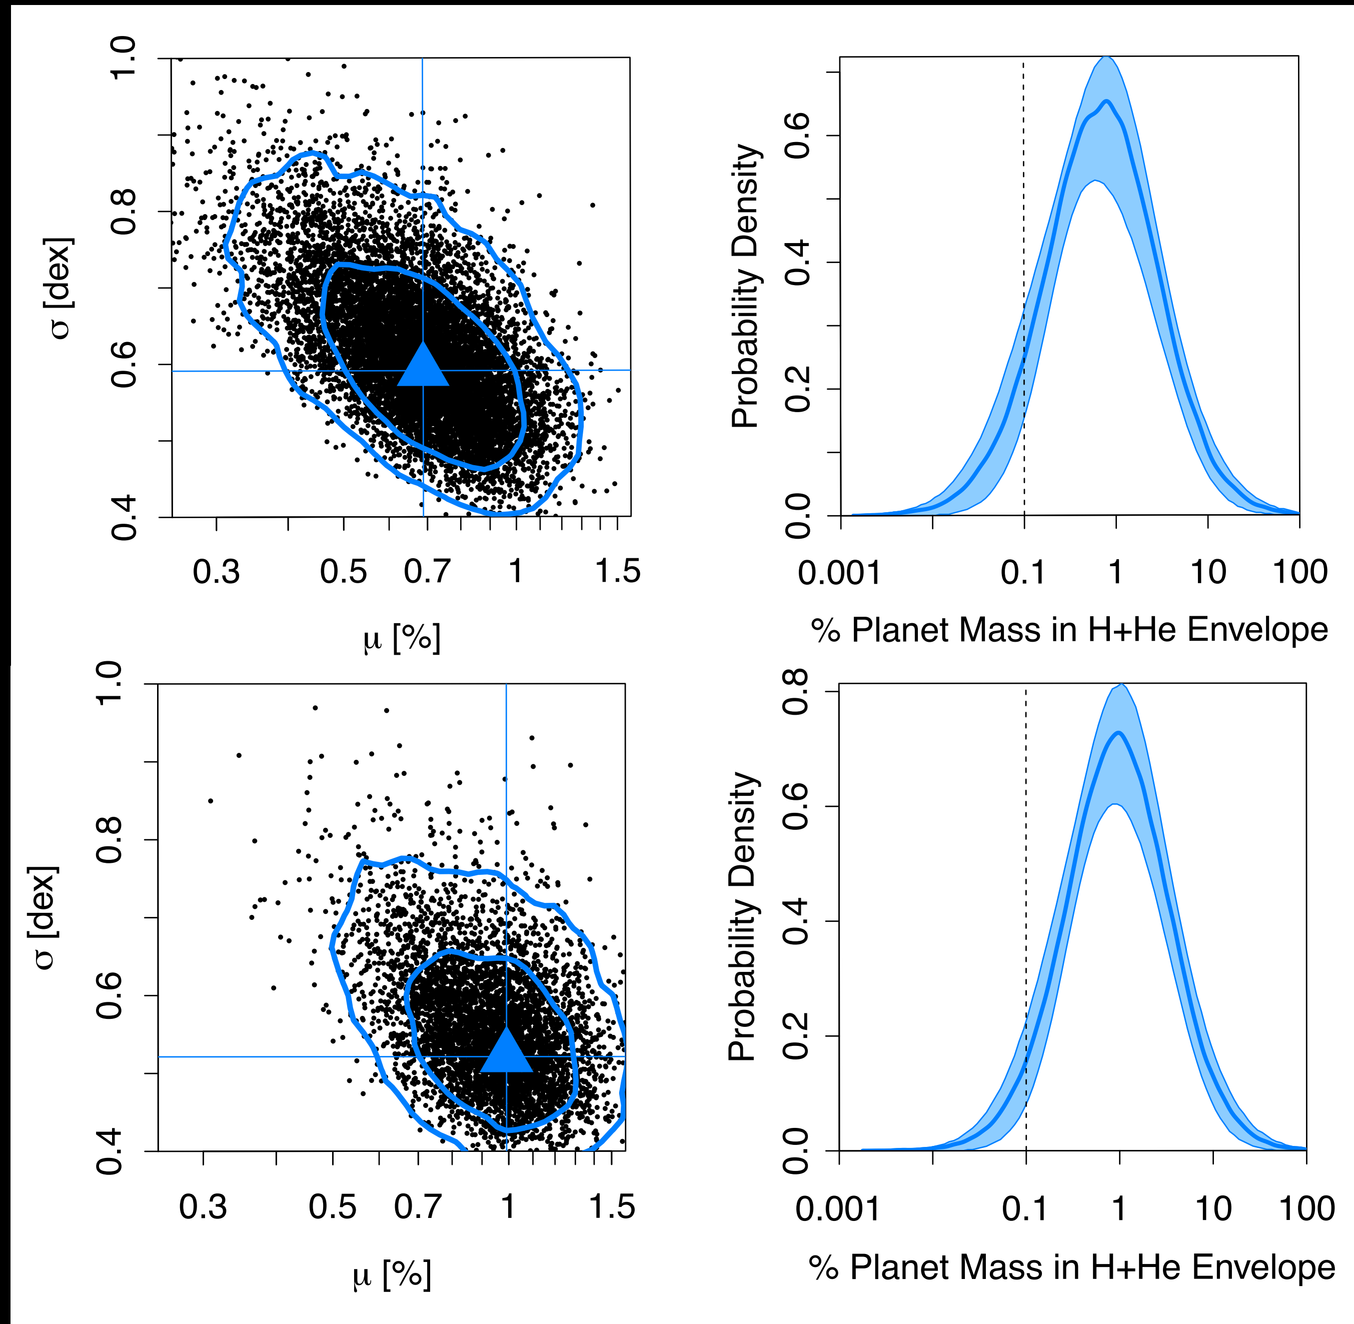

γ = 2

One last sanity check