Bayesian, Fiducial, and Frequentist (BFF): Best Friends Forever?

Xiao-Li Meng

Department of Statistics, Harvard University

- Liu & Meng (2106) **There Is Individualized Treatment. Why Not Individualized Inference?** *Annual Review of Statistics and Its Application*, 3: 79-111
- Liu & Meng (2014). **A Fruitful Resolution To Simpson's Paradox via Multi-Resolution Inference.** *The American Statistician*, 68: 17-29.
- Meng (2014). **A Trio of Inference Problems That Could Win You a Nobel Prize in Statistics (if you help fund it)**. *In the Past, Present, and Future of Statistical Science (Eds: X. Lin, et. al.)*, 535-560.

What is *inference*? Katie's answer ...

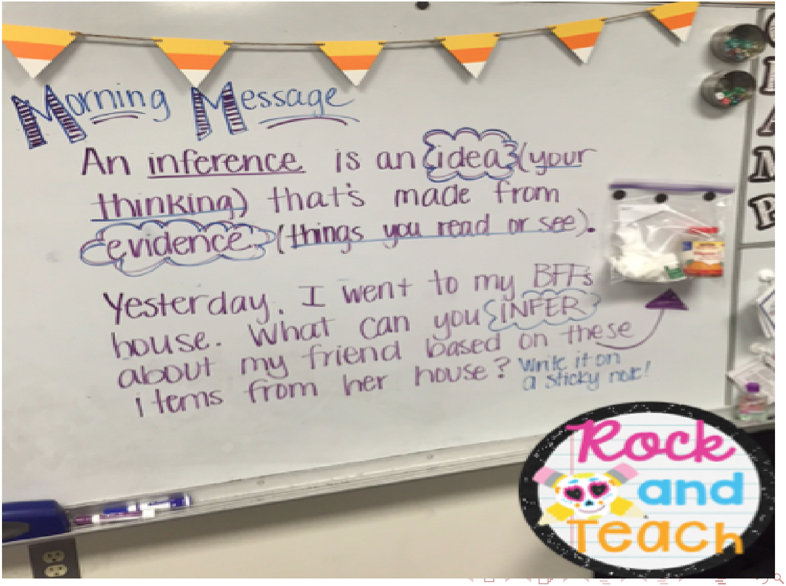

But what is *Statistical/Probabilistic* Inference?

- An ultimate intellectual game: **"to guess wisely and to guess meaningfully the errors in our guesses."** (*XL-Files*, Oct 2015)

But what is *Statistical/Probabilistic* Inference?

- An ultimate intellectual game: **"to guess wisely and to guess meaningfully the errors in our guesses."** (*XL-Files*, Oct 2015)
- Impossible to access exact errors, but **a full spectrum** of possibilities for accessing **probabilistic errors**.

But what is *Statistical/Probabilistic* Inference?

- An ultimate intellectual game: **"to guess wisely and to guess meaningfully the errors in our guesses."** (*XL-Files*, Oct 2015)
- Impossible to access exact errors, but **a full spectrum** of possibilities for accessing **probabilistic errors**.
- Balancing the **degree of inexactness (Relevance)** & the **reliance on assumptions (Robustness).**

But what is *Statistical/Probabilistic* Inference?

- An ultimate intellectual game: **"to guess wisely and to guess meaningfully the errors in our guesses."** (*XL-Files*, Oct 2015)
- Impossible to access exact errors, but **a full spectrum** of possibilities for accessing **probabilistic errors**.
- Balancing the **degree of inexactness (Relevance)** & the **reliance on assumptions (Robustness).**

Pure Frequentist (Fully unconditional)

Most Robust but Least Relevant

But what is *Statistical/Probabilistic* Inference?

- An ultimate intellectual game: **"to guess wisely and to guess meaningfully the errors in our guesses."** (*XL-Files*, Oct 2015)
- Impossible to access exact errors, but **a full spectrum** of possibilities for accessing **probabilistic errors**.
- Balancing the **degree of inexactness (Relevance)** & the **reliance on assumptions (Robustness).**

Pure Frequentist (Fully unconditional)

Most Robust but Least Relevant

Pure Bayesian (Fully conditional)

Most Relevant but Least Robust

But what is *Statistical/Probabilistic* Inference?

- An ultimate intellectual game: **"to guess wisely and to guess meaningfully the errors in our guesses."** (*XL-Files*, Oct 2015)
- Impossible to access exact errors, but **a full spectrum** of possibilities for accessing **probabilistic errors**.
- Balancing the **degree of inexactness (Relevance)** & the **reliance on assumptions (Robustness).**

Pure Frequentist (Fully unconditional)

Most Robust but Least Relevant

Pure Bayesian (Fully conditional)

Most Relevant but Least Robust

But life is about *compromise*:

Conditional frequentist, Objective Bayesian, Fiducial ...

It all depends on which *Replications* you want ...

It all depends on which *Replications* you want ...

Statistical Model via Stochastic Representation

$$\underbrace{D}_{Data} = G(\underbrace{\theta}_{Signal}, \underbrace{U}_{Noise}) \qquad (S)$$

It all depends on which *Replications* you want ...

Statistical Model via Stochastic Representation

$$\underbrace{D}_{Data} = G(\underbrace{\theta}_{Signal}, \underbrace{U}_{Noise}) \qquad (S)$$

Ex: $D = \{X_1, \ldots, X_n\}$, where

$$X_i = \theta + U_i, \quad U_i \stackrel{\text{iid}}{\sim} N(0,1),$$

and $U = \{U_i, i = 1, \ldots, n\}$ represents "God's Uncertainty"

It all depends on which *Replications* you want ...

Statistical Model via Stochastic Representation

$$\underbrace{D}_{Data} = G(\underbrace{\theta}_{Signal}, \underbrace{U}_{Noise}) \qquad (S)$$

Ex: $D = \{X_1, \ldots, X_n\}$, where

$$X_i = \theta + U_i, \quad U_i \stackrel{\text{iid}}{\sim} N(0, 1),$$

and $U = \{U_i, i = 1, \ldots, n\}$ represents "God's Uncertainty"

- Frequentist: Fix parameter θ, vary D

It all depends on which *Replications* you want ...

Statistical Model via Stochastic Representation

$$\underbrace{D}_{Data} = G(\underbrace{\theta}_{Signal}, \underbrace{U}_{Noise}) \qquad (S)$$

Ex: $D = \{X_1, \ldots, X_n\}$, where

$$X_i = \theta + U_i, \quad U_i \overset{\text{iid}}{\sim} N(0,1),$$

and $U = \{U_i, i = 1, \ldots, n\}$ represents "God's Uncertainty"

- Frequentist: Fix parameter θ, vary D
- Bayesian: Fix data D, vary θ

It all depends on which *Replications* you want ...

Statistical Model via Stochastic Representation

$$\underbrace{D}_{Data} = G(\underbrace{\theta}_{Signal}, \underbrace{U}_{Noise}) \qquad (S)$$

Ex: $D = \{X_1, \ldots, X_n\}$, where

$$X_i = \theta + U_i, \quad U_i \overset{\text{iid}}{\sim} N(0,1),$$

and $U = \{U_i, i = 1, \ldots, n\}$ represents "God's Uncertainty"

- Frequentist: Fix parameter θ, vary D
- Bayesian: Fix data D, vary θ
- Fiducial: Fix neither, but vary U, subject to the constraint (S) (or implied constraints with $A(U)$ fixed)

The differences are in the replications ...

$$\begin{tikzcd} \theta \arrow[d, no head] \\ D \end{tikzcd}$$

The differences are in the replications ...

$$\begin{tikzcd} \theta \arrow[d, dash] \\ D \end{tikzcd}$$

Frequentist Inference

$$\begin{tikzcd} & & \theta \arrow[dll] \arrow[dl] \arrow[d] \arrow[dr] \arrow[drr] & & \\ D'_1 & D'_2 & D & \cdots & D'_\infty \end{tikzcd}$$

$$p(D'|\theta)$$

The differences are in the replications ...

Frequentist Inference

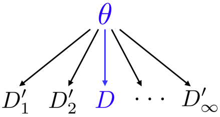

$$p(D'|\theta)$$

Bayesian Inference

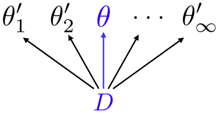

$$p(\theta'|D)$$

$$\propto p(D|\theta')\pi_0(\theta')$$

The differences are in the replications ...

$$\begin{tikzcd} \theta \arrow[d, dash] \\ D \end{tikzcd}$$

Frequentist Inference

$$\begin{tikzcd} & & \theta \arrow[dll] \arrow[dl] \arrow[d] \arrow[dr] \arrow[drr] & & \\ D'_1 & D'_2 & D & \cdots & D'_\infty \end{tikzcd}$$

$$p(D'|\theta)$$

Bayesian Inference

$$\begin{tikzcd} \theta'_1 & \theta'_2 & \theta & \cdots & \theta'_\infty \\ & & D \arrow[ull] \arrow[ul] \arrow[u] \arrow[ur] \arrow[urr] & & \end{tikzcd}$$

$$p(\theta'|D)$$

$$\propto p(D|\theta')\pi_0(\theta')$$

Fiducial Inference

$$\begin{tikzcd} \theta'_1 \arrow[d, leftrightarrow] & \theta'_2 \arrow[d, leftrightarrow] & \theta \arrow[d, leftrightarrow] & \cdots & \theta'_\infty \arrow[d, leftrightarrow] \\ D'_1 & D'_2 & D & \cdots & D'_\infty \end{tikzcd}$$

$$p(D', \theta'|A(U))$$

$$= p(D'|\theta', A(U))\pi(\theta')$$

Illustrate BFF for $X \sim N(\theta, 1)$

Frequentist

Illustrate BFF for $X \sim N(\theta, 1)$

Frequentist

Bayesian

Illustrate BFF for $X \sim N(\theta, 1)$

Frequentist

Bayesian

Fiducial

Illustrate BFF for $X \sim N(\theta, 1)$

Frequentist

Bayesian

Fiducial

Sampling Dist.
$X|\theta \sim N(\theta, 1)$

Illustrate BFF for $X \sim N(\theta, 1)$

Frequentist

Sampling Dist.
$X|\theta \sim N(\theta, 1)$

Bayesian

+ Prior Dist.
$\pi_0(\theta) \propto 1$

Fiducial

Illustrate BFF for $X \sim N(\theta, 1)$

Frequentist	Bayesian	Fiducial
Sampling Dist. $X\|\theta \sim N(\theta, 1)$	+ Prior Dist. $\pi_0(\theta) \propto 1$	God's U Dist. $X - \theta = U \sim N(0, 1)$

Illustrate BFF for $X \sim N(\theta, 1)$

Frequentist	Bayesian	Fiducial
Sampling Dist. $X \mid \theta \sim N(\theta, 1)$	+ Prior Dist. $\pi_0(\theta) \propto 1$	God's U Dist. $X - \theta = U \sim N(0, 1)$

Confidence Interval
$(X - z_p, X + z_p)$

Illustrate BFF for $X \sim N(\theta, 1)$

Frequentist

Sampling Dist.
$X|\theta \sim N(\theta, 1)$

Confidence Interval
$(X - z_p, X + z_p)$

Bayesian

+ Prior Dist.
$\pi_0(\theta) \propto 1$

Posterior Dist.
$\theta|X \sim N(X, 1)$

Fiducial

God's U Dist.
$X - \theta = U \sim N(0, 1)$

Illustrate BFF for $X \sim N(\theta, 1)$

Frequentist

Sampling Dist.
$X|\theta \sim N(\theta, 1)$

Confidence Interval
$(X - z_p, X + z_p)$

Bayesian

$+$ Prior Dist.
$\pi_0(\theta) \propto 1$

Posterior Dist.
$\theta|X \sim N(X, 1)$

Fiducial

God's U Dist.
$X - \theta = U \sim N(0, 1)$

Fiducial Dist.
$\theta = X + U \sim N(X, 1)$

Illustrate BFF for $X \sim N(\theta, 1)$

Frequentist

Sampling Dist.
$X|\theta \sim N(\theta, 1)$

Bayesian

$+$ Prior Dist.
$\pi_0(\theta) \propto 1$

Posterior Dist.
$\theta|X \sim N(X, 1)$

Fiducial

God's U Dist.
$X - \theta = U \sim N(0, 1)$

Fiducial Dist.
$\theta = X + U \sim N(X, 1)$

Confidence Interval
$(X - z_p, X + z_p)$

Posterior Interval
$(X - z_p, X + z_p)$

Illustrate BFF for $X \sim N(\theta, 1)$

Frequentist	Bayesian	Fiducial
Sampling Dist. $X\|\theta \sim N(\theta, 1)$	+ Prior Dist. $\pi_0(\theta) \propto 1$	God's U Dist. $X - \theta = U \sim N(0, 1)$
	Posterior Dist. $\theta\|X \sim N(X, 1)$	Fiducial Dist. $\theta = X + U \sim N(X, 1)$
Confidence Interval $(X - z_p, X + z_p)$	Posterior Interval $(X - z_p, X + z_p)$	Fiducial Interval $(X - z_p, X + z_p)$

Illustrate BFF for $X \sim N(\theta, 1)$

Frequentist	Bayesian	Fiducial
Sampling Dist. $X \mid \theta \sim N(\theta, 1)$	+ Prior Dist. $\pi_0(\theta) \propto 1$	God's U Dist. $X - \theta = U \sim N(0, 1)$
Confidence Dist. $N(X, 1)$	Posterior Dist. $\theta \mid X \sim N(X, 1)$	Fiducial Dist. $\theta = X + U \sim N(X, 1)$
↑	↓	↓
Confidence Interval $(X - z_p, X + z_p)$	Posterior Interval $(X - z_p, X + z_p)$	Fiducial Interval $(X - z_p, X + z_p)$

Illustrate BFF for $X \sim N(\theta, 1)$

Frequentist	Bayesian	Fiducial
Sampling Dist. $X\|\theta \sim N(\theta, 1)$	+ Prior Dist. $\pi_0(\theta) \propto 1$	God's U Dist. $X - \theta = U \sim N(0, 1)$
Confidence Dist. $N(X, 1)$	Posterior Dist. $\theta\|X \sim N(X, 1)$	Fiducial Dist. $\theta = X + U \sim N(X, 1)$
↑	↓	↓
Confidence Interval $(X - z_p, X + z_p)$	Posterior Interval $(X - z_p, X + z_p)$	Fiducial Interval $(X - z_p, X + z_p)$

Generate for $i = 1, \ldots$
$X_i|\theta \sim N(\theta, 1)$, then
$(X_i - 1.96, X_i + 1.96)$
covers θ 95% of times

Illustrate BFF for $X \sim N(\theta, 1)$

Frequentist	Bayesian	Fiducial
Sampling Dist. $X \mid \theta \sim N(\theta, 1)$	+ Prior Dist. $\pi_0(\theta) \propto 1$	God's U Dist. $X - \theta = U \sim N(0, 1)$
Confidence Dist. $N(X, 1)$ ↑	Posterior Dist. $\theta \mid X \sim N(X, 1)$ ↓	Fiducial Dist. $\theta = X + U \sim N(X, 1)$ ↓
Confidence Interval $(X - z_p, X + z_p)$	Posterior Interval $(X - z_p, X + z_p)$	Fiducial Interval $(X - z_p, X + z_p)$
Generate for $i = 1, \ldots$ $X_i \mid \theta \sim N(\theta, 1)$, then $(X_i - 1.96, X_i + 1.96)$ covers θ 95% of times	Generate for $i = 1, \ldots$ $\theta_i \mid X \sim N(X, 1)$, then $(X - 1.96, X + 1.96)$ covers θ_i 95% of times	

Illustrate BFF for $X \sim N(\theta, 1)$

Frequentist	Bayesian	Fiducial
Sampling Dist. $X \mid \theta \sim N(\theta, 1)$	$+$ Prior Dist. $\pi_0(\theta) \propto 1$	God's U Dist. $X - \theta = U \sim N(0, 1)$
Confidence Dist. $N(X, 1)$	Posterior Dist. $\theta \mid X \sim N(X, 1)$	Fiducial Dist. $\theta = X + U \sim N(X, 1)$
$\uparrow$	$\downarrow$	$\downarrow$
Confidence Interval $(X - z_p, X + z_p)$	Posterior Interval $(X - z_p, X + z_p)$	Fiducial Interval $(X - z_p, X + z_p)$
Generate for $i = 1, \ldots$ $X_i \mid \theta \sim N(\theta, 1)$, then $(X_i - 1.96, X_i + 1.96)$ covers θ 95% of times	Generate for $i = 1, \ldots$ $\theta_i \mid X \sim N(X, 1)$, then $(X - 1.96, X + 1.96)$ covers θ_i 95% of times	Generate for $i = 1, \ldots$ $\theta_i \sim$ **any** $\pi(\theta)$, & $X_i \mid \theta_i \sim N(\theta_i, 1)$, then $(X_i - 1.96, X_i + 1.96)$ covers θ_i 95% of times

Finding the Right "Control Population": Treating Data as Your Patient

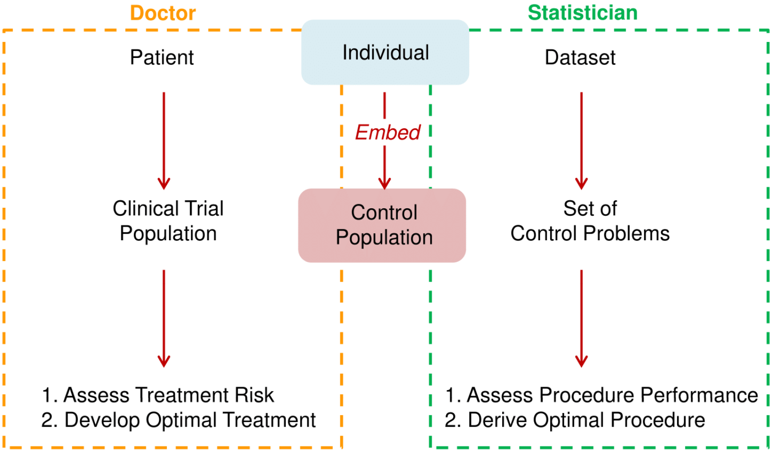

The Inevitable Statistical "Bootstrap": Creating Internal Replications

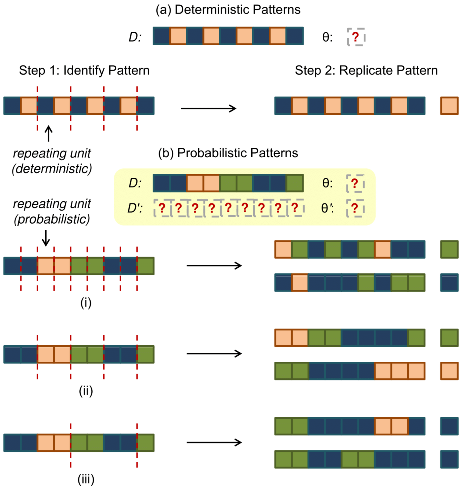

Relevant Controls/Replications are always needed

Method Type	Error on Actual Problem Δ	Average Error over Relevant Controls $\bar{\Delta}'$	References
Point Estimate *Goal*: Give our best guess, $\hat{\theta}$, for value of θ.	$\mathbb{L}(\theta,\hat{\theta})$ *Loss*: Specify how "far" θ is from $\hat{\theta}$ via loss function.	*Risk*: The average loss of an estimator over control problems (D', θ').	Robinson 1979b Rukhin 1988, Lu and Berger 1989, Fourdrinier and Wells 2012
Set Estimate *Goal*: Identify set, $C(D)$, of likely values for θ.	$\mathbb{I}(\theta \notin C(D))$ * *Coverage*: Does our set contain the true value of θ?	*Non-Coverage Probability*: Proportion of times a set estimate, e.g. interval estimate, fails to contain the true value of θ'.	Casella 1992, Goutis and Casella 1995, Robinson 1979a, Berger 1988
Hypothesis Test *Goal*: Should we reject a null hypothesis, H_0, based on evidence from data?	$\mathbb{I}(\hat{T} \neq T)$ *Type I or II Error*: Do we falsely reject or falsely accept H_0?	*Error Probability*: The test's rates of false rejection and false acceptance when applied to control problems.	Hwang et al. 1992, Berger et al. 1994, Berger 2003

* $\mathbb{I}$(*statement*) denotes the indicator function: it equals 1 if the statement in parentheses is true and 0 otherwise.

Multi-resolution Replications

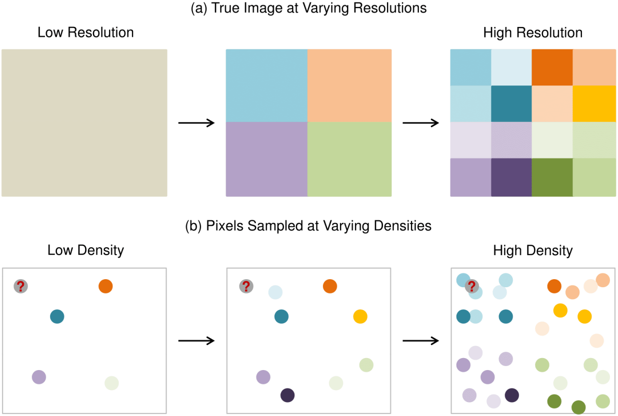

The Problem Gets Easier But My Intervals Get Longer ?!

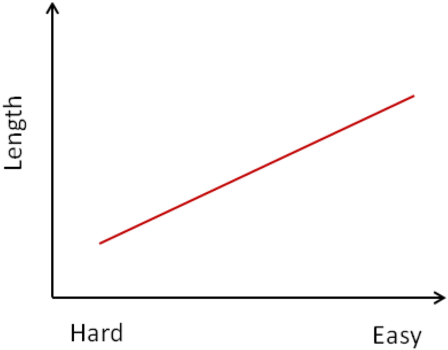

Precision as Function of Multiple Features (Basu 1964)

(X_i, Y_i) bivariate standard normal with unknown correlation θ

Precision as Function of Multiple Features (Basu 1964)

(X_i, Y_i) bivariate standard normal with unknown correlation θ

- Fact 1: X_i, Y_i marginally ancillary, not jointly ancillary.

Precision as Function of Multiple Features (Basu 1964)

(X_i, Y_i) bivariate standard normal with unknown correlation θ

- Fact 1: X_i, Y_i marginally ancillary, not jointly ancillary.
- Fact 2: As $||X||$ or $||Y||$ increases, precision for θ increases.

Precision as Function of Multiple Features (Basu 1964)

(X_i, Y_i) bivariate standard normal with unknown correlation θ

- Fact 1: X_i, Y_i marginally ancillary, not jointly ancillary.
- Fact 2: As $||X||$ or $||Y||$ increases, precision for θ increases.

Option 1: Evaluate uncertainty of $\hat{\theta}$ (MLE) *unconditionally*. Construct pivot (using inverse CDF) and invert into CI.

Precision as Function of Multiple Features (Basu 1964)

(X_i, Y_i) bivariate standard normal with unknown correlation θ

- Fact 1: X_i, Y_i marginally ancillary, not jointly ancillary.
- Fact 2: As $||X||$ or $||Y||$ increases, precision for θ increases.

Option 1: Evaluate uncertainty of $\hat{\theta}$ (MLE) *unconditionally*. Construct pivot (using inverse CDF) and invert into CI.

- Achieves exact, unconditional coverage.

Precision as Function of Multiple Features (Basu 1964)

(X_i, Y_i) bivariate standard normal with unknown correlation θ

- Fact 1: X_i, Y_i marginally ancillary, not jointly ancillary.
- Fact 2: As $||X||$ or $||Y||$ increases, precision for θ increases.

Option 1: Evaluate uncertainty of $\hat{\theta}$ (MLE) *unconditionally*. Construct pivot (using inverse CDF) and invert into CI.

- Achieves exact, unconditional coverage.

Option 2: Evaluate uncertainty of $\hat{\theta}$ *conditional* on $||X||$.

- But what about the effect of $||Y||$ on precision?

A Heterogeneous Population of Datasets

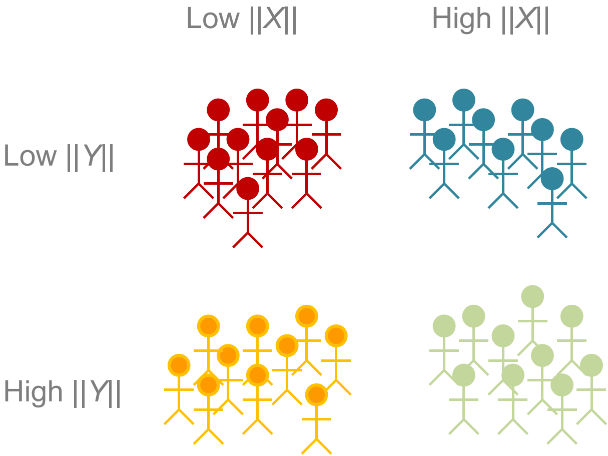

Here's Where Resolution Helps Us Reason...

A Regression Perspective

- As $||X||$ increases, precision of $\hat{\theta}$ increases.

Here's Where Resolution Helps Us Reason...

A Regression Perspective

- As $||X||$ increases, precision of $\hat{\theta}$ increases.
- As $||Y||$ increases, precision of $\hat{\theta}$ increases.

Here's Where Resolution Helps Us Reason...

A Regression Perspective

- As $||X||$ increases, precision of $\hat{\theta}$ increases.
- As $||Y||$ increases, precision of $\hat{\theta}$ increases.
- The first order effects of $||X||$ and $||Y||$ on precision are robust to assumptions about θ.

Here's Where Resolution Helps Us Reason...

A Regression Perspective

- As $||X||$ increases, precision of $\hat{\theta}$ increases.
- As $||Y||$ increases, precision of $\hat{\theta}$ increases.
- The first order effects of $||X||$ and $||Y||$ on precision are robust to assumptions about θ.

But when we condition on $||X||$ *and* $||Y||$ ...

- We also model second order effect: how $||X||$ and $||Y||$ *together* affect data precision (their interaction).

Here's Where Resolution Helps Us Reason...

A Regression Perspective

- As $||X||$ increases, precision of $\hat{\theta}$ increases.
- As $||Y||$ increases, precision of $\hat{\theta}$ increases.
- The first order effects of $||X||$ and $||Y||$ on precision are robust to assumptions about θ.

But when we condition on $||X||$ *and* $||Y||$ ...

- We also model second order effect: how $||X||$ and $||Y||$ *together* affect data precision (their interaction).
- Second order effect (interaction term) is not robust to prior assumptions about θ.

Here's Where Resolution Helps Us Reason...

A Regression Perspective

- As $||X||$ increases, precision of $\hat{\theta}$ increases.
- As $||Y||$ increases, precision of $\hat{\theta}$ increases.
- The first order effects of $||X||$ and $||Y||$ on precision are robust to assumptions about θ.

But when we condition on $||X||$ *and* $||Y||$ ...

- We also model second order effect: how $||X||$ and $||Y||$ *together* affect data precision (their interaction).
- Second order effect (interaction term) is not robust to prior assumptions about θ.
- How to account for first order effects while ignoring second order effects and do so in a *principled* way?

Fiducial's Pivotal Idea (Fraser 68, Hannig 09)

God's U Always Exists

Represent data as $X = g(\theta; U)$ where $U \sim p(U)$ is known.

$$\text{Normal}: \qquad X = \theta + U \qquad U \sim N(0,1)$$

Fiducial's Pivotal Idea (Fraser 68, Hannig 09)

God's U Always Exists

Represent data as $X = g(\theta; U)$ where $U \sim p(U)$ is known.

$$\text{Normal}: \qquad X = \theta + U \qquad U \sim N(0, 1)$$

Such representations even exist in cases where pivots do not:

$$\text{Bernoulli}: \quad X = I(U < \theta) \qquad U \sim \text{Unif}[0, 1].$$

Fiducial's Pivotal Idea (Fraser 68, Hannig 09)

God's U Always Exists

Represent data as $X = g(\theta; U)$ where $U \sim p(U)$ is known.

$$\text{Normal}: \quad X = \theta + U \qquad U \sim N(0,1)$$

Such representations even exist in cases where pivots do not:

$$\text{Bernoulli}: \quad X = I(U < \theta) \qquad U \sim \text{Unif}[0,1].$$

Fiducial Procedure

1. Make a "post-data" inference for U without involving θ by ignoring a part or all data: e.g., *pretend* $U|X \sim N(0,1)$.

Fiducial's Pivotal Idea (Fraser 68, Hannig 09)

God's U Always Exists

Represent data as $X = g(\theta; U)$ where $U \sim p(U)$ is known.

$$\text{Normal}: \quad X = \theta + U \qquad U \sim N(0,1)$$

Such representations even exist in cases where pivots do not:

$$\text{Bernoulli}: \quad X = I(U < \theta) \qquad U \sim \text{Unif}[0,1].$$

Fiducial Procedure

1. Make a "post-data" inference for U without involving θ by ignoring a part or all data: e.g., *pretend* $U|X \sim N(0,1)$.
2. Convert inference for U into inference for θ by inverting $X = g(\theta; U)$ to obtain $\theta = h(U; X)$:

$$\text{E.g.}: \quad \theta = X - U \sim N(X,1).$$

(X_i, Y_i) bivariate normal with mean 0, var 1 and correlation θ.

- Reduce to sufficient statistics: $S_1 = \sum_i (X_i + Y_i)^2$ and $S_2 = \sum_i (X_i - Y_i)^2$.

Fiducial Inference for Bivariate Normal

(X_i, Y_i) bivariate normal with mean 0, var 1 and correlation θ.

- Reduce to sufficient statistics: $S_1 = \sum_i (X_i + Y_i)^2$ and $S_2 = \sum_i (X_i - Y_i)^2$.
- Representation: $S_1 = 4(1+\theta)^2 Q_1$ and $S_2 = 4(1-\theta)^2 Q_2$ where Q_i are i.i.d. $\chi^2_{(n)}$.

Fiducial Inference for Bivariate Normal

(X_i, Y_i) bivariate normal with mean 0, var 1 and correlation θ.

- Reduce to sufficient statistics: $S_1 = \sum_i (X_i + Y_i)^2$ and $S_2 = \sum_i (X_i - Y_i)^2$.
- Representation: $S_1 = 4\,(1+\theta)^2\, Q_1$ and $S_2 = 4\,(1-\theta)^2\, Q_2$ where Q_i are i.i.d. $\chi^2_{(n)}$.
- Inference for Q_1, Q_2: Impute Q_1 and Q_2 conditional on

$$\sqrt{\frac{S_1}{4Q_1}} + \sqrt{\frac{S_2}{4Q_2}} - 2 = 0$$

and $Q_i \geq S_i/16$ for $i = 1, 2$.

(X_i, Y_i) bivariate normal with mean 0, var 1 and correlation θ.

- Reduce to sufficient statistics: $S_1 = \sum_i (X_i + Y_i)^2$ and $S_2 = \sum_i (X_i - Y_i)^2$.
- Representation: $S_1 = 4(1+\theta)^2 Q_1$ and $S_2 = 4(1-\theta)^2 Q_2$ where Q_i are i.i.d. $\chi^2_{(n)}$.
- Inference for Q_1, Q_2: Impute Q_1 and Q_2 conditional on

$$\sqrt{\frac{S_1}{4Q_1}} + \sqrt{\frac{S_2}{4Q_2}} - 2 = 0$$

  and $Q_i \geq S_i/16$ for $i = 1, 2$.
- Inference for θ: Given Q_1, Q_2, let $\theta = \sqrt{\frac{S_1}{4Q_1}} - 1$.

Checking Coverage and Length Conditioning on $||x||$

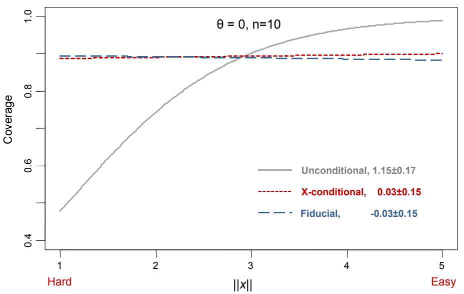

Checking Coverage and Length Conditioning on $||x||$

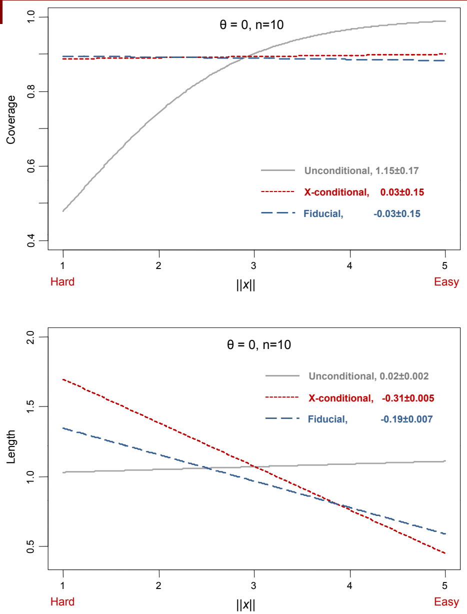

Checking Coverage and Length Conditioning on $||y||$

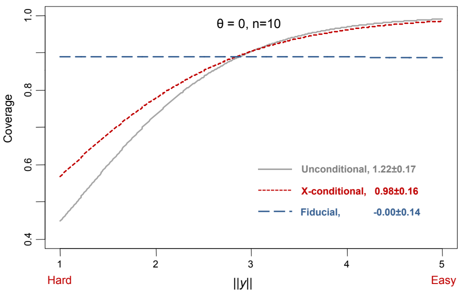

Checking Coverage and Length Conditioning on $||y||$

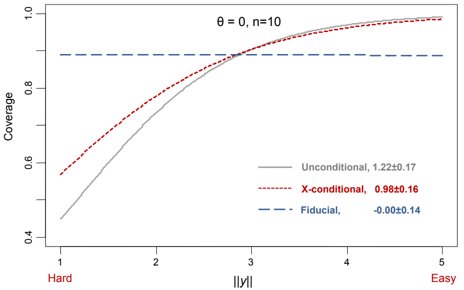

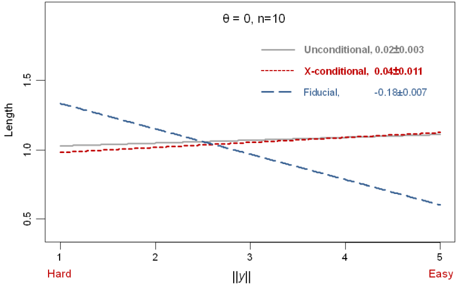

A Fundamental Principle of Statistical Inference: Bias-Variance or Relevant-Robust Trade-off

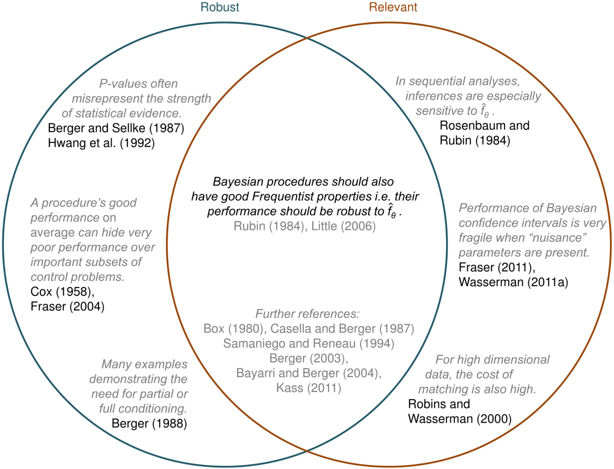

A Unified Picture of BFF (and Inference)?

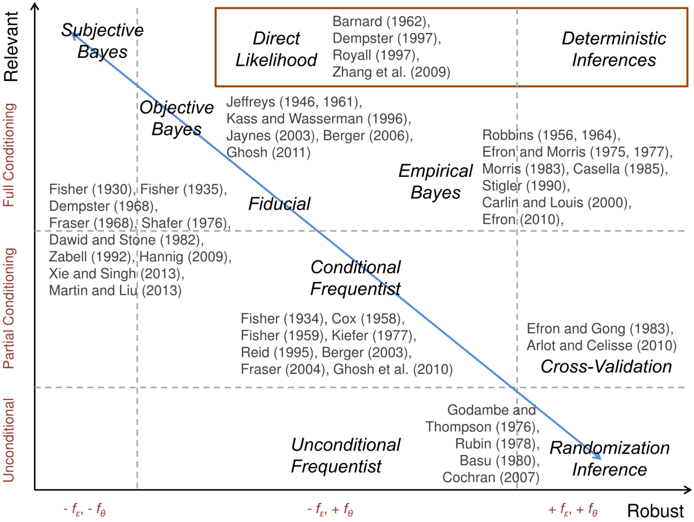

Let's be BFF, not merely FWB ...