Bayesian Balancing of (Hybrid) Supply–Use Tables with MCMC

Towards an uncertainty-aware BONSAI footprint database

Simon Schulte

University of Freiburg (Germany)

IIOA conference Seville 2026

25.06.2026

When complex IO models … become simple decisions

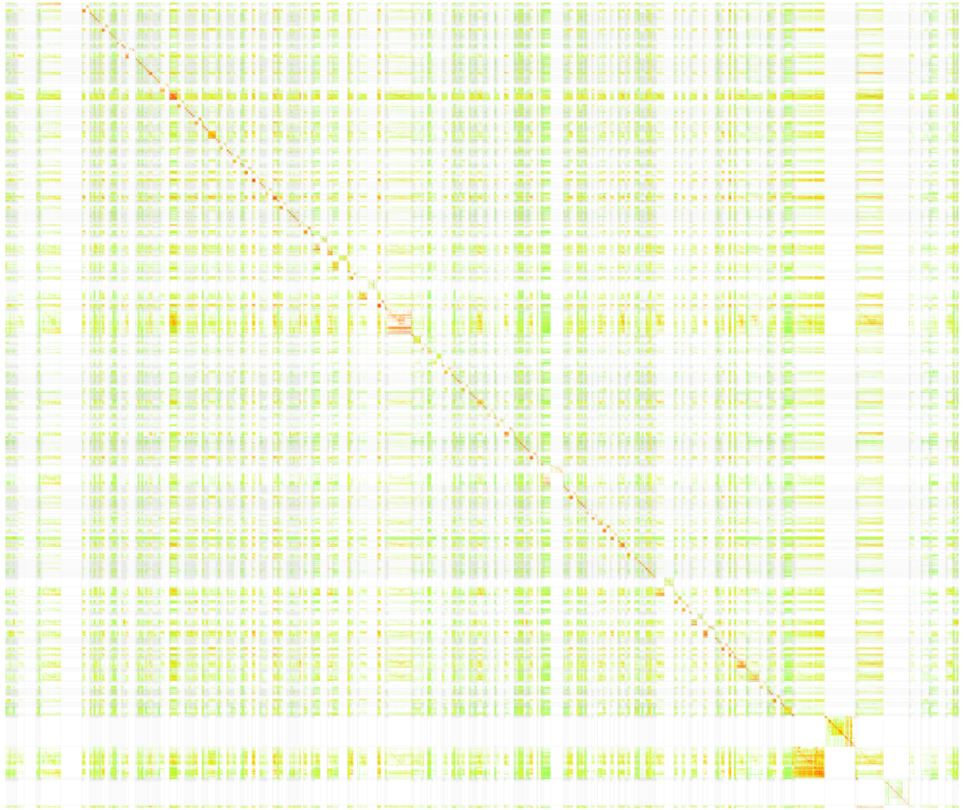

source: worldmrio.com/heatmap/

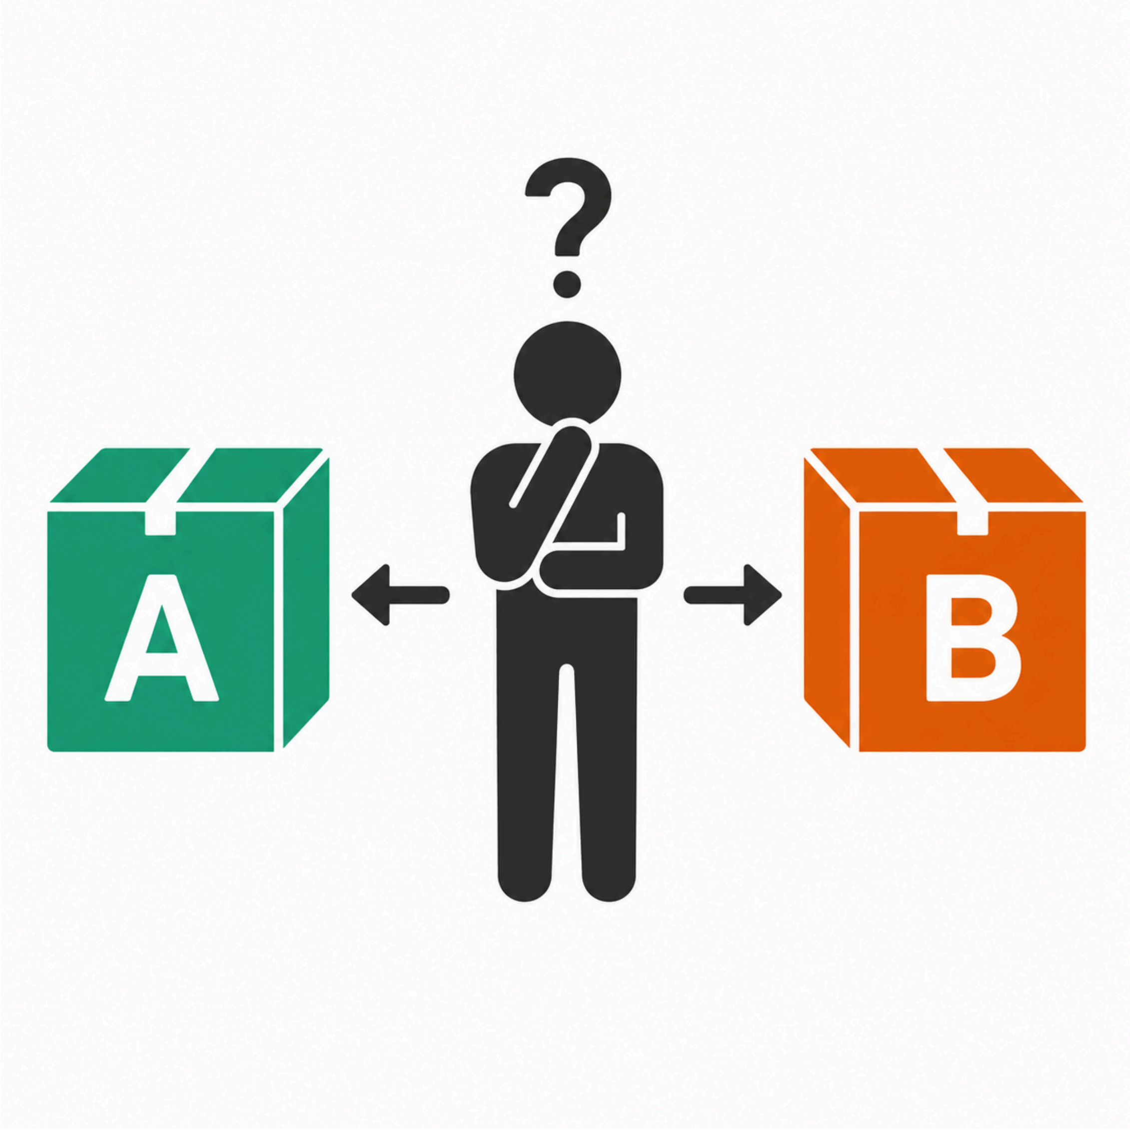

A simple question: is A < B?

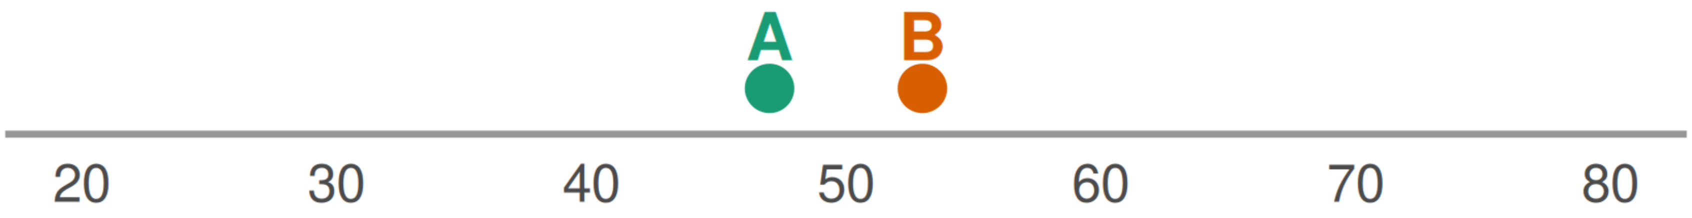

- Based on the point estimates alone, the answer looks obvious.
- But point estimates hide uncertainty and dependence.

Add uncertainty

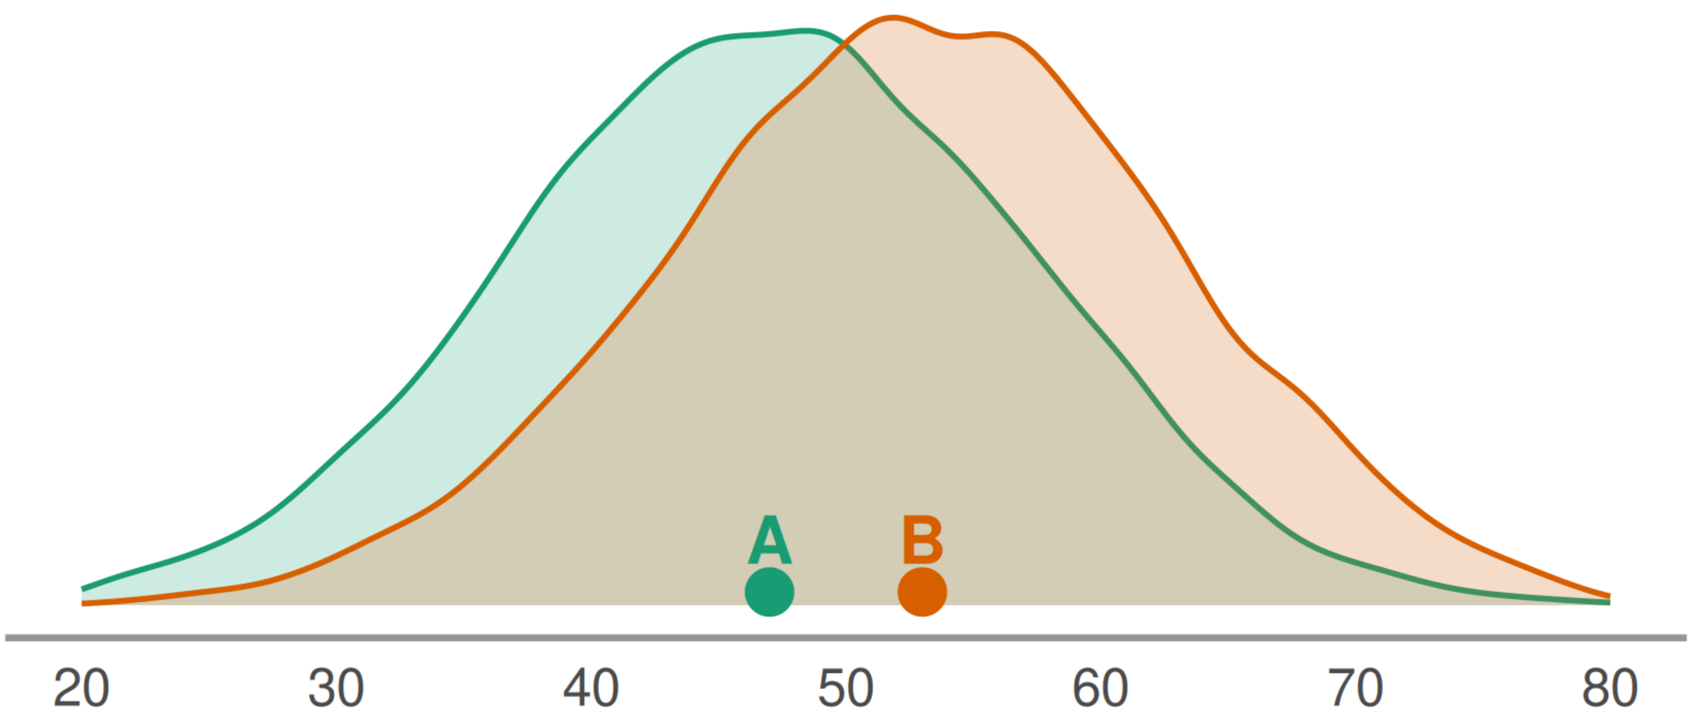

- The means differ, but the marginals overlap.
- So the question becomes probabilistic: what is `P(A < B)`?

Same marginals, different dependence

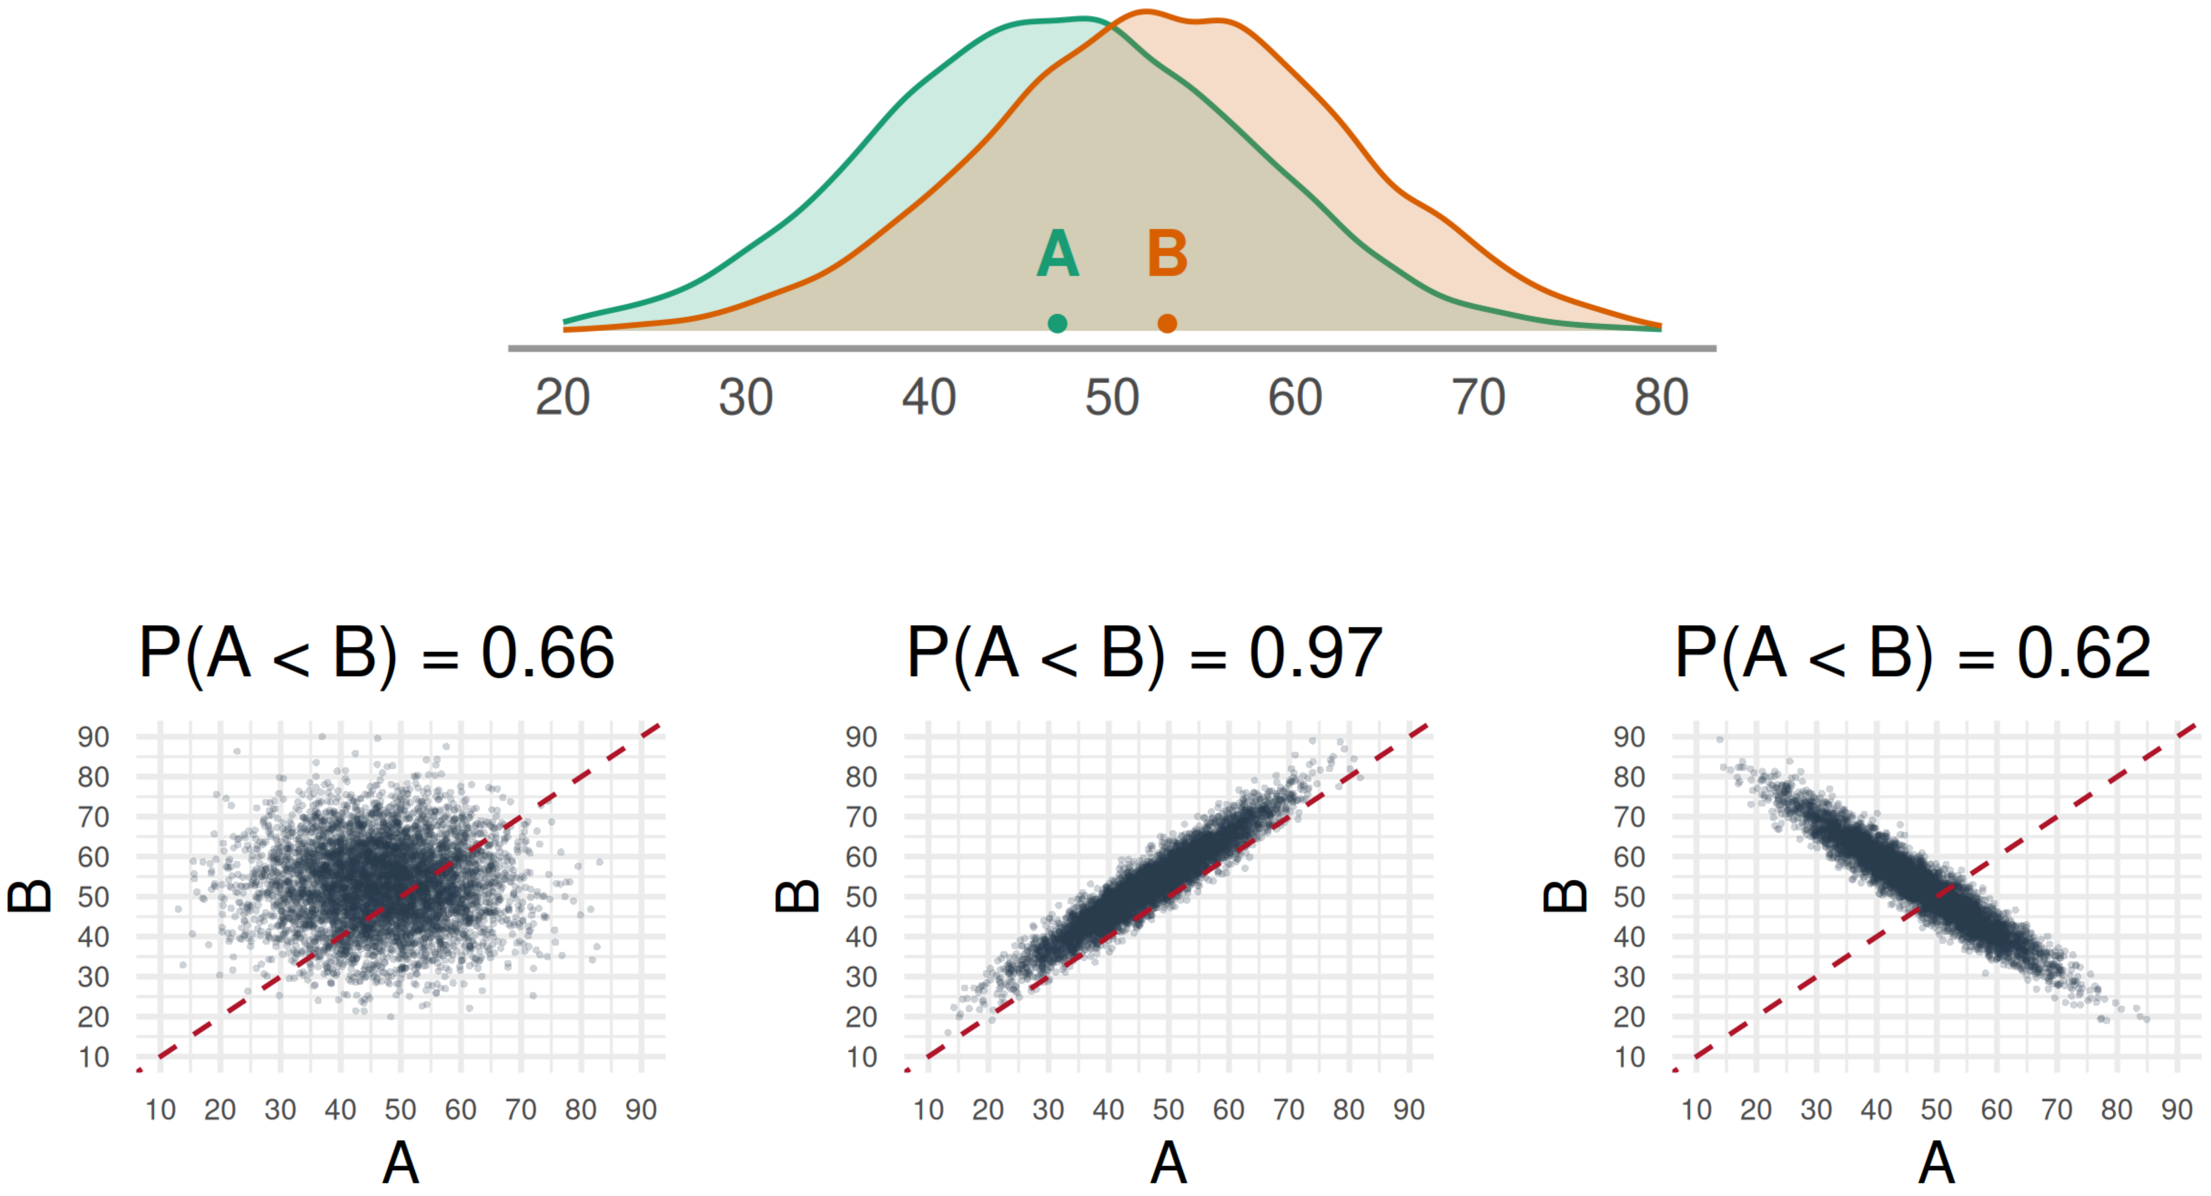
A
B
20
30
40
50
60
70
80
P(A < B) = 0.66
P(A < B) = 0.97
P(A < B) = 0.62
B
A
10
20
30
40
50
60
70
80
90

Why IO/SUT results are not independent

- IOTs/SUTs are compiled from heterogeneous, incomplete and sometimes conflicting data sources
- need for **balancing** (reconciliation) to satisfy key accounting and consistency constraints: supply = demand, input = output
- accounting constraints induce **dependencies** between flows: more output requires more input, more demand requires more supply
- Standard balancing (RAS, KRAS, WLS, cross-entropy, …) yields **a single reconciled table** (often without uncertainty and dependencies)

The challenge

We need a balancing method that:

1. can use heterogeneous **reliability information** in the balancing step
2. provides the **uncertainty for the balanced table**
3. accounts for **dependence induced by constraints**
4. supports for **downstream uncertainty propagation** (e.g. $P(\mathrm{A}) > P(\mathrm{B})$)
5. **scales** to the size of BONSAI

BONSAI

- Open climate footprint database: bonsai.uno
- **Hybrid Supply-Use** system with multiple property layers
- Layers are coupled via **conversion factors** (e.g., €/t, TJ/t, €/TJ)
- dimensions: ~1700 products × ~1500 activities × 49 regions → **>1.8M non-zero flows**

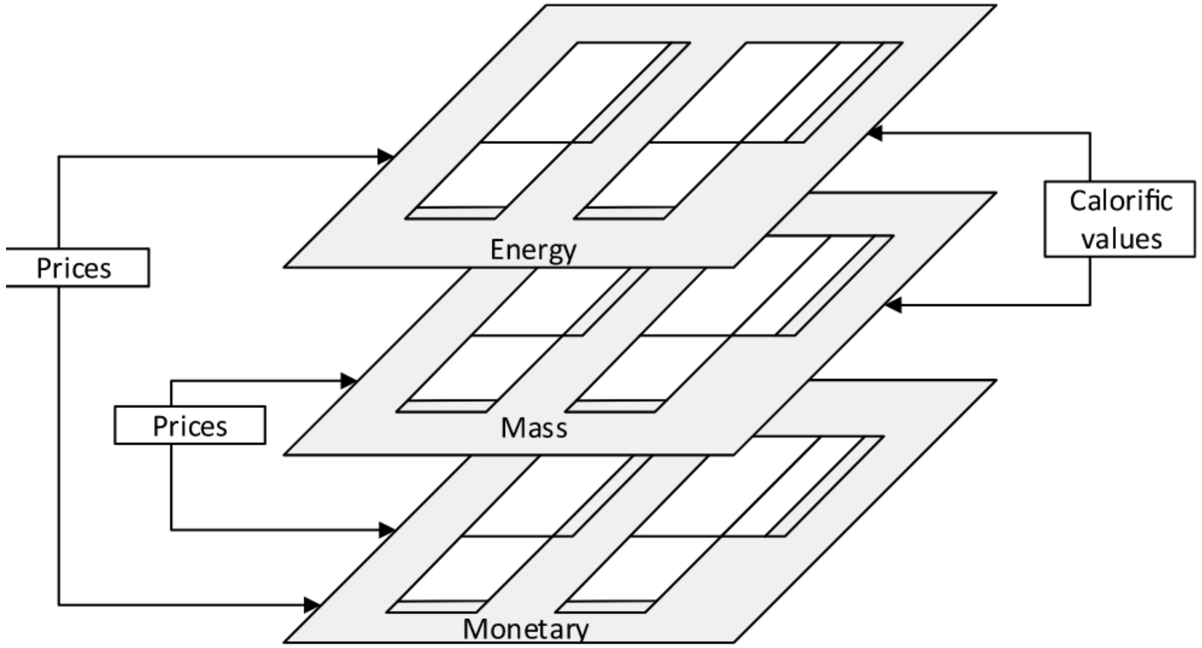

source: Merciai (2019)

The Bayes' Theorem

$$p(\mathrm{SUT} \mid \mathrm{constraints}) = \frac{p(\mathrm{constraints} \mid \mathrm{SUT}) \times p(\mathrm{SUT})}{p(\mathrm{constraints})}$$

- $p(\mathrm{SUT})$ = **Prior distribution**
- $p(\mathrm{constraints} \mid \mathrm{SUT})$= **Likelihood**
- $p(\mathrm{SUT} \mid \mathrm{constraints})$ = **Posterior distribution**
- $p(\mathrm{constraints})$ = Normalising constant

=> everything (supply/use flows, conversion factors, constraints) is a probability distribution!

$$p(\mathrm{SUT} \mid \mathrm{constraints}) = \frac{p(\mathrm{constraints} \mid \mathrm{SUT}) \times \underline{p(\mathrm{SUT})}}{p(\mathrm{constraints})}$$

Prior distribution:

- Our prior belief about the Supply & Use flows, conversion factors and transfer coefficients,
- centered on the initial (unbalanced) SUT,
- uncertainty/variance: ad-hoc approach for proof-of-concept
  - flows subject to small imbalances -> rel. small variance,
  - flows subject to large imbalances -> rel. large variance
  - clamping to $[0.1, 0.5]$

$$p(\text{SUT} \mid \text{constraints}) = \frac{\underline{p(\text{constraints} \mid \text{SUT})} \times p(\text{SUT})}{p(\text{constraints})}$$

Likelihood:

- How well a candidate SUT satisfies the accounting identities/constraints. SUTs that violate constraints get low probability.
- Constraints used:
  - product & activity balances (per region, per layer)
  - cross-layer consistency via conversion factors
  - transfer-coefficient consistency

$$p(\mathrm{SUT} \mid \mathrm{constraints}) = \frac{\underline{p(\mathbf{constraints} \mid \mathbf{SUT})} \times p(\mathrm{SUT})}{p(\mathrm{constraints})}$$

$$p(\mathrm{SUT} \mid \mathrm{constraints}) = \frac{p(\mathrm{constraints} \mid \mathrm{SUT}) \times p(\mathrm{SUT})}{p(\mathrm{constraints})}$$

Posterior distribution:

- The distribution of balanced SUTs that are close to the initial data & consistent with all structural constraints.
- we cannot compute the posterior in a closed-form
- use MCMC (Monte Carlo Markov Chains) to approximate posterior

MCMC to approximate posterior

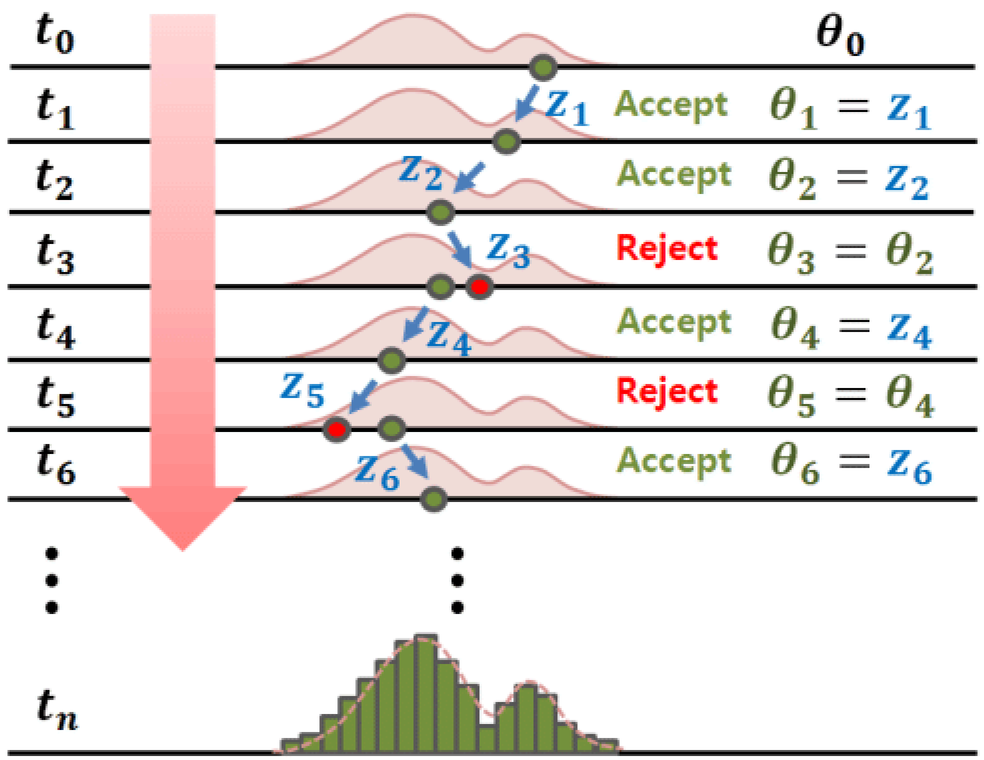

1. Find starting parameter position
2. Propose to jump from that position to somewhere else
3. Evaluate: Is it a "good" place to jump or not?
   $\mathrm{P(accept)} = \mathrm{P(proposal)} \;/\; \mathrm{P(current)}$
4. Each accepted state is recorded
5. By running enough iterations we will approach the posterior distribution

1

Implementation & Computation

- Implemented in Stan (using NUTS)
- Successfully tested at BONSAI scale (>1.8 Million flows, multi-layer)
- run on HPC: 4 parallel chains and 500-1000 iterations per chain required on the order of 19 to 44 hours and 24 GB of RAM
- processing the Stan output files with the current implementation requires around 200 GB of RAM

Results

1a) how much does balancing move the table?

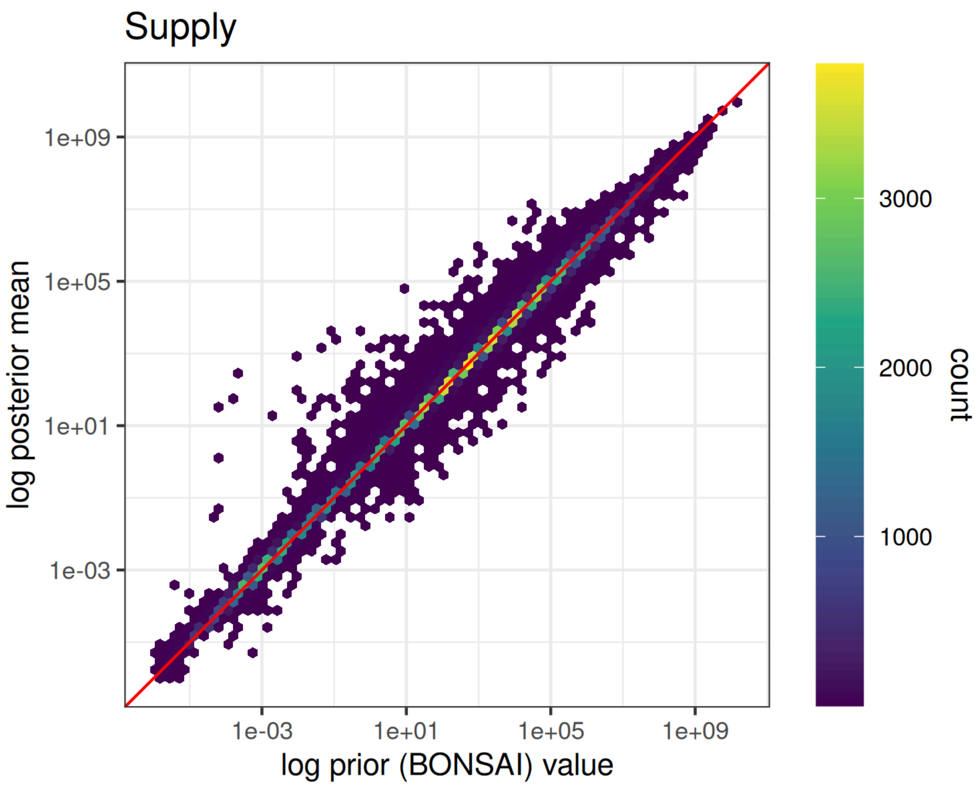

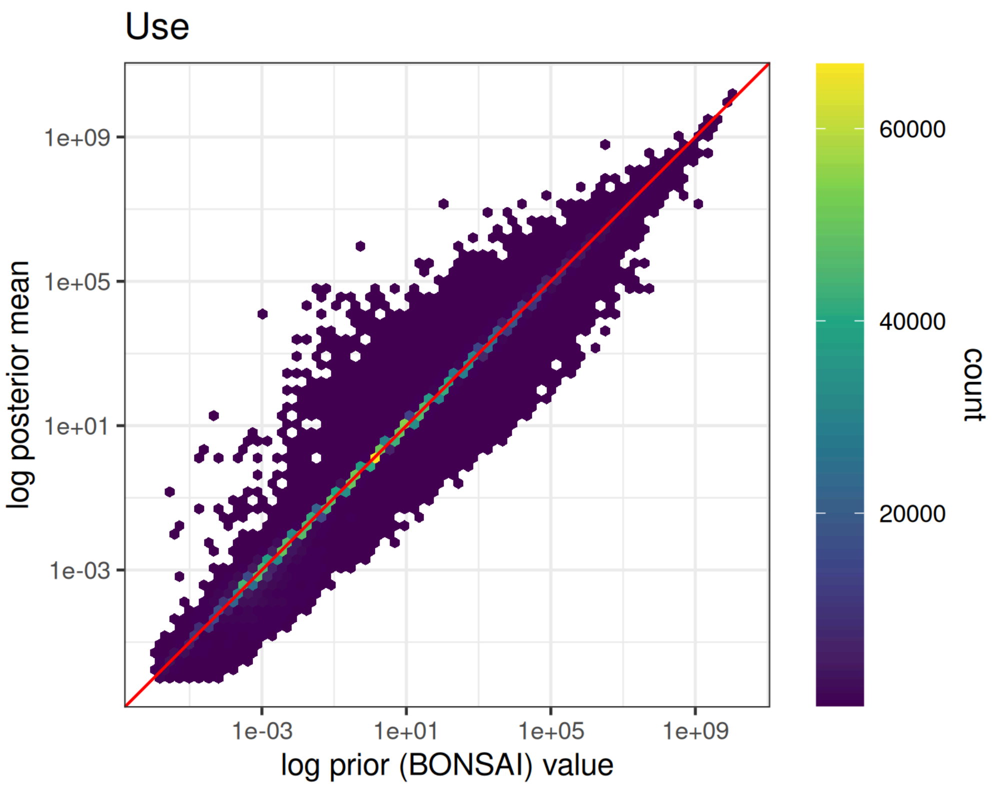

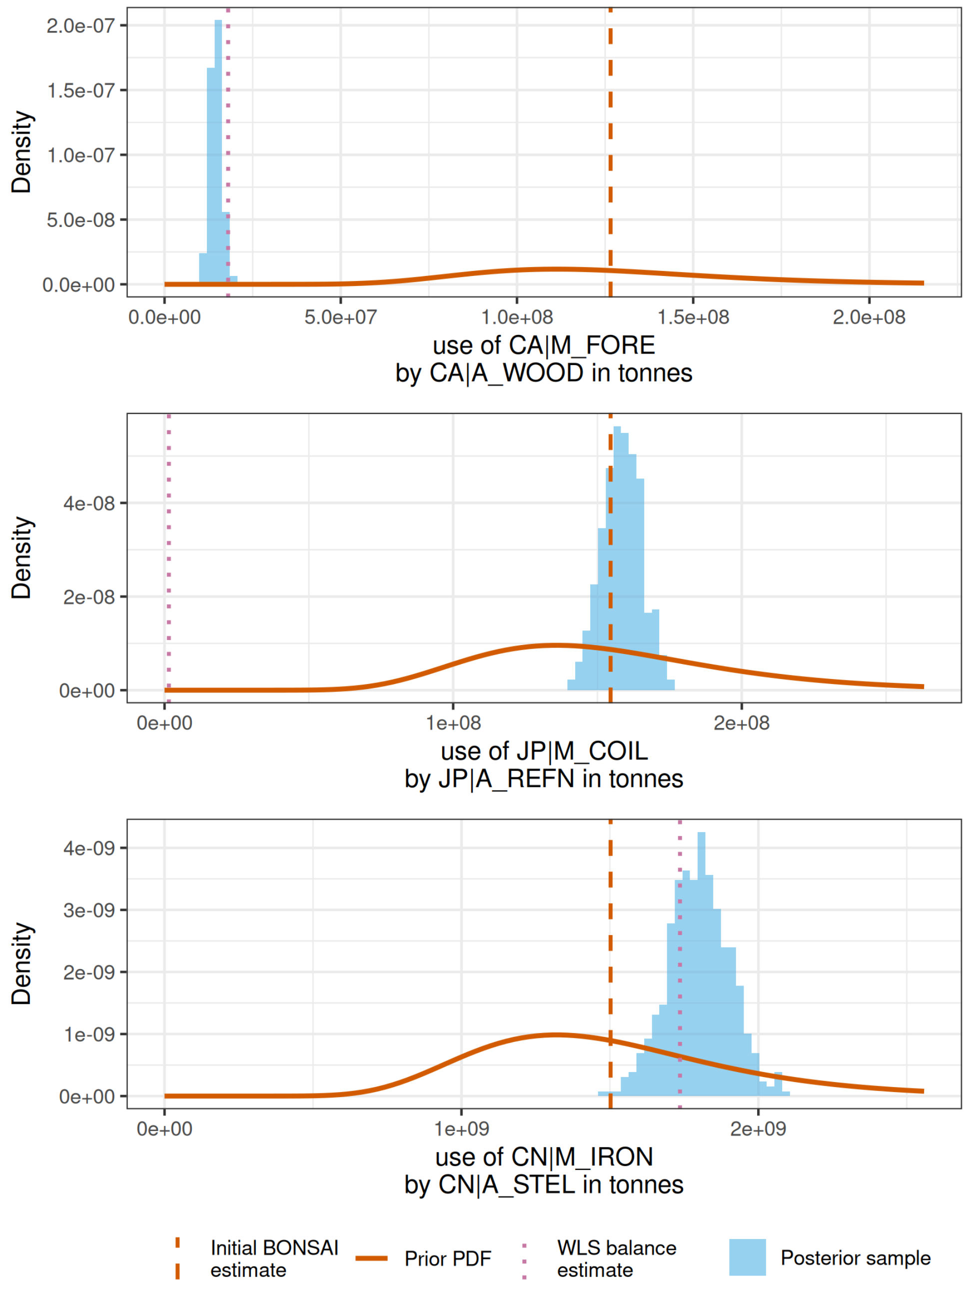
Density
2.0e-07
1.5e-07
1.0e-07
5.0e-08
0.0e+00
0.0e+00
5.0e+07
1.0e+08
1.5e+08
2.0e+08
use of CA|M_FORE
by CA|A_WOOD in tonnes
Density
4e-08
2e-08
0e+00
0e+00
1e+08
2e+08
use of JP|M_COIL
by JP|A_REFN in tonnes
Density
4e-09
3e-09
2e-09
1e-09
0e+00
0e+00
1e+09
2e+09
use of CN|M_IRON
by CN|A_STEL in tonnes
Initial BONSAI estimate
Prior PDF
WLS balance estimate
Posterior sample

2) uncertainty changes (selectively)

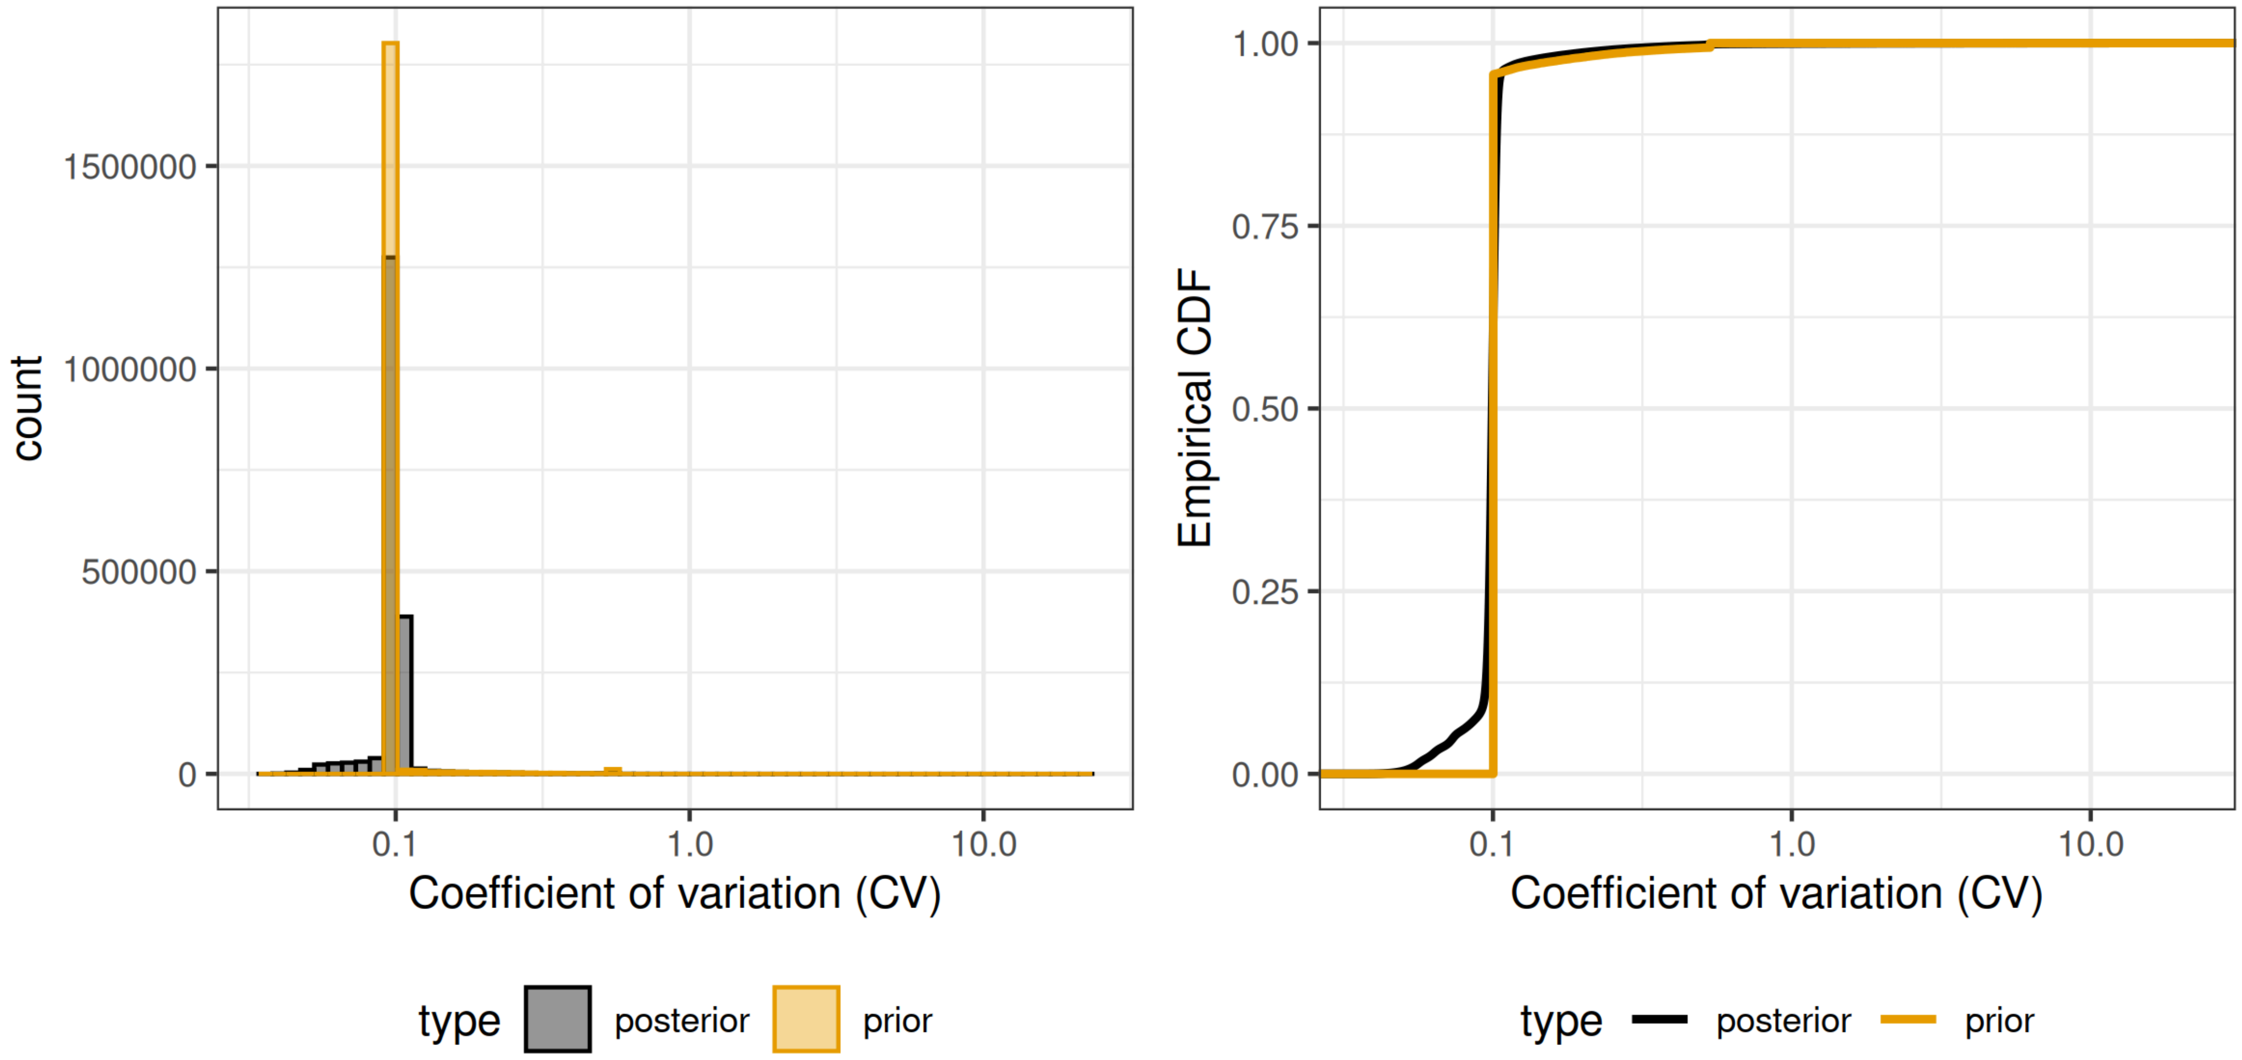

3) Before & After balancing

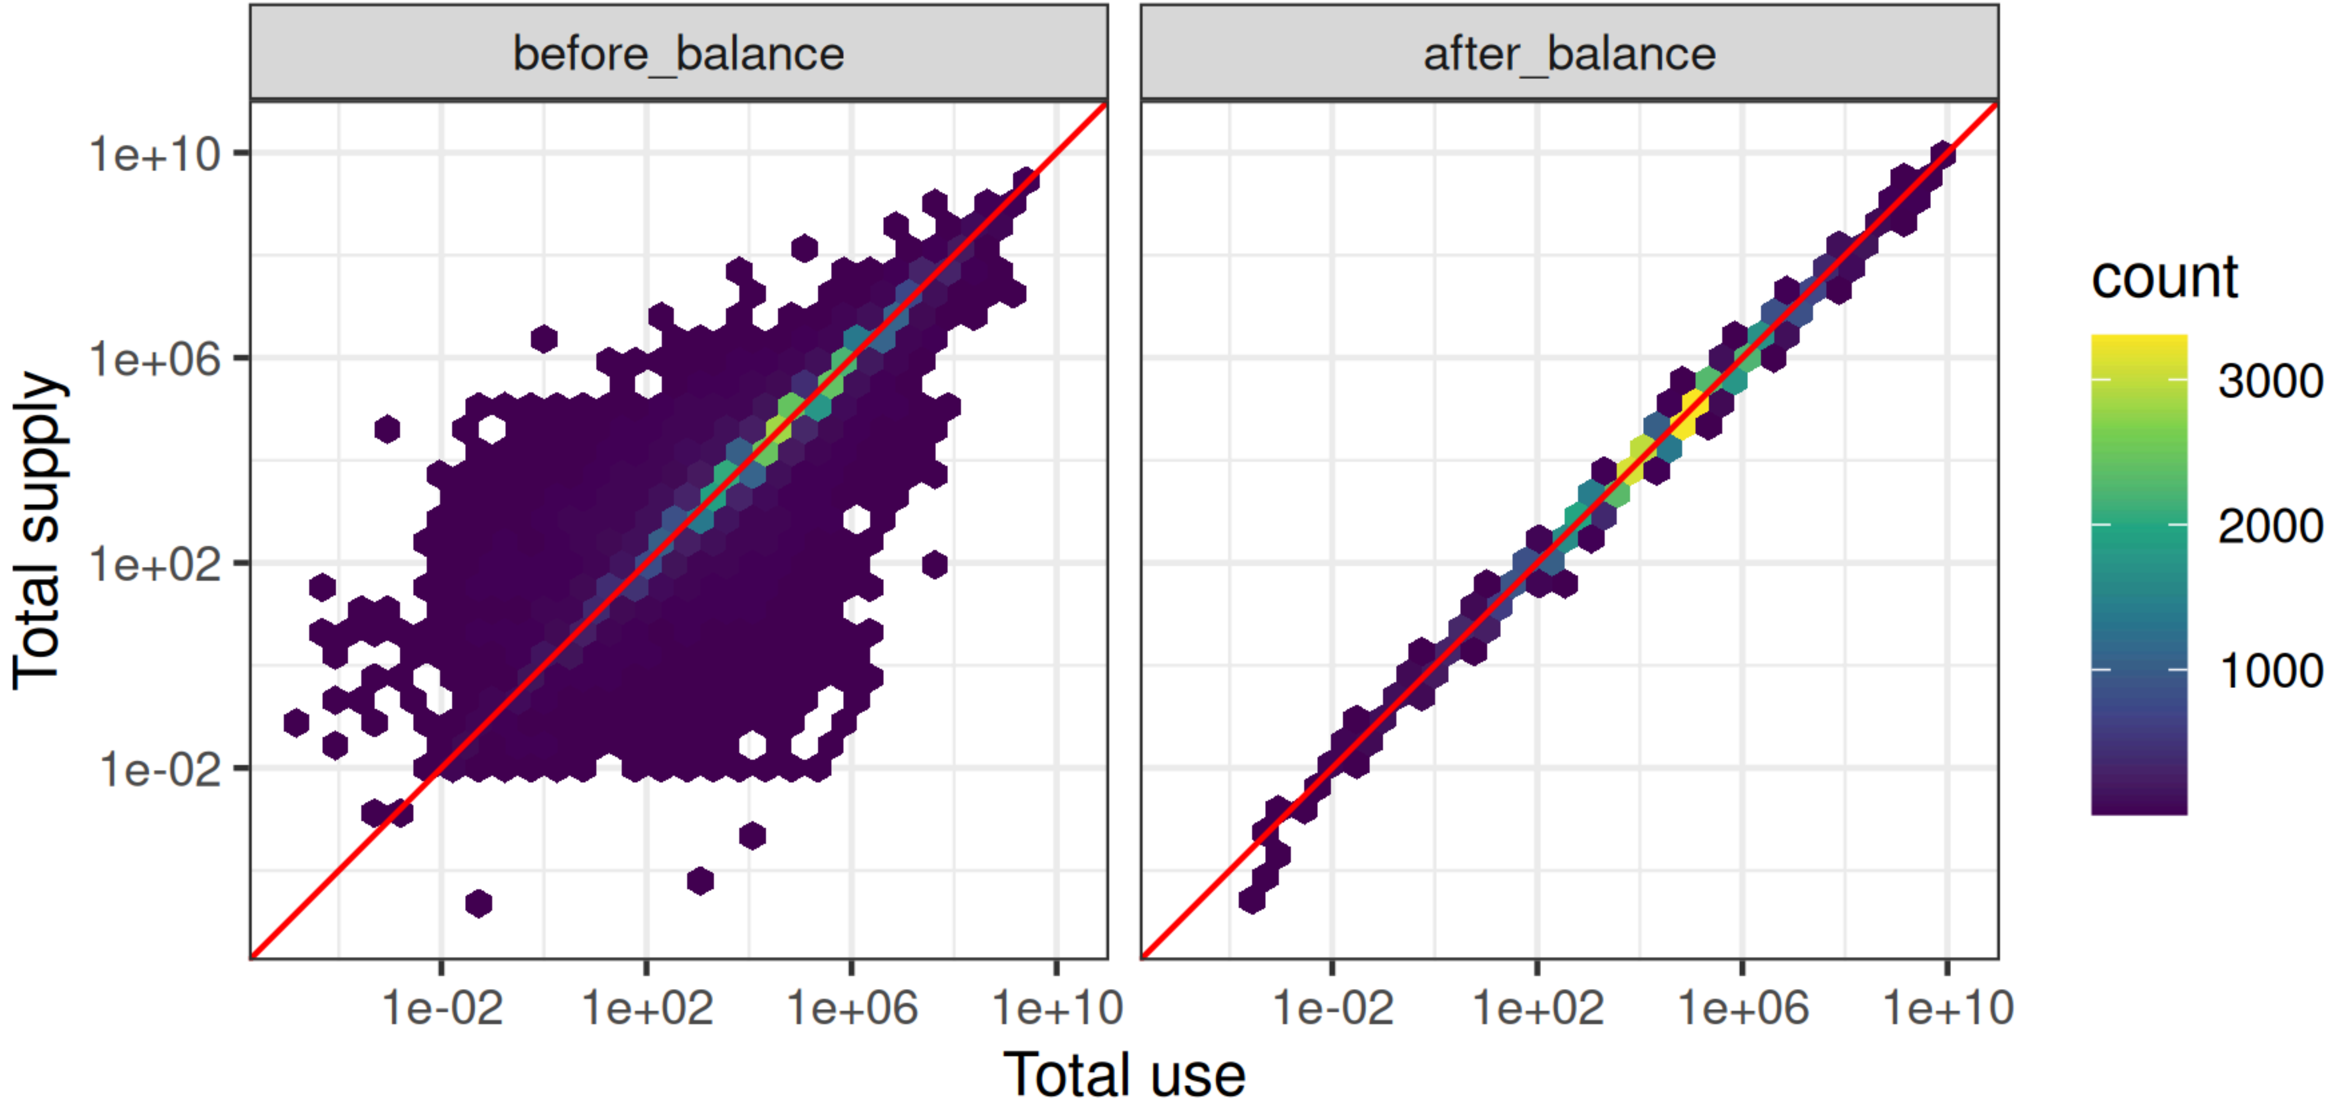

Product constraint [tonnes]

4) posterior covariance

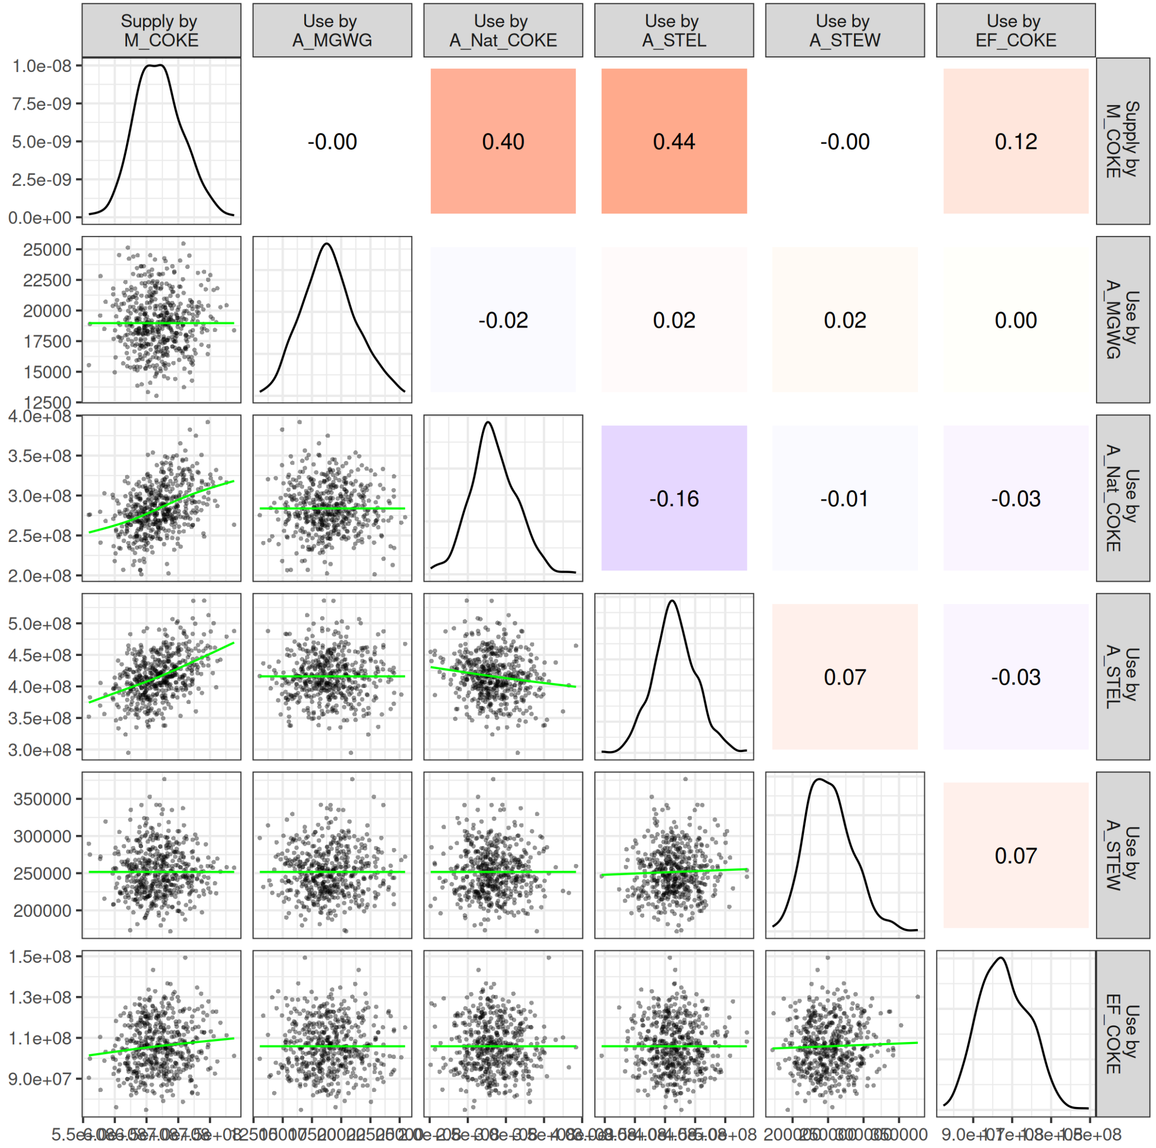

- M_COKE: Chinese coke market
- A_Nat_COKE: natural coke production
- A_STEL: steel manufacturing

> **Note**
>
> Due to use of soft constraints that still allow for deviation from the constraints, we expect stronger dependencies when tightening the allowed log-ratio deviation from the constraints, which are currently set relatively moderate rel_tol_* between 0.05 and 0.2

Limitations of the current prototype

- prior variances largely data-driven based on prior imbalances (due to computational reasons)
- global (coarse) uncertainty parameters for some linkages (conversion between layers, transfer coefficients)
- soft constraints require care (tolerances matter)
- computational optimisation: post-processing/memory pipeline, MCMC

Outlook

- Empirically grounded prior elicitation (source/flow specific)
- tighter tolerances
- testing Bayesian approximations (e.g. Pathfinder, Variational Inference)
- generalizations: zero handling, time-series (temporal consistency & information transfer)
- create a usable tool

From:

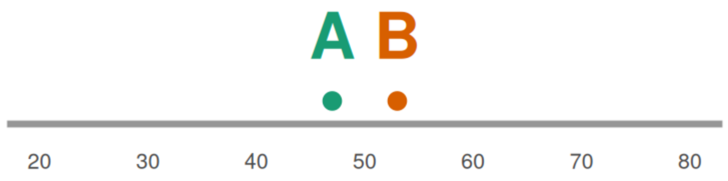

$A \quad \text{or} \quad B?$

Point estimates hide uncertainty and dependencies.

To:

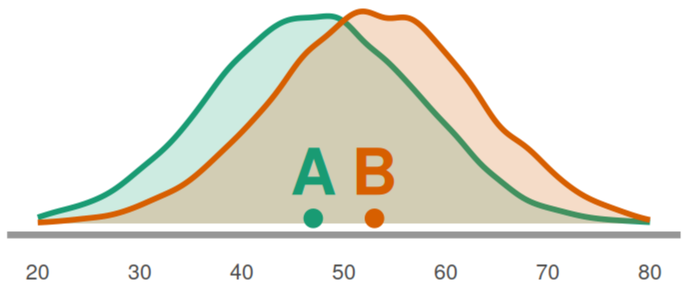

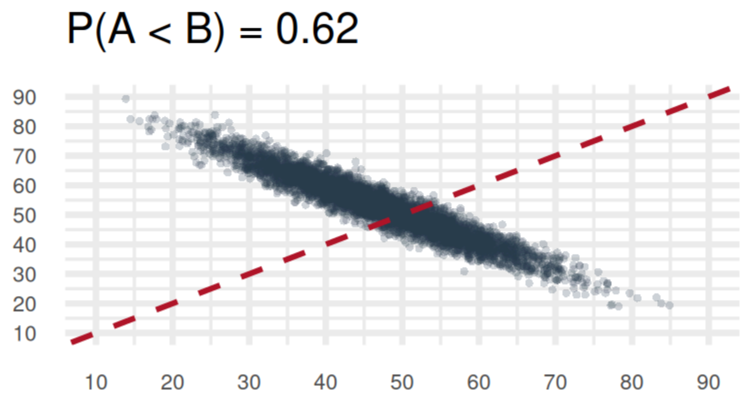

Bayesian balancing provides

- a posterior ensemble of balanced SUTs
- uncertainty intervals and covariances
- diagnostics of data–constraint tensions
- direct propagation to decision probabilities:

$$P(\text{footprint}_A < \text{footprint}_B \mid \text{data, constraints})$$

Thank you

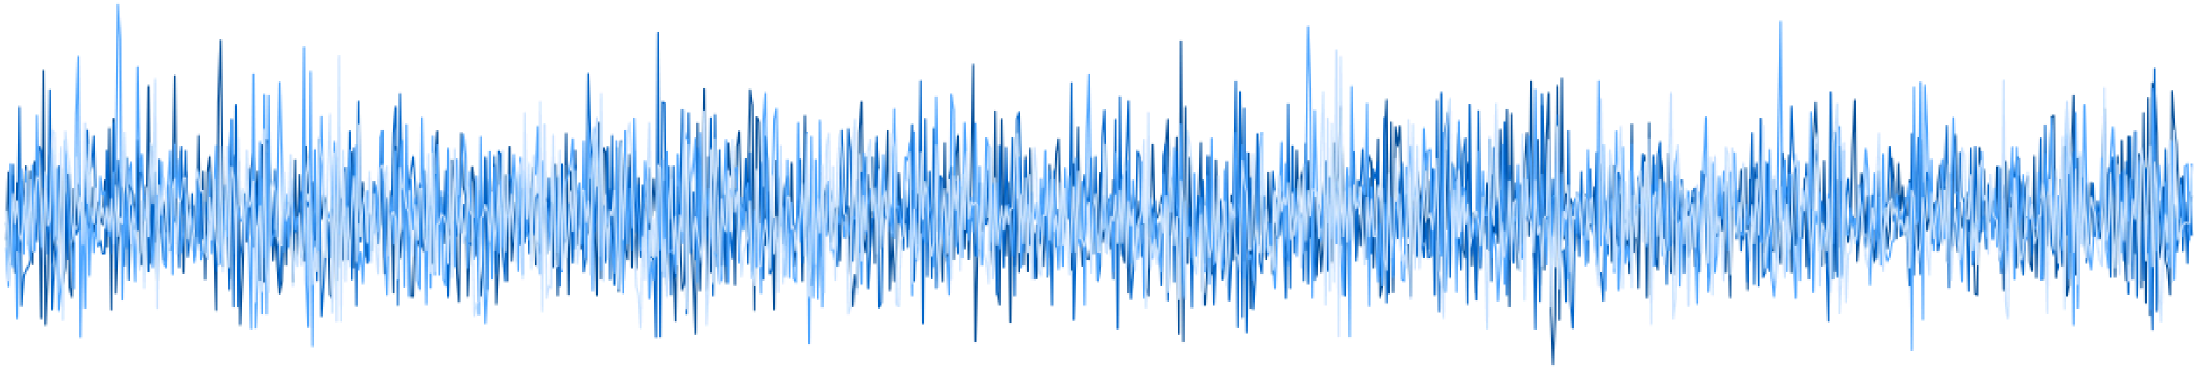

Contact: simon.schulte@indecol.uni-freiburg.de

👈 Slides:
https://simschul.github.io/2026_IIOA_conference_Sevilla/

Working paper: https://freidok.uni-freiburg.de/data/279579 👉

Discussion

1. What tolerances (=deviations from constraints) would you consider acceptable for the respective constraints?
2. What kind of uncertainty info can statistical offices provide (to be encoded as priors)?
3. IO practitioners:
    - What would you need to make use of such an uncertainty-aware SUT? (data: full samples, covariances, …; tools: integration in pymrio/MARIO? stand-alone library?)
    - Which challenges do you see for practitioners?

Additional slides

The model

Model variables (1)

Supply & Use flows:

$$\log x_i^{(\mathrm{sup},l)} \sim \mathcal{N}\left(\mu_i^{(\mathrm{sup},l)},\ \sigma_i^{(\mathrm{sup},l)\ 2}\right)$$

$$\log x_i^{(\mathrm{use},l)} \sim \mathcal{N}\left(\mu_i^{(\mathrm{use},l)},\ \sigma_i^{(\mathrm{use},l)\ 2}\right)$$

Conversion factors:

$$x_i^{(cf,l_1 \to l_2)} \sim \mathcal{N}(y_i^{(cf,l_1 \to l_2)},\ \sigma_i^{(cf,l_1 \to l_2) 2})$$

Model variables (2)

Transfer coefficient:

$$x_i^{(\mathrm{tc})} \sim \mathrm{Beta}(\alpha_i^{(\mathrm{tc})},\ \beta_i^{(\mathrm{tc})})$$

$$y_i^{(\mathrm{tc})} \sim \mathcal{N}(x_i^{(\mathrm{tc})},\ \sigma_i^{(\mathrm{tc})2})$$

Prior variances (Supply & Use flows) (1)

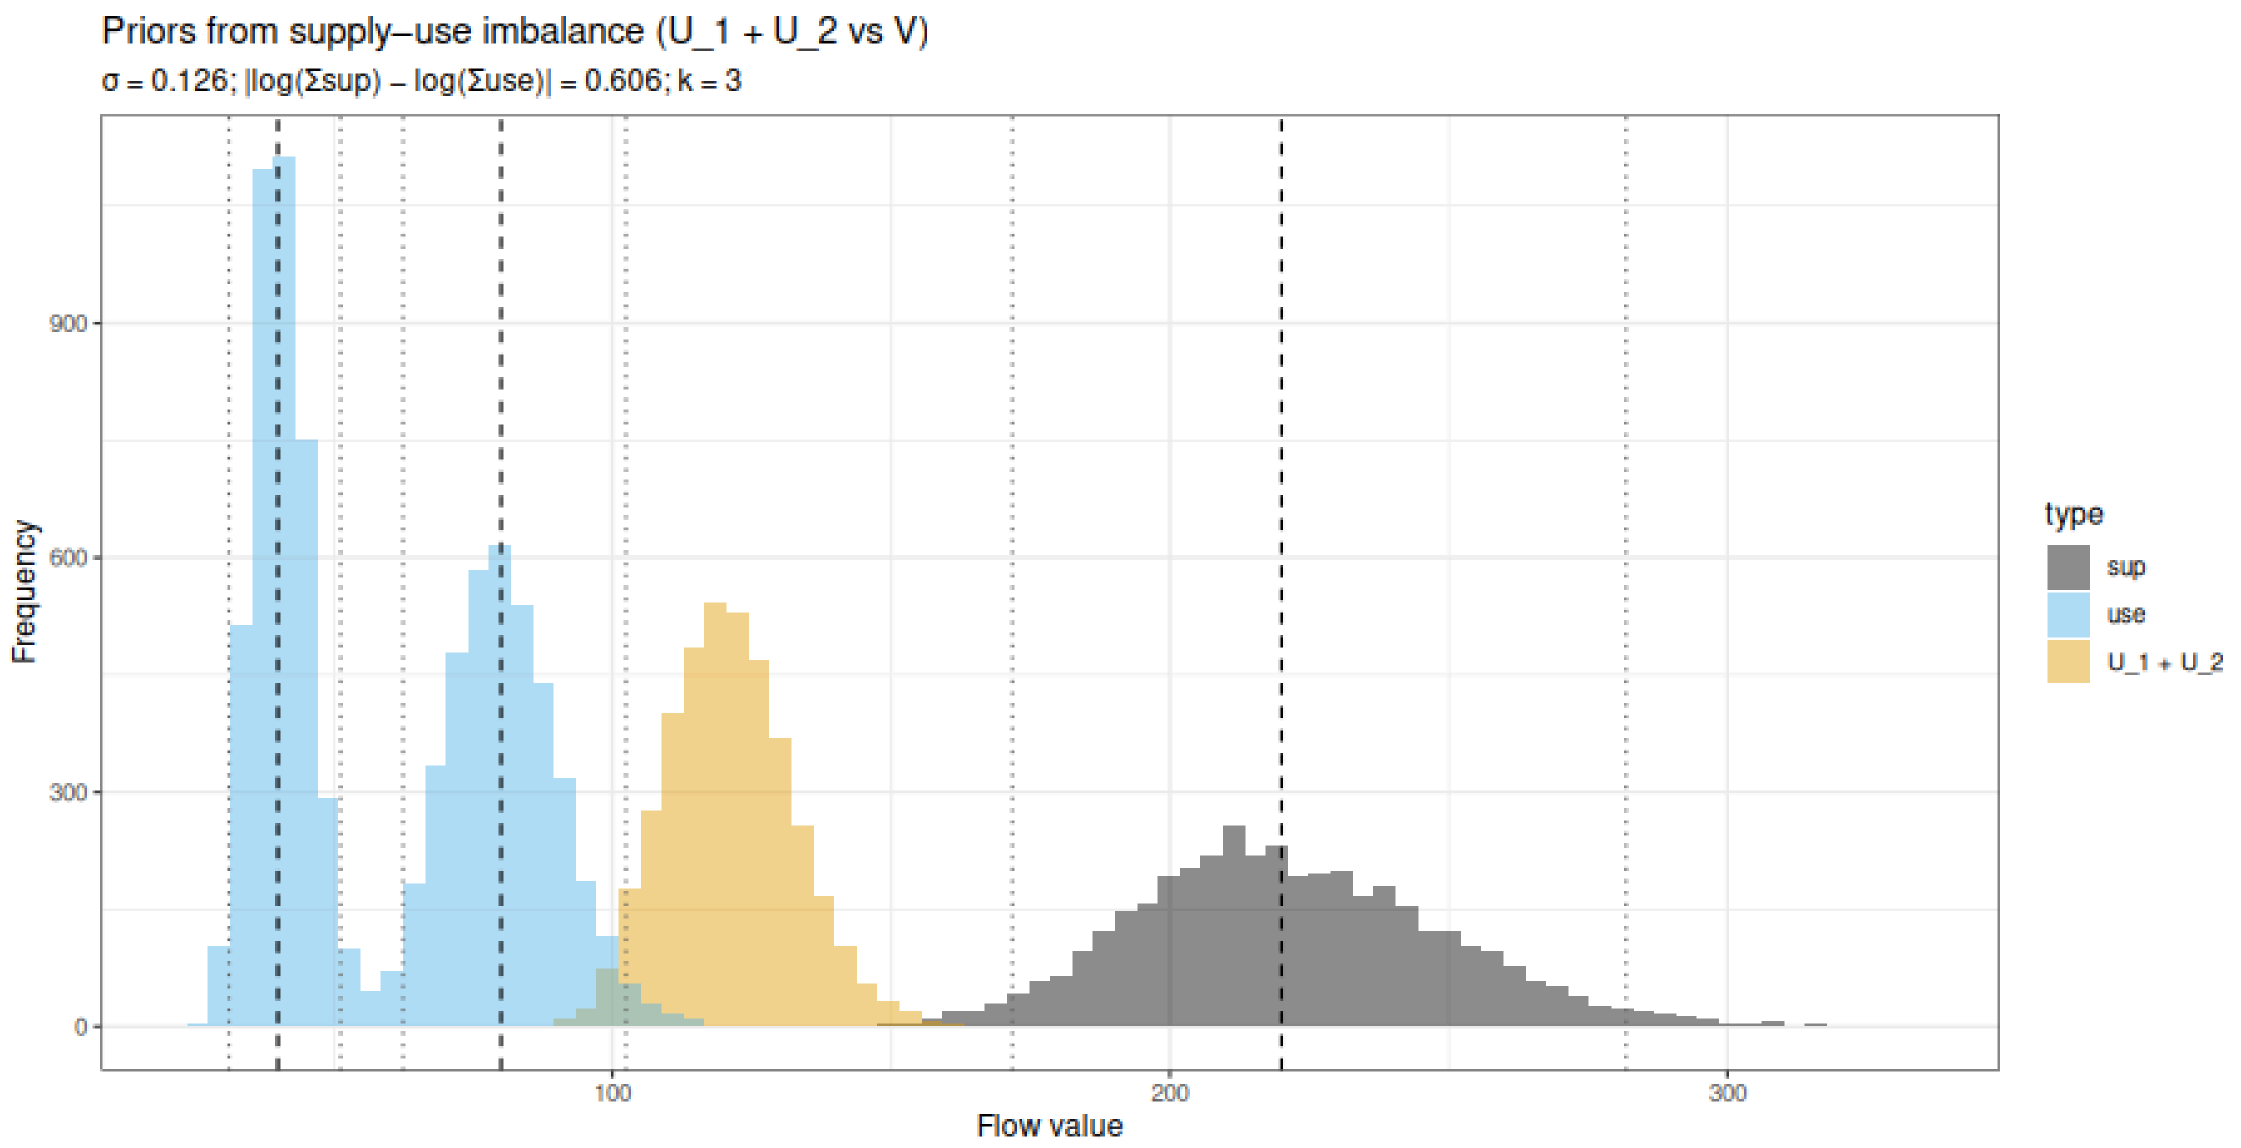

Prior variances (Supply & Use flows) (2)

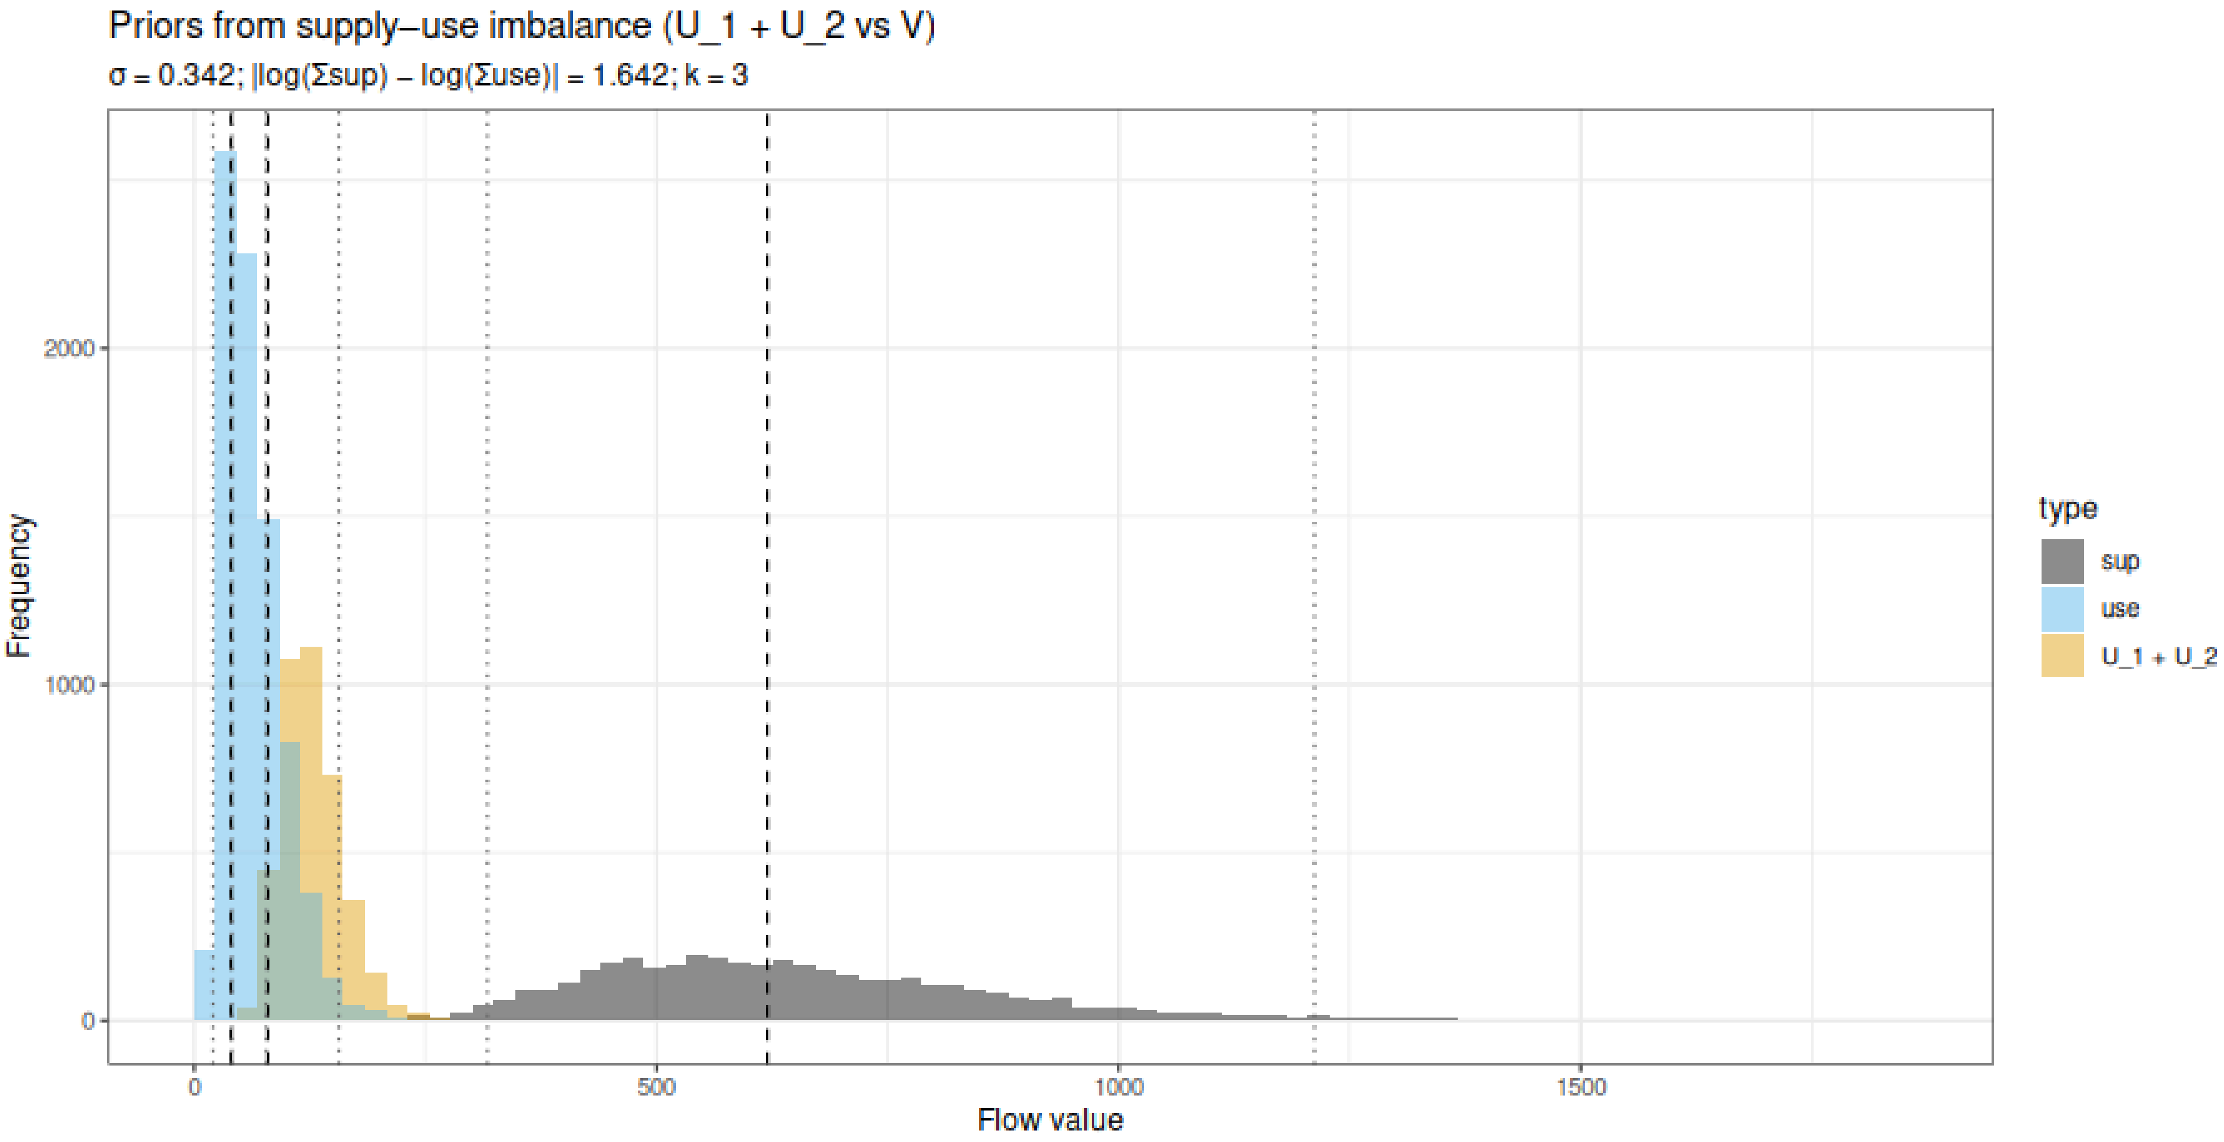

Prior variances (Supply & Use flows) (3)

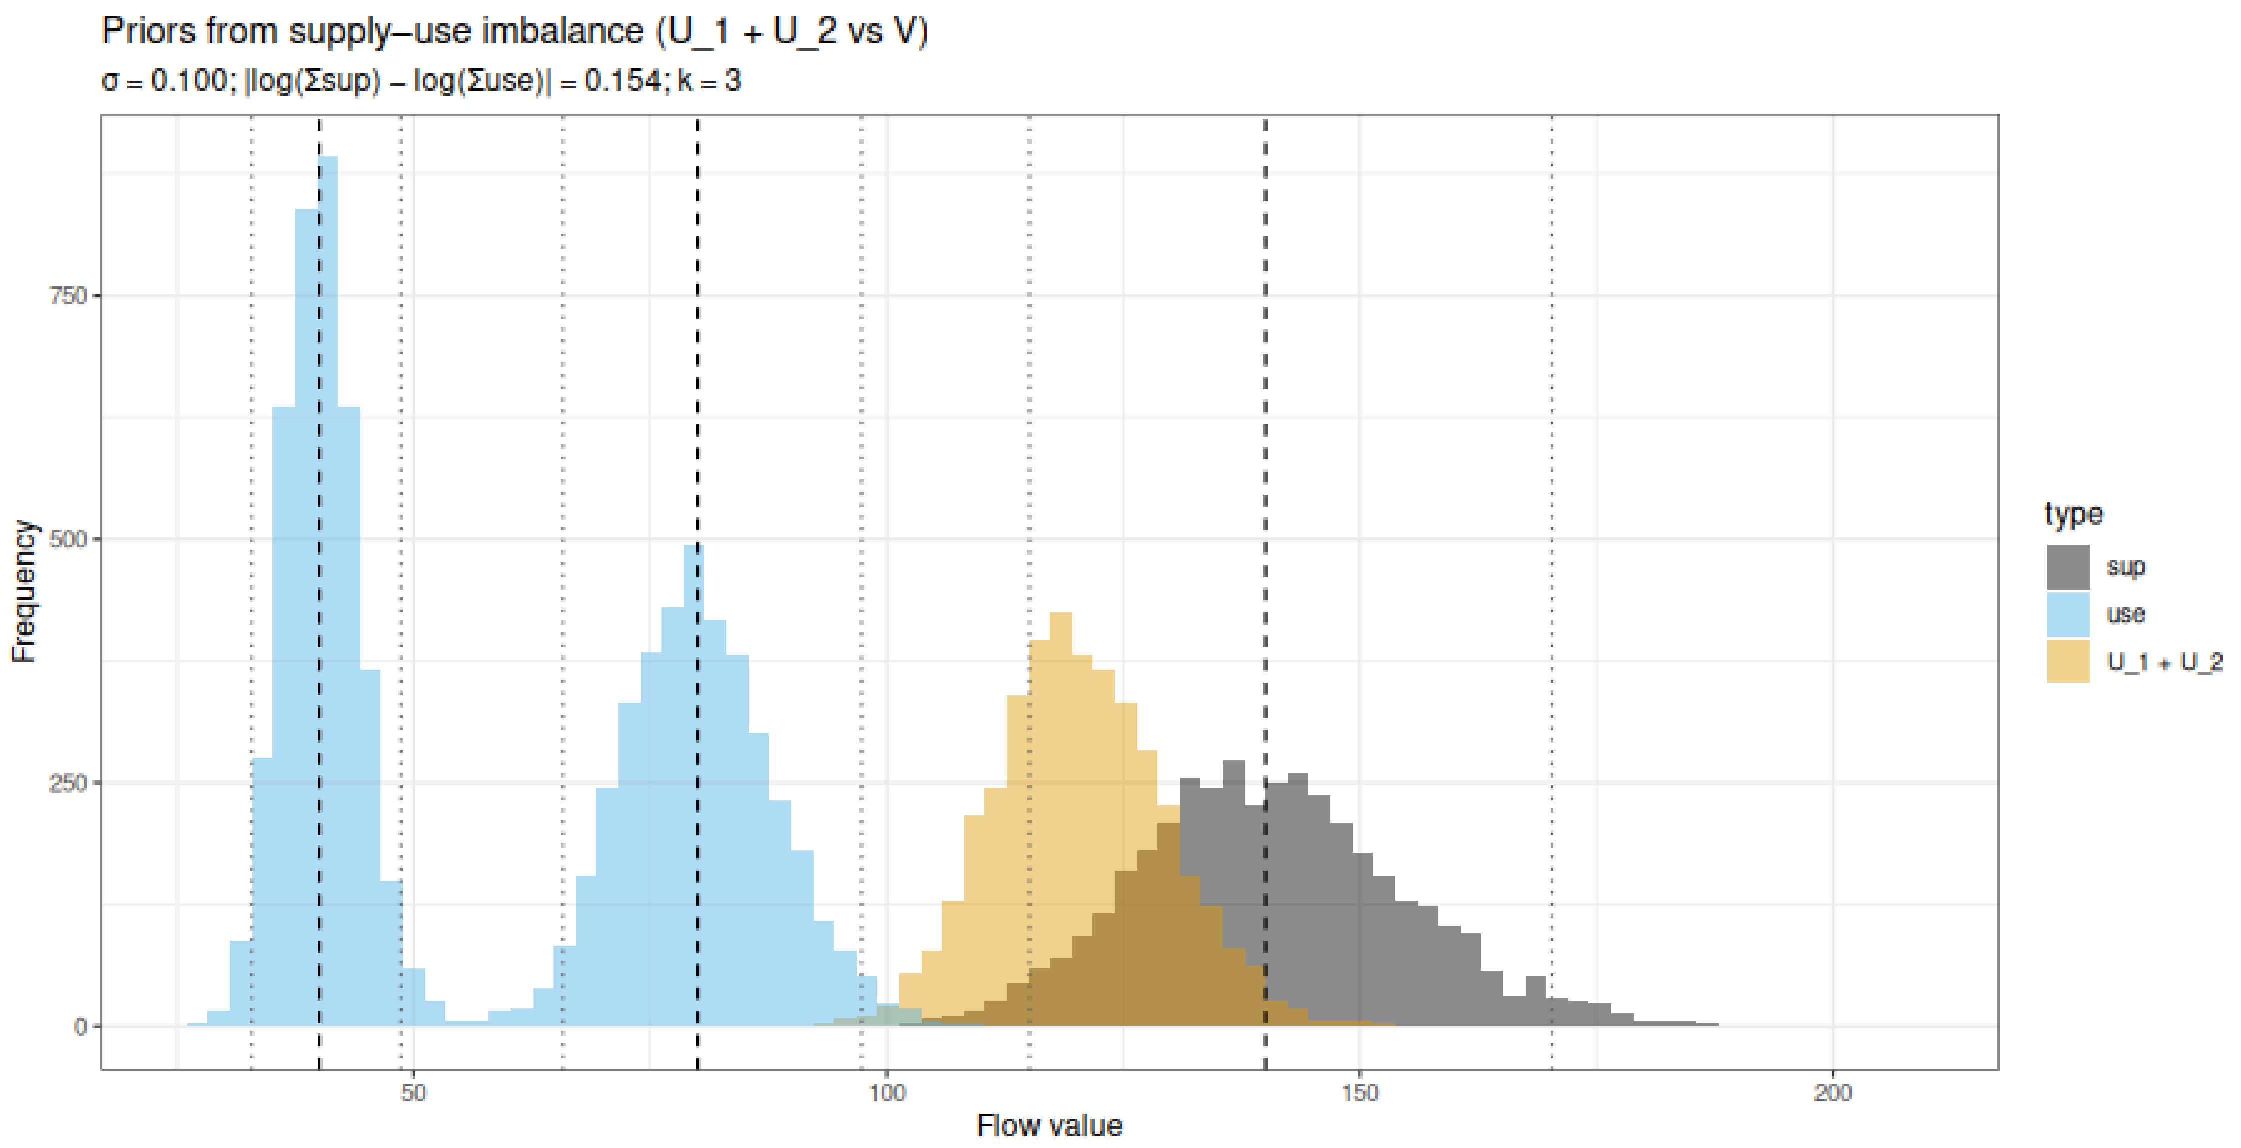

1. **Quantify the Imbalance:** We compute the log-difference Δ for each accounting balance (product or activity) at the product-region level (pr) and activity region level (ar), respectively:

$$\Delta_{\mathrm{pr},j}^{(l)} = \log(s_{\mathrm{pr},j}^{(l)} + \epsilon) - \log(u_{\mathrm{pr},j}^{(l)} + \epsilon) \tag{20}$$

$$\Delta_{\mathrm{ar},k}^{(l)} = \log(i_{\mathrm{ar},k}^{(l)} + \epsilon) - \log(o_{\mathrm{ar},k}^{(l)} + \epsilon), \tag{21}$$

where $s_{\mathrm{pr},j}^{(l)}$ and $u_{\mathrm{pr},j}^{(l)}$ is the total supply/use by product-region j (e.g. coal from Denmark) and property layer l; and $i_{\mathrm{ar},k}^{(l)}$ and $o_{\mathrm{ar},k}^{(l)}$ is the total input/output by activity-region k (e.g. coke production in Germany) and property layer l.
2. **Derive the Noise Level:** Next, to each of these observed imbalances $\Delta_{\mathrm{pr},j}^{(l)}$ and $\Delta_{\mathrm{ar},k}^{(l)}$ there are $m_{\mathrm{pr},j}^{(l)}/m_{\mathrm{ar},k}^{(l)}$ contributing flows. For instance, to the imbalance $\Delta_{\mathrm{pr,German\_steel}}^{(\mathrm{tonnes})}$ the contributing flows are "use of steel by German car manufacturing", "use of steel by German machinery manufacturing" and "supply of steel by German steel manufacturing", amongst others. We assume that these observed imbalances $\Delta_{(\cdot)}$ each sit at the edge of a 95% confidence interval generated by the sum of m noisy rows. This allows us to back-calculate the required σ's for each contributing flow:

$$\sigma_i^{(\mathrm{pr},l)} \approx \frac{|\Delta_{\mathrm{pr},j}^{(l)}|}{1.96 \cdot \sqrt{2} \cdot \sqrt{m_j}}, \tag{22}$$

$$\sigma_i^{(\mathrm{ar},l)} \approx \frac{|\Delta_{\mathrm{ar},k}^{(l)}|}{1.96 \cdot \sqrt{2} \cdot \sqrt{m_k}}, \tag{23}$$

where $\sigma_i^{(\mathrm{pr},l)}/\sigma_i^{(\mathrm{ar},l)}$ are the flow-specific variances.
Here, the $\sqrt{m}$ term accounts for the Central Limit Theorem: as more flows contribute to a balance, individual flow uncertainties contribute less to the aggregate uncertainty.
3. For all flows that are subject to more than one constraint, take the largest relevant $\sigma_i = \max\{\sigma_i^{(\mathrm{pr},l)}, \sigma_i^{(\mathrm{ar},l)}\}$ (most conservative) and pass it to the model as the log-scale prior for that flow.
4. **Clamping:** To prevent the priors from becoming physically unrealistic (too wide) or numerically unstable (too narrow), the derived σ_i is clamped to a conservative range $[sd_{\min}, sd_{\max}]$.

Product & Activity constraints

Product constraint:

$$\log \frac{s_{\mathrm{pr},j}^{(l)} + \epsilon}{l_{\mathrm{pr},j}^{(l)} + \epsilon} \sim \mathcal{N}(0,\ \mathrm{rel\_tol}_{\mathrm{product}}^2)$$

Activity constraint:

$$\log \frac{o_{\mathrm{ar},k}^{(l)} + \epsilon}{i_{\mathrm{ar},k}^{(l)} + \epsilon} \sim \mathcal{N}(0,\ \mathrm{rel\_tol}_{\mathrm{activity}}^2)$$

Conversion factor consistency

$$\mathrm{converted}_i^{(\mathrm{sup},l_2)} = x_i^{(\mathrm{sup},l_1)} \cdot x_i^{(cf,l_1\to l_2)},$$

$$\mathrm{reference}_i^{(\mathrm{sup},l_2)} = x_i^{(\mathrm{sup},l_2)}$$

(and analogously for use flows $x^{(\mathrm{use},l)}$).

Soft log-ratio constraint:

$$\log \frac{\mathrm{converted}_i^{(\cdot)} + \epsilon}{\mathrm{reference}_i^{(\cdot)} + \epsilon} \sim \mathcal{N}\left(0,\ \mathrm{rel\_tol}^2_{\mathrm{conv\_fac}}\right)$$

Transfer coefficient consistency

- First, we have to aggregate supply and use to the product-activity-region (PAR) level (e.g.~total use of steel in car manufacturing in Germany, or supply of cars from car manufacturing in Germany) using aggregation matrices $\mathbf{A}_{par}^{(\cdot)}$:

$$\mathbf{x}_{\text{par}}^{(\text{use},\text{tonnes})} = \mathbf{A}_{\text{par}}^{(\text{use},\text{tonnes})} \mathbf{x}^{(\text{use},\text{tonnes})} ; \mathbf{x}_{\text{par}}^{(\text{sup},\text{tonnes})} = \mathbf{A}_{\text{par}}^{(\text{sup},\text{tonnes})} \mathbf{x}^{(\text{sup},\text{tonnes})}$$

- Second, we compute the expected supply from the transfer coefficients by multiplying the transfer coefficient record with the respective use-flow at the PAR-level:

$$x_{\text{par},i}^{(\text{exp\_sup},\text{tonnes})} = x_i^{(\text{tc})} \cdot x_{\text{par},i}^{(\text{use},\text{tonnes})}$$

- Third, we again apply a soft log-ratio constraint so that the observed PAR-level supplies in tonnes must approximately match expected supplies implied via transfer coefficients:

$$\log \frac{x_{\text{par},i}^{(\text{exp\_sup},\text{tonnes})} + \epsilon}{x_{\text{par},i}^{(\text{sup},\text{tonnes})} + \epsilon} \sim \mathcal{N}\left(0, \text{rel\_tol}_{\text{tc}}^2\right)$$

Before/After balancing (product balance, Mill. Euro)

Before/After balancing (product balance, TJ)

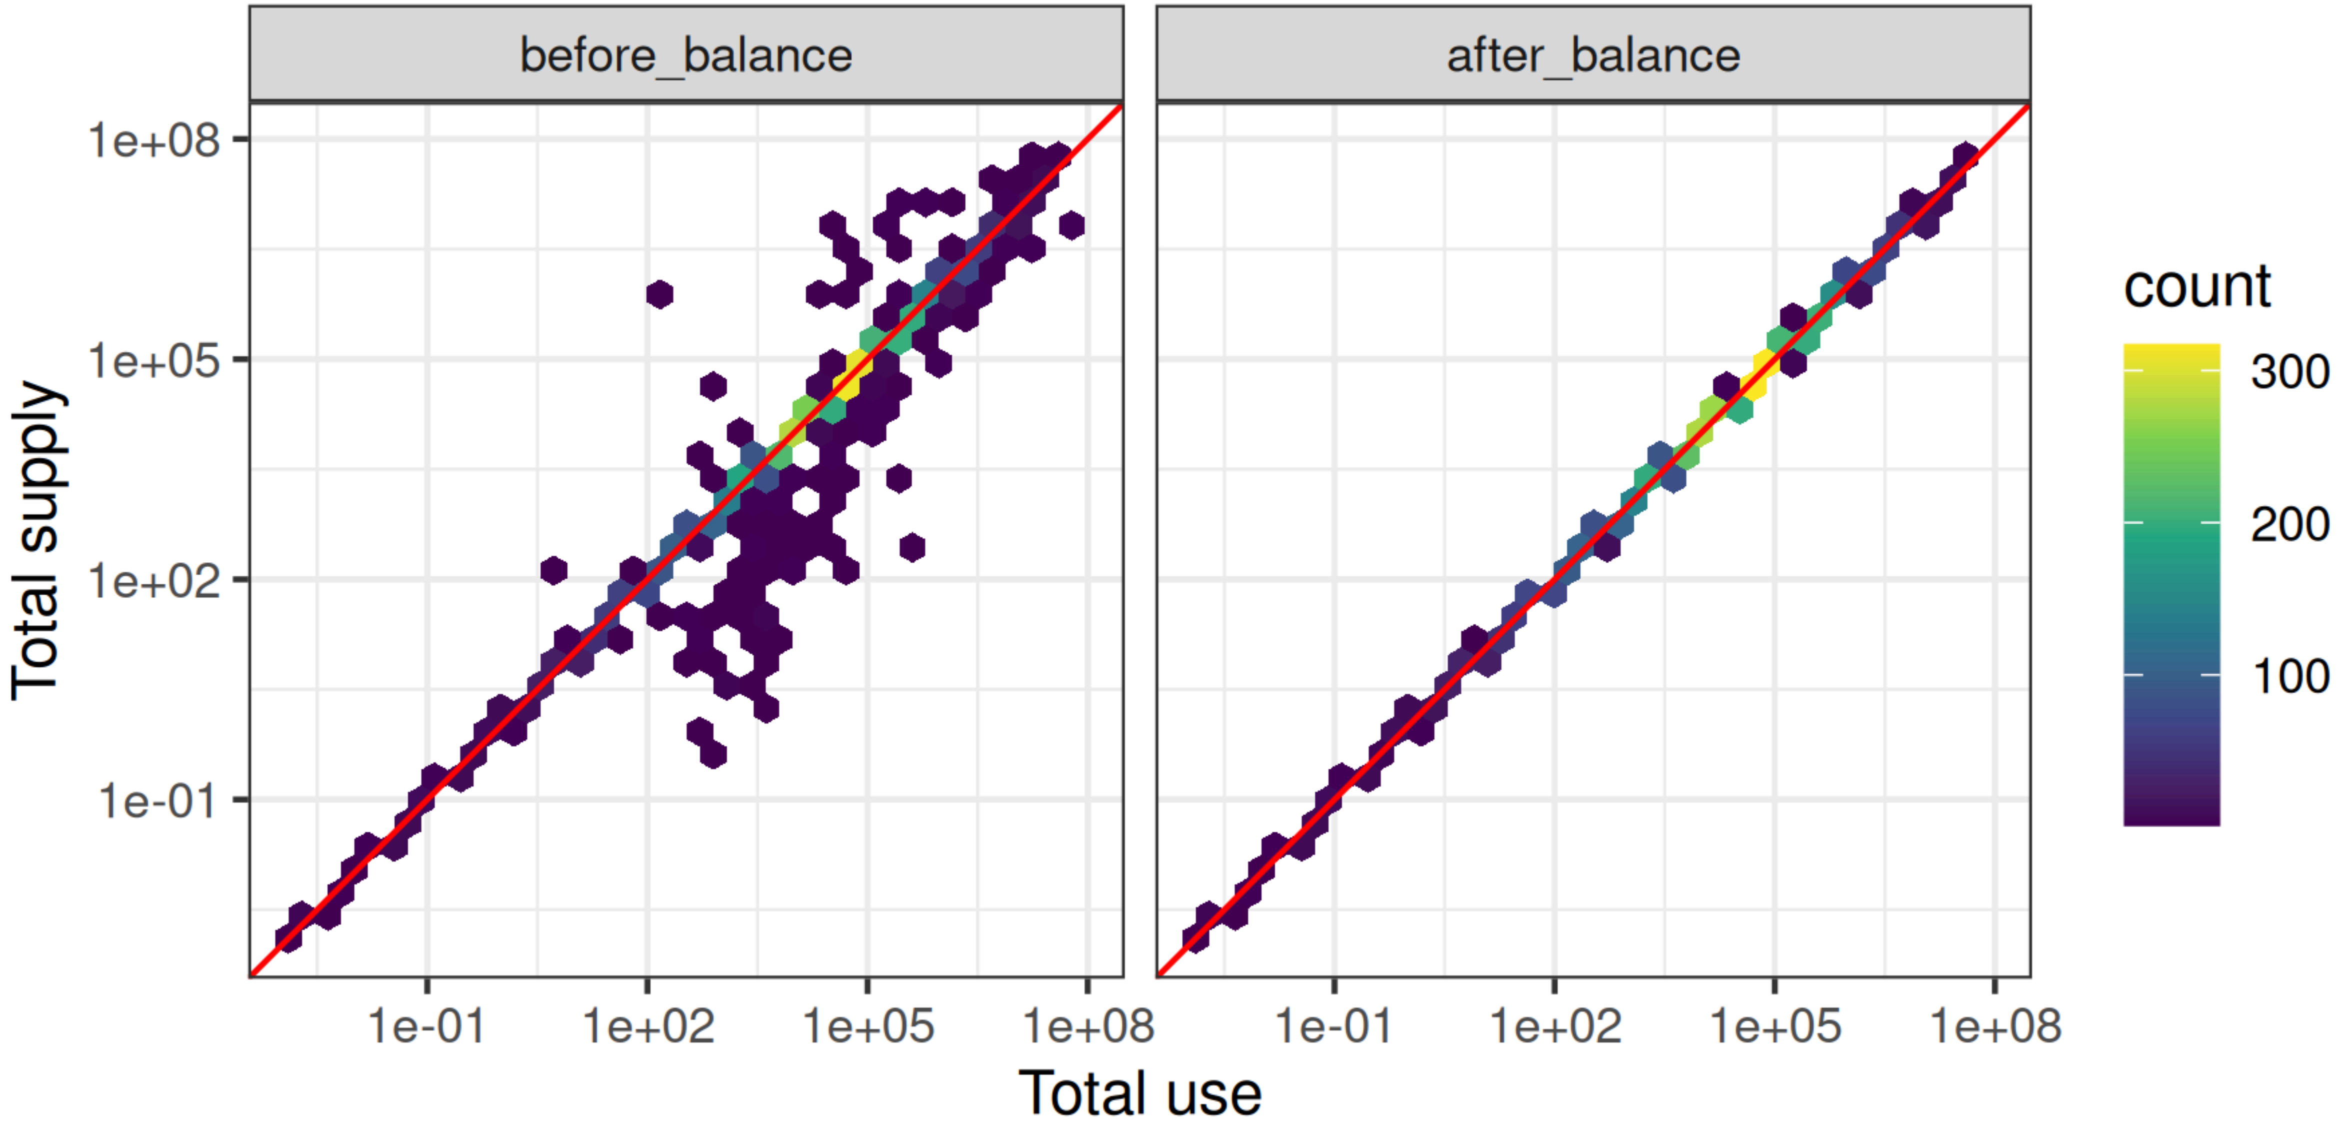

Before/After balancing (activity balance, tonnes)

Before/After balancing (activity balance, Mill. Euro)

Before/After balancing (activity balance, TJ)

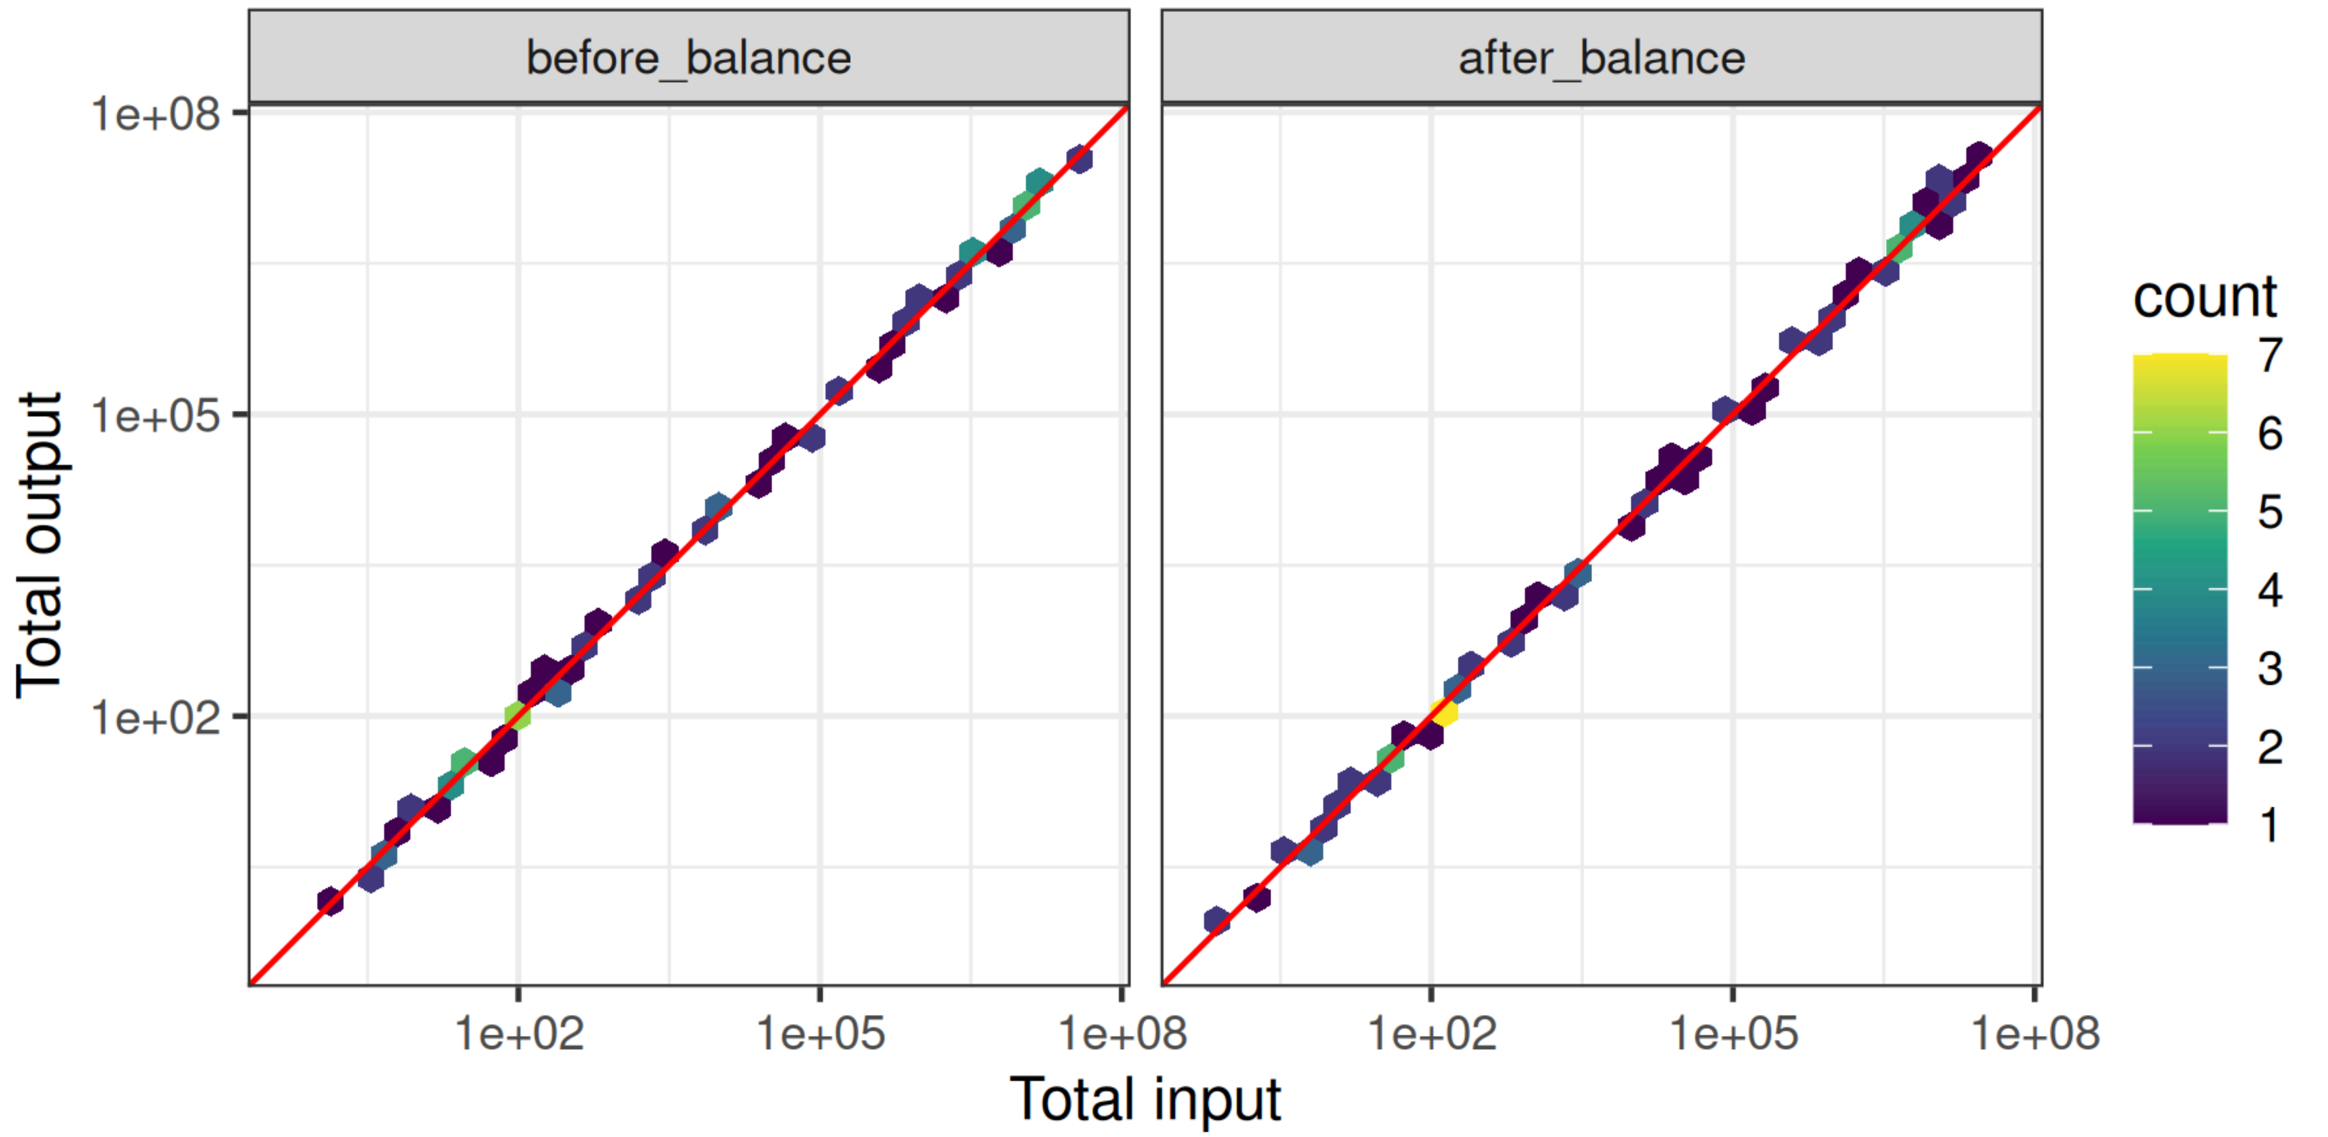

Comparison with deterministic WLS

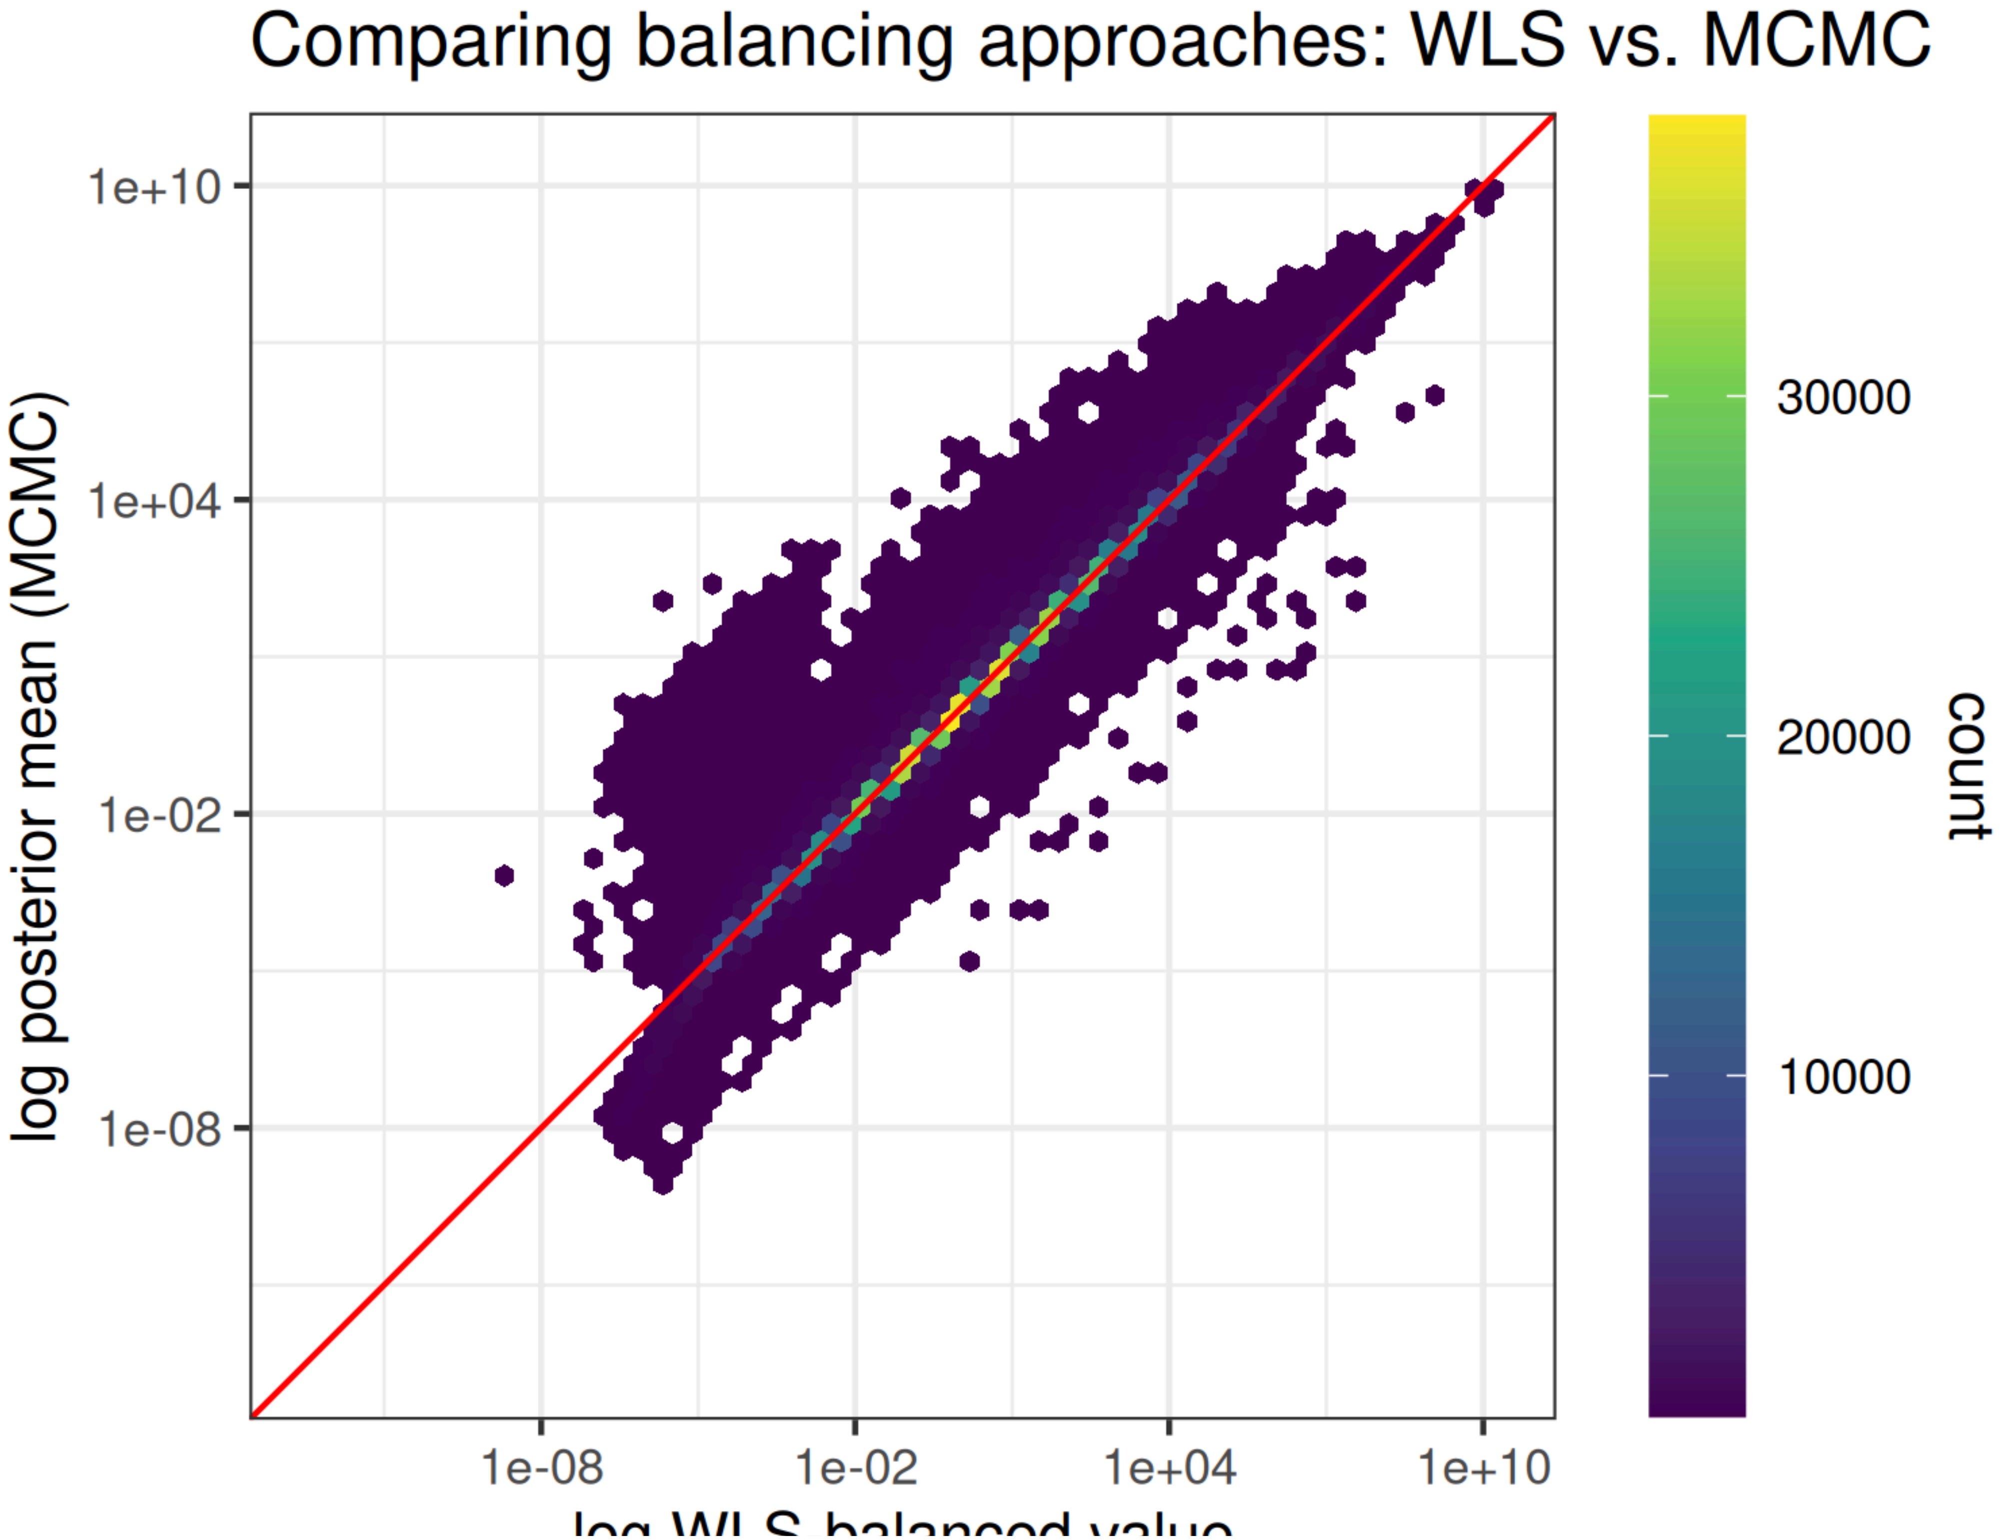
Comparing balancing approaches: WLS vs. MCMC
log posterior mean (MCMC)
1e+10
1e+04
1e-02
1e-08
1e-08
1e-02
1e+04
1e+10
count
30000
20000
10000