P&G data&modelingsciences

BAYESIALAB 7TH ANNUAL BayesiaLab Conference DURHAM 2019
NORTH CAROLINA BIOTECHNOLOGY CENTER

Understanding Your Customer Using the “Most Relevant Explanations” (MRE) in BayesiaLab

Michael L. Thompson, Ph.D.

Research Fellow, Corporate R&D, Data & Modeling Sciences

Procter & Gamble Company

Oct. 11, 2019

Overview

- Why & What are Some Applications?
- What is “Most Relevant Explanation”, MRE?
- How to Use It in BayesiaLab?
- Example: Understanding Customers

Why?
Motivation…

- We need actionable insights – where can we intervene?
- We need to understand how our machine learning systems behave.
- We need to distinguish classes/groups by the most discriminating traits.
- We need to discover “interesting” cases in our Big Data – representative of key phenomena.

What are Some Applications?

- **Explainable AI**
  - **Recommenders:** Why did our system recommend that option to that person?
    Evidence, E: person profile; Hypothesis, H: option
  - **Classifiers:** Why did our Deep Learning Neural Net classifier classify that person as "high-risk"?
    E: person profile; H:assigned class = "high-risk"
- **Data-based Exemplar Identification**
  - **Database Discovery:** Which cases in my data-base most strongly exemplify the traits of interest?
    E: traits; {H}: ranked cases
- **Causal Inference for Intervention**
  - **Fault & Medical Diagnosis:** What is the underlying cause of the symptoms we're observing?
    E: symptoms; {H}: ranked causes

What is "Most Relevant Explanation"?

Yuan, C., et al.
Most relevant explanations in Bayesian networks, *J. AI Research*, 2011

- **Generalized Bayes Factors, GBF(H_i;E):** Given evidence E, contrast hypothetical scenarios H_1, H_2, … of any kind and complexity (Turing & Good).
  - ***Not*** Bayes Factor, BF(M_i;E): Given evidence E, compare models M_0, M_1,… (i.e., likelihood ratio) → out-of-favor in modern model selection:
    e.g., A. Gelman, LOOIC for mod. sel.; sparsity-inducing priors for var. sel.
- **Most Relevant Explanation, MRE({H};E):** Given evidence E and a BBN describing the relevant universe, exhaustively search through the space of possible hypotheses in that universe and rank them by **GBF(H_i;E)**. (Yuan et al.)
  - The hypotheses with the top K GBF(H_i;E) are called the "KRE".

See also:
Good, I.J., "Weight of Evidence: A Brief Survey", *Bayesian Statistics 2*, 1985
Tenenbaum, J. & Griffith, T., The Rational Basis of Representativeness, *Proc. 23rd Ann. Conf. Cog. Sci. Soc 2001*.
Fitelson, B., Likelihoodism, Bayesianism, and Relational Confirmation, *Synthese*, 2007.

Foundations: Generalized Bayes Factor GBF(H;E)

GBF(H;E) = P(E|H)/P(E|≠H) = O(H|E)/O(H)

Weight of Evidence, WE(H;E)=log(GBF(H;E))

Odds Ratio O(H|E)/O(H)

- Example: Product Recommender
  - E = Hi-Income shopper profile
  - H1 = Bread, H2 = Artisan cheeses
- O(H|E): Odds of buying product H amongst shoppers like E
- O(H): Odds of buying product H amongst ALL shoppers.
- Contrasting two hypotheses H1 & H2:

$$\frac{\mathrm{O}(Art.cheeses|Hi-income)}{\mathrm{O}(Art.cheeses)} \gg \frac{\mathrm{O}(Bread|Hi-income)}{\mathrm{O}(Bread)}$$

"Likelihood" Ratio P(E|H)/P(E|≠H)

- Example: Medical Diagnosis
  - E1 = runny nose, E2 = Xray+
  - H1 = common cold, H2 = pneumonia
- P(E1|H): Probability of runny nose amongst patients with ailment H
- P(E1|≠H): Probability of runny nose amongst patients not having ailment H (i.e., ALL alternatives: those who are healthy AND those with all other ailments)
- Contrasting two evidence scenarios E1 & E2:

$$\frac{\mathrm{P}(Xray+|pneumonia)}{\mathrm{P}(Xray+|\neq pneumonia)} \gg \frac{\mathrm{P}(runny\ nose|pneumonia)}{\mathrm{P}(runny\ nose|\neq pneumonia)}$$

$$\gg$$

$$\frac{\mathrm{P}(Xray+|cold)}{\mathrm{P}(Xray+|\neq cold)} \ll \frac{\mathrm{P}(runny\ nose|cold)}{\mathrm{P}(runny\ nose|\neq cold)}$$

Contrasting two hypotheses H1 & H2

Interpretation of GBF(H;E)

Tip: When interpreting GBF(H;E), opt for the narrative (odds ratio vs. "likelihood ratio") that corresponds to the more natural (causal) "generative story".

Classifiers/Recommenders:
"Shopper E buys Product H", so condition on E:
If $\mathrm{GBF}(H;E) = O(H|E)/O(H) \gg 1$.
Interpretation: "The **odds in favor** of purchasing Product H is much higher amongst shoppers like E than it is amongst all shoppers."

Diagnosis:
"Disease H manifests as Symptom E", so condition on H:
If $\mathrm{GBF}(H;E) = P(E|H)/P(E|\neq H) \gg 1$.
Interpretation: "The **probability of observing** Symptom E is much higher amongst patients with Disease H than it is amongst patients without Disease H."

$$\frac{\mathrm{O}(Art.cheeses|Hi{-}income)}{\mathrm{O}(Art.cheeses)} \gg 1$$

$$\frac{\mathrm{P}(Xray{+}|pneumonia)}{\mathrm{P}(Xray{+}|{\neq}pneumonia)} \gg 1$$

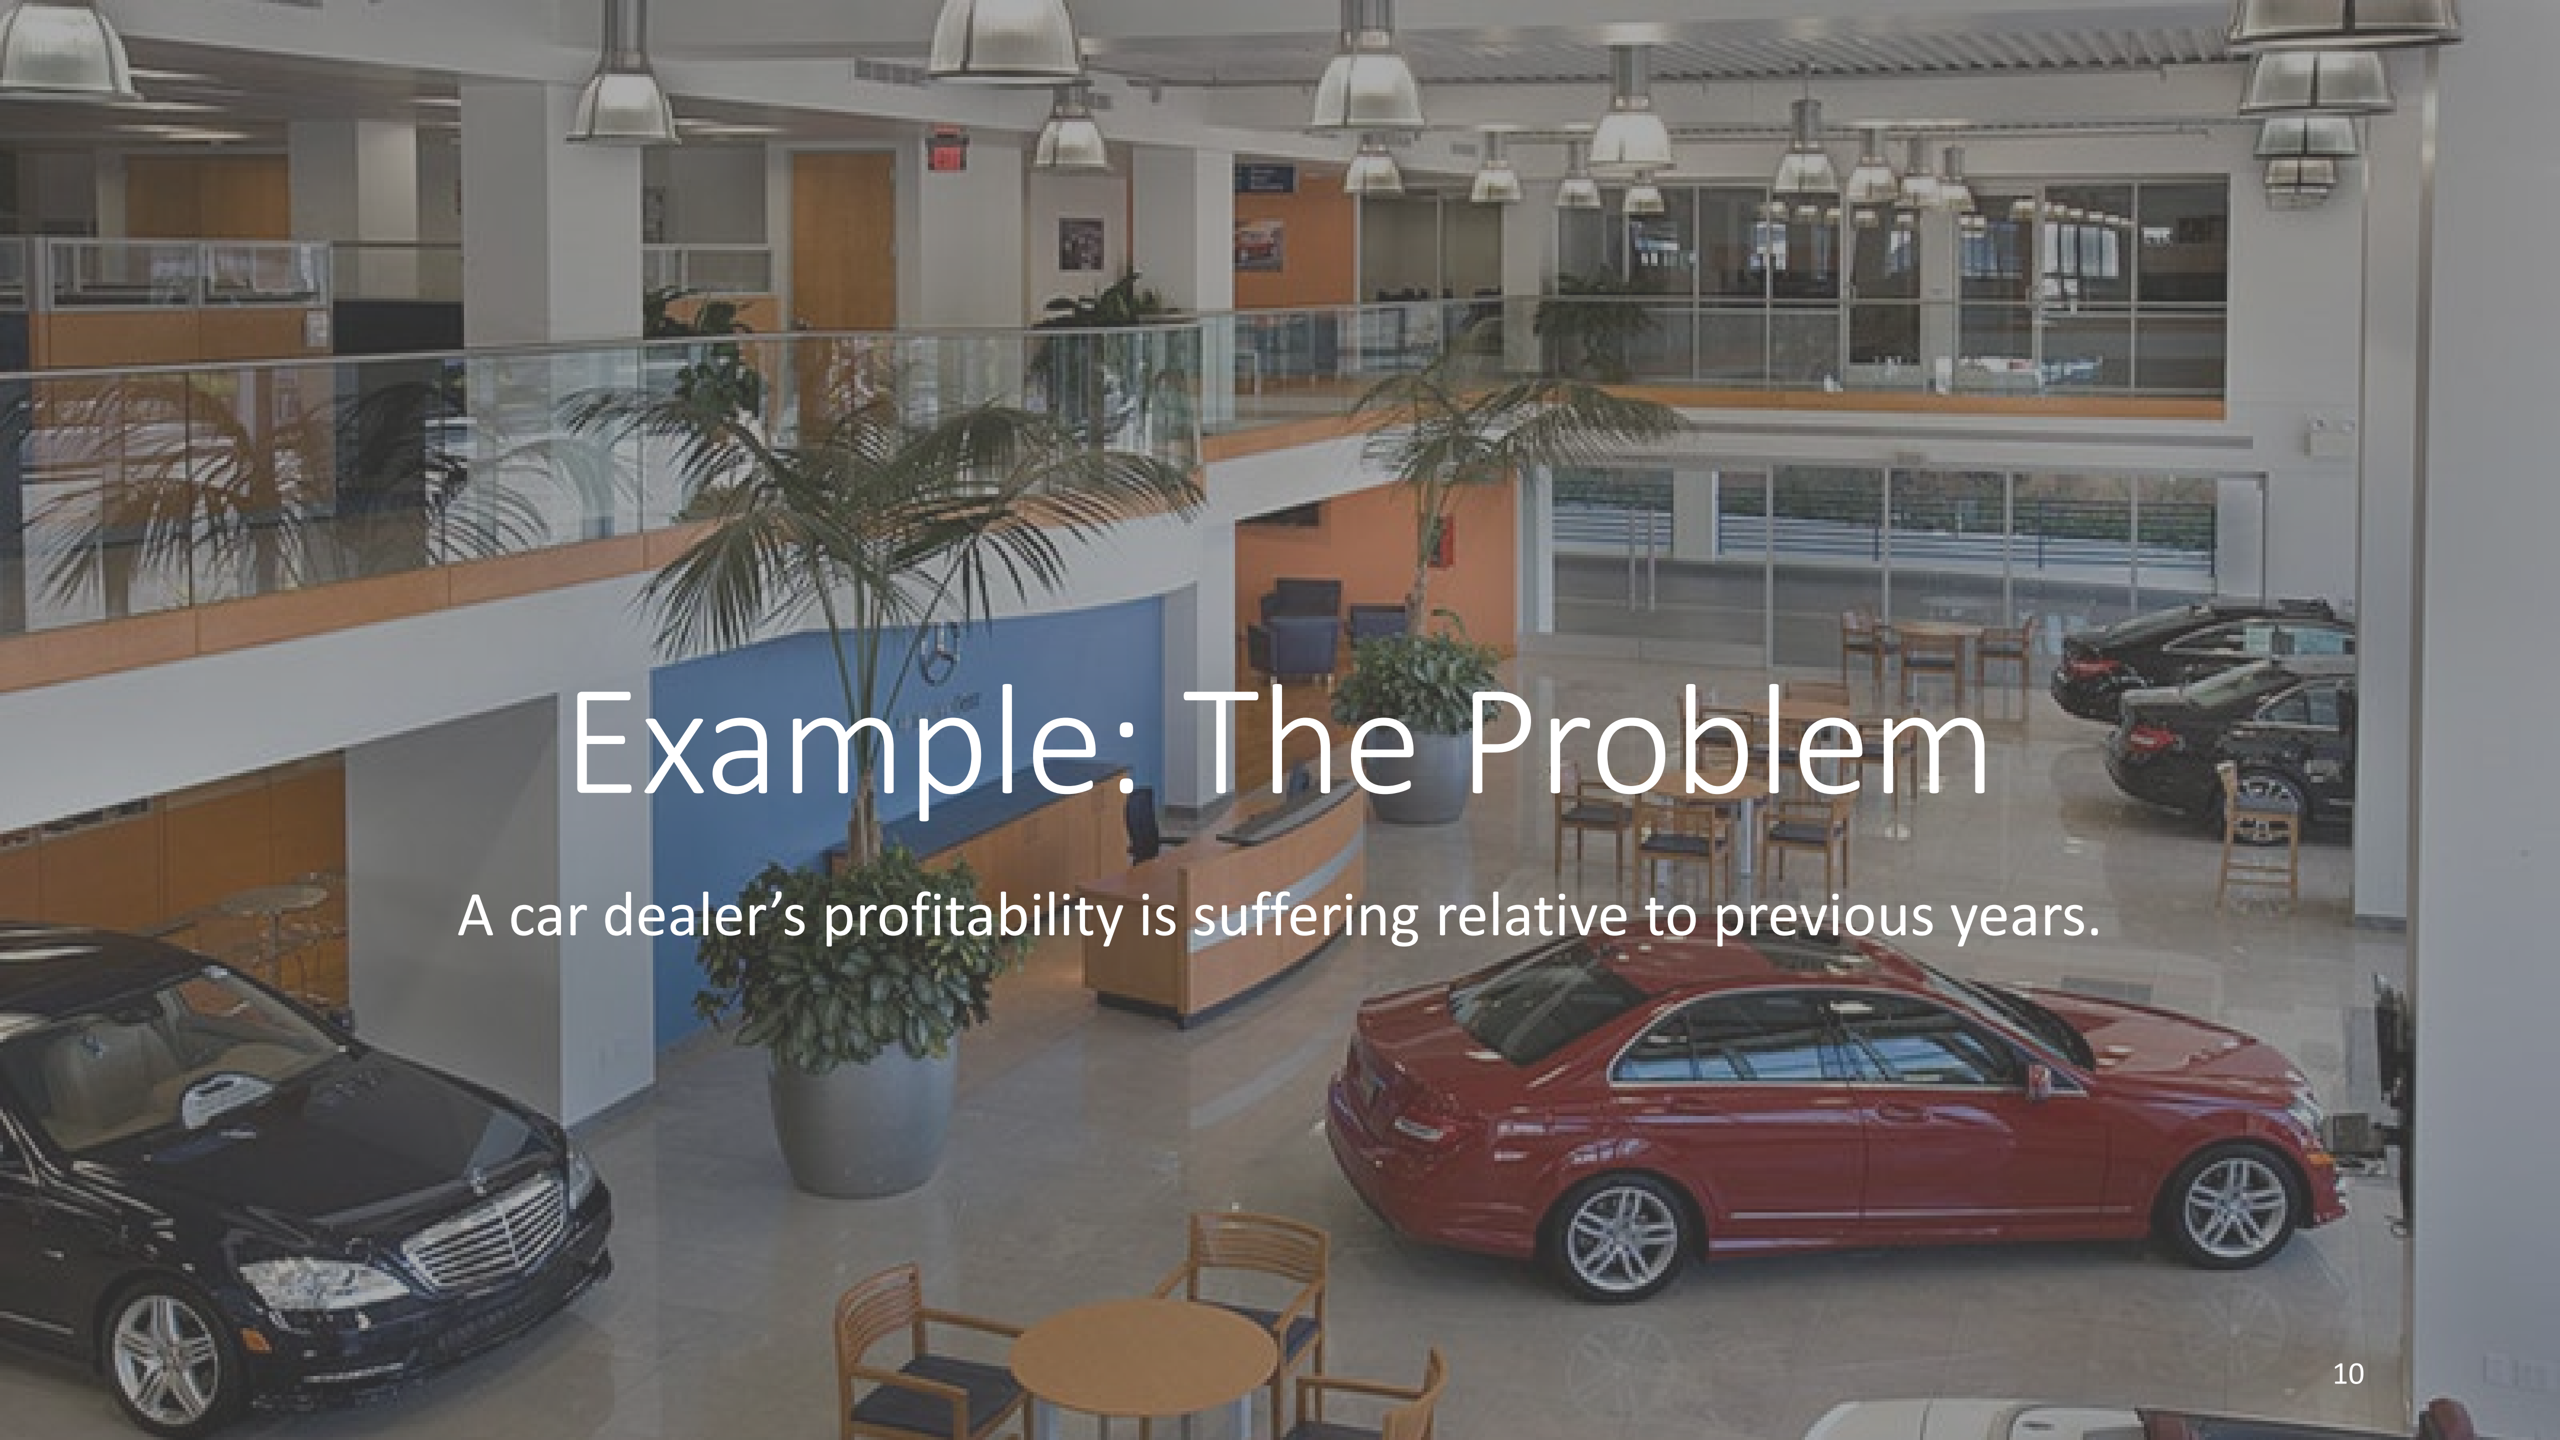

Example: The Problem

A car dealer’s profitability is suffering relative to previous years.

Example: The Questions

- What can the dealer change to improve profitability?
    - What intervention gives the dealer the "biggest bang for the buck"?
    - What distinguishes the customers who are the dealer's greatest profit opportunity vs. all other customers?
    - How do the most satisfied customers differ from the least satisfied customers?

Example: The Data & Models

- **Data:**
  - Survey response data from several hundred customers
    - Demographics, Attitudes/Agreement, Relationship exclusivity, Satisfaction w/Service, Satisfaction w/Vehicle
  - Customer history
    - Services used, Vehicles purchased, Recency & Frequency
- **Models:**
  - Probabilistic Structural Equations Model (PSEM) ***capturing domain knowledge & data*** in the form of a Bayesian belief network (BBN) built using BayesiaLab
  - Associative BBN capturing observational (correlative) relationships amongst all survey responses & customer history data

Model: Associative Bayesian Belief Network

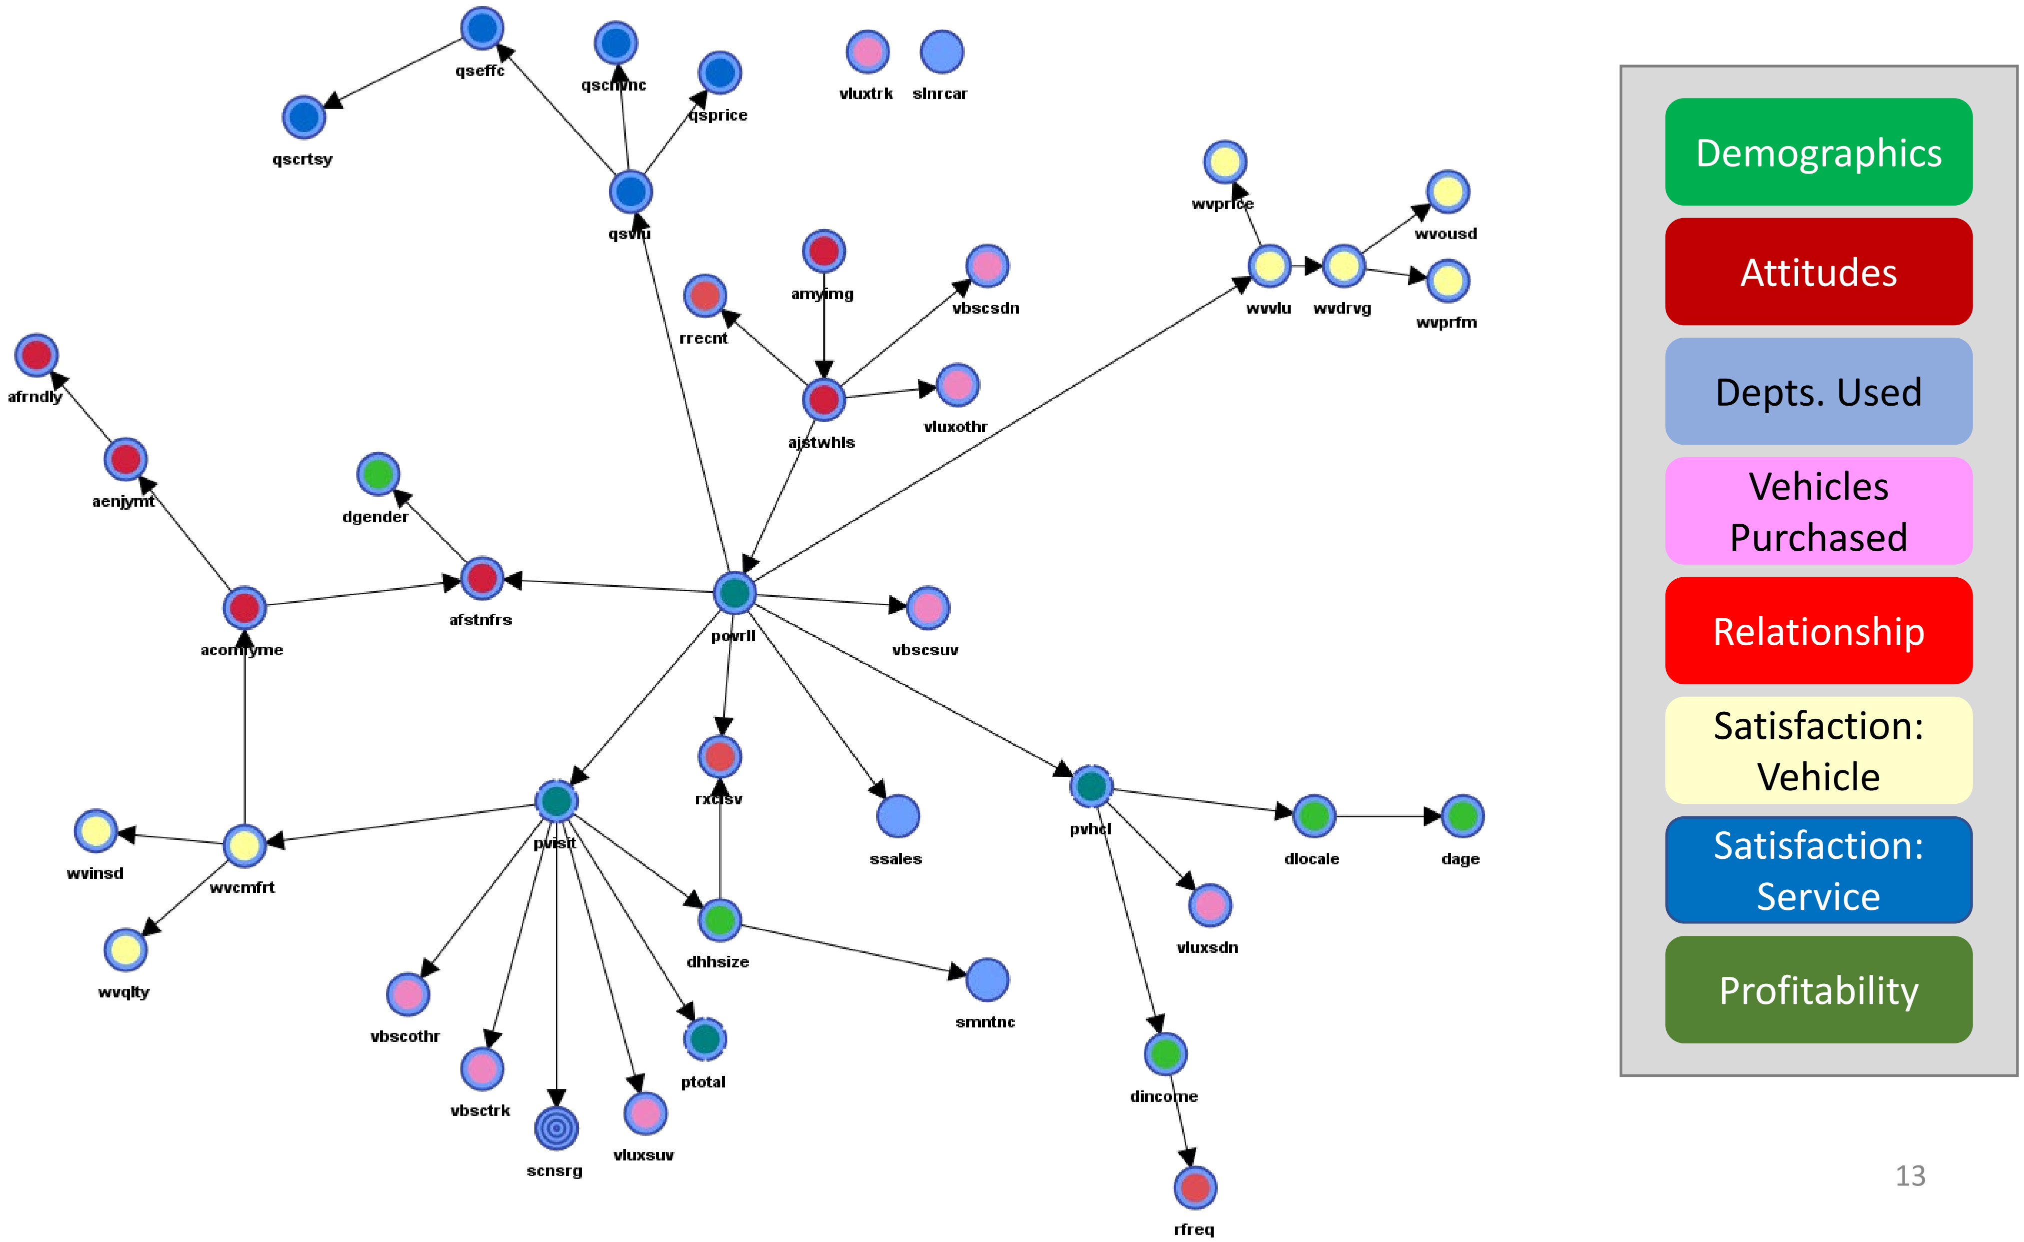

Example: The Insights

- What distinguishes the customers who are the dealer’s most profitable vs. those who are the dealer’s least profitable?

How to use MRE in BayesiaLab?

- In Validation Mode:
  - Analysis → Report → Evidence → Most Relevant Explanations

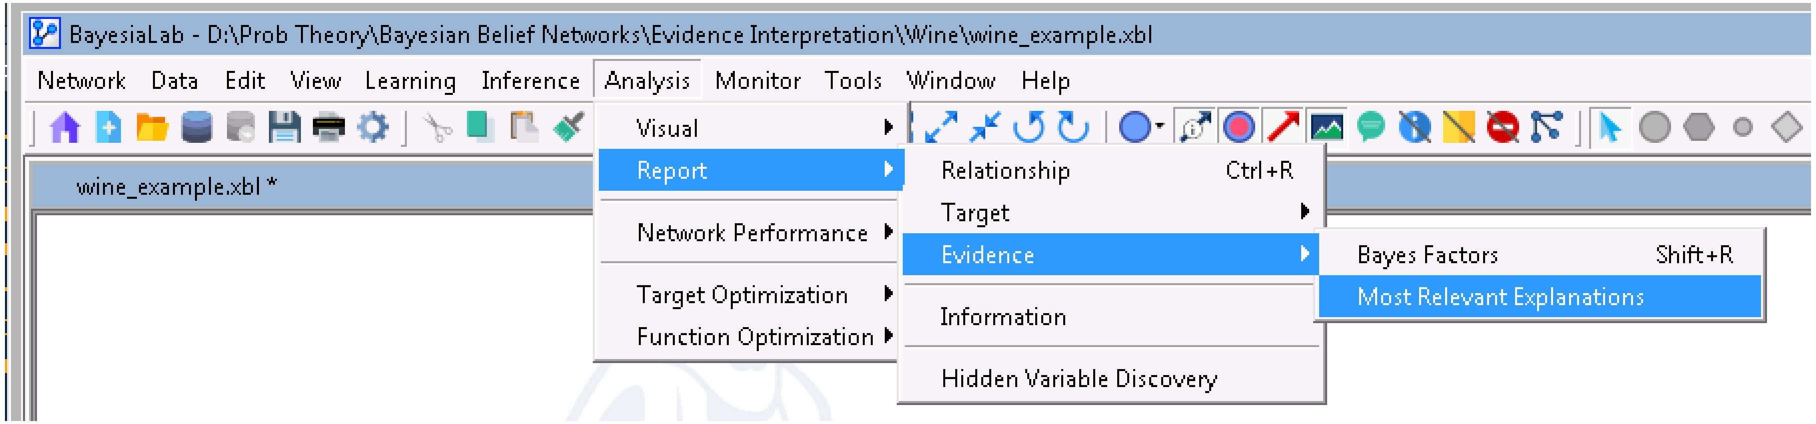

Steps in BayesiaLab v. 8.1.3

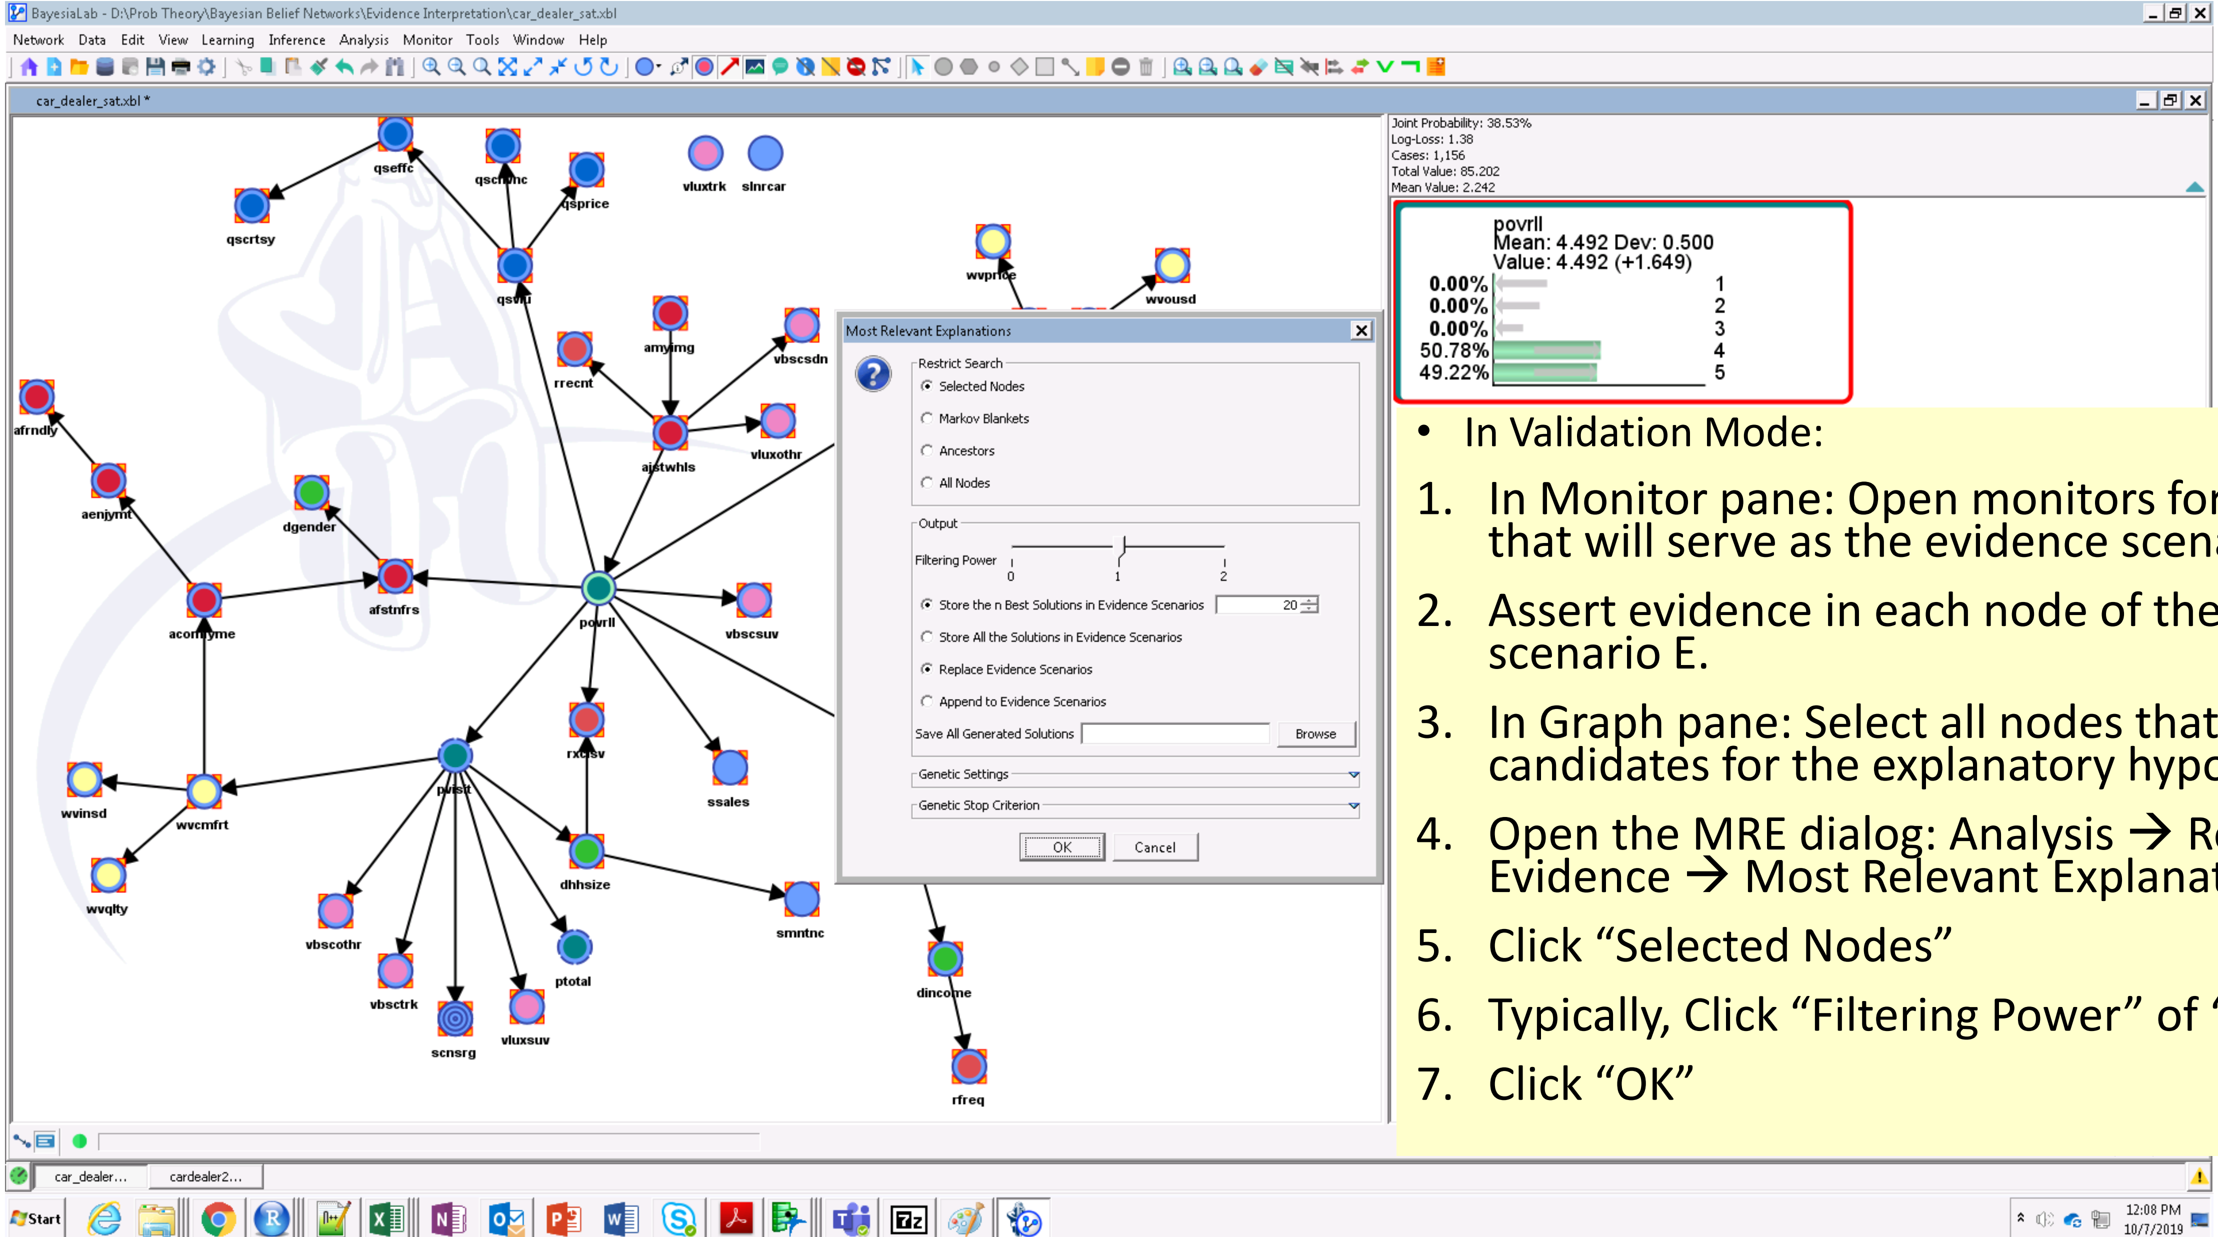

- In Validation Mode:

1. In Monitor pane: Open monitors for all nodes that will serve as the evidence scenario E.
2. Assert evidence in each node of the evidence scenario E.
3. In Graph pane: Select all nodes that will serve as candidates for the explanatory hypotheses H.
4. Open the MRE dialog: Analysis → Report → Evidence → Most Relevant Explanations
5. Click "Selected Nodes"
6. Typically, Click "Filtering Power" of "1".
7. Click "OK"

MRE for High-Profit

For example: The probability of finding “Hi-Profit” amongst customers who did NOT buy a “Basic SUV” is ~6 times higher than it is amongst customers who did buy a “Basic SUV”.

Analysis Context		
povrll	Profit: OVERALL	I {1: 0.00%, 2: 0.00%, 3: 0.00%, 4: 100.00%, 5: 100.00%}

Evidence, E

Best Solutions

Attitude: Just Wheels (ajstwhls)	Satisfaction, Service: Value (qsvlu)	Depts. Used: Sales (ssales)	Vehicle Purchased: Basic SUV (vbscsuv)	Satisfaction, Vehicle: Value (wvvlu)	Demog.: HH Size (dhhsize)	Attitude: Comfy Me (acomfyme)	Relationship: Exclusivity (rxclsv)	Attitude: Fast & Furious (afstnfrs)	MRE size	Weight of Evidence	General Bayes Factor	Likelihood $P(e \mid h)$	Posterior Odds $O(h \mid e)$	Posterior Probabiity $P(h \mid e)$
			No						1	2.5	5.8	47%	2.1E+01	95%
Strongly Disagree									1	1.6	3.0	99%	3.1E-01	24%
Disagree									1	1.6	3.0	93%	4.4E-01	31%
	Excellent			Excellent					2	1.4	2.6	97%	3.8E-02	4%
	Excel									1.3	2.5	96%	3.4E-02	3%
										1.3	2.5	94%	6.8E-02	6%
	Excel									1.3	2.5	94%	6.0E-02	6%
										1.3	2.5	93%	8.5E-02	8%
										1.3	2.5	91%	1.2E-01	11%
	Very									1.3	2.5	92%	9.8E-02	9%
	Excel									1.3	2.5	92%	7.7E-02	7%
	Excel									1.3	2.5	91%	1.1E-01	10%
	Excel									1.3	2.5	86%	2.2E-01	18%
										1.3	2.5	91%	8.0E-02	7%
Neither Agree nor Disagree										1.3	2.4	91%	4.7E-02	4%
										1.3	2.4	83%	2.4E-01	19%
										1.3	2.4	89%	7.0E-02	7%
	Very									1.2	2.4	84%	1.7E-01	14%
										1.2	2.4	84%	1.4E-01	12%
										1.2	2.3	83%	1.7E-01	14%

Interpretation [edit]

Wikipedia: “Bayes Factor: Interpretation”

only considers evidence *against* it. Harold Jeffreys gave a scale for interpretation of K:[9]

K	dHart	bits	Strength of evidence
$< 10^0$	0	—	Negative (supports M_2)
10^0 to $10^{1/2}$	0 to 5	0 to 1.6	Barely worth mentioning
$10^{1/2}$ to 10^1	5 to 10	1.6 to 3.3	Substantial
10^1 to $10^{3/2}$	10 to 15	3.3 to 5.0	Strong
$10^{3/2}$ to 10^2	15 to 20	5.0 to 6.6	Very strong
$> 10^2$	> 20	> 6.6	Decisive

Hypotheses, H_i

MRE for High-Profitability Customers (P(E)=39%)

Analysis Context		
povrll	Profit: OVERALL	I {1: 0.00%, 2: 0.00%, 3: 0.00%, 4: 100.00%, 5: 100.00%}

"bits" "K"

Best Solutions

Attitude: Just Wheels (ajstwhls)	Satisfaction, Service: Value (qsvlu)	Depts. Used: Sales (ssales)	Vehicle Purchased: Basic SUV (vbscsuv)	Satisfaction, Vehicle: Value (wvvlu)	Demog.: HH Size (dhhsize)	Attitude: Comfy Me (acomfyme)	Relationship: Exclusivity (rxclsv)	Attitude: Fast & Furious (afstnfrs)	MRE size	Weight of Evidence	Generalized Bayes Factor	Likelihood P(e|h)	Posterior Odds O(h|e)	Posterior Probability P(h|e)
			No						1	2.5	5.8	47%	2.1E+01	95%
Strongly Disagree									1	1.6	3.0	99%	3.1E-01	24%
Disagree									1	1.6	3.0	93%	4.4E-01	31%
	Excellent			Excellent					2	1.4	2.6	97%	3.8E-02	4%
	Excel									1.3	2.5	96%	3.4E-02	3%
										1.3	2.5	94%	6.8E-02	6%
	Excel									1.3	2.5	94%	6.0E-02	6%
										1.3	2.5	93%	8.5E-02	8%
										1.3	2.5	91%	1.2E-01	11%
	Very									1.3	2.5	92%	9.8E-02	9%
	Excel									1.3	2.5	92%	7.7E-02	7%
	Excel									1.3	2.5	91%	1.1E-01	10%
	Excel									1.3	2.5	86%	2.2E-01	18%
										1.3	2.5	91%	8.0E-02	7%
Neither Agree nor Disagree										1.3	2.4	91%	4.7E-02	4%
										1.3	2.4	83%	2.4E-01	19%
										1.3	2.4	89%	7.0E-02	7%
	Very									1.2	2.4	84%	1.7E-01	14%
										1.2	2.4	84%	1.4E-01	12%
										1.2	2.3	83%	1.7E-01	14%

Interpretation [edit]

Wikipedia: "Bayes Factor: Interpretation"

only considers evidence *against* it. Harold Jeffreys gave a scale for interpretation of *K*:[9]

K	dHart	bits	Strength of evidence
$< 10^0$	0	—	Negative (supports M_2)
10^0 to $10^{1/2}$	0 to 5	0 to 1.6	Barely worth mentioning
$10^{1/2}$ to 10^1	5 to 10	1.6 to 3.3	Substantial
10^1 to $10^{3/2}$	10 to 15	3.3 to 5.0	Strong
$10^{3/2}$ to 10^2	15 to 20	5.0 to 6.6	Very strong
$> 10^2$	> 20	> 6.6	Decisive

MRE for Low-Profitability Customers (P(E)=47%)

Analysis Context		
povrll	Profit: OVERALL	I {1: 100.00%, 2: 100.00%, 3: 0.00%, 4: 0.00%, 5: 0.00%}

Best Solutions

Attitude: Just Wheels (ajstwhls)	Satisfaction, Service: Value (qsvlu)	Depts. Used: Sales (ssales)	Vehicle Purchased: Basic SUV (vbscsuv)	Satisfaction, Vehicle: Value (wvvlu)	Demog.: HH Size (dhhsize)	Attitude: Comfy Me (acomfyme)	Relationship: Exclusivity (rxclsv)	Attitude: Fast & Furious (afstnfrs)	MRE size	Weight of Evidence	Generalized Bayes Factor	Likelihood P(e|h)	Posterior Odds O(h|e)	Posterior Probability P(h|e)
Strongly Agree									1	2.2	4.5	93%	2.7E+00	73%
							low		1	1.6	2.9	79%	1.9E+00	65%
					3				1	1.2	2.4	74%	1.4E+00	58%
		No							1	1.2	2.2	59%	3.9E+00	79%
			Yes						1	1.1	2.2	82%	6.0E-01	38%
						Strongly Disagree		Agree	2	1.1	2.1	100%	5.8E-03	1%
						Neither Agree nor Disagree		Strongly Disagree	2	1.1	2.1	100%	7.0E-04	0%
	Poor								1	1.1	2.1	91%	2.1E-01	17%
	Fair								1	1.0	2.0	77%	6.6E-01	40%
				Poor					1	0.9	1.9	86%	7.8E-02	7%
						Neither Agree nor Disagree		Neither Agree nor Disagree	2	0.9	1.8	82%	1.4E-01	12%
				Fair					1	0.8	1.8	72%	4.5E-01	31%
						Disagree		Agree	2	0.7	1.6	72%	1.1E-01	10%
								Neither Agree nor Disagree	1	0.6	1.5	64%	5.4E-01	35%
								Disagree	1	0.4	1.4	61%	2.4E-01	20%
				Good		Neither Agree nor Disagree			2	0.3	1.3	59%	1.4E-01	12%
				Good					1	0.3	1.3	55%	6.1E-01	38%
						Disagree			1	0.2	1.2	54%	2.0E-01	17%
						Neither Agree nor Disagree			1	0.2	1.1	51%	4.9E-01	33%
	Good					Agree			2	0.1	1.1	50%	1.3E-01	11%

Conclusions

- MRE is useful when you need to seek insights into **"Why?"**
  - Must carefully consider what assertions to make for evidence scenarios E and what candidate variables to be included in explanatory hypotheses H.
    (Also, useful for "Least Representative Hypotheses": i.e., what's atypical or anomalous wrt the evidence.)
  - Always be sure to also consider how probable the hypothesis is, P(H|E), and the likelihood P(E|H).

- MRE is based upon the Generalized Bayes Factor GBF(H;E)
  - Connections to Information Theory (A. Turing & I.J. Good: KLD); to optimal learning in Cognitive Science (T. Griffith & J. Tenenbaum); to philosophy of science (B. Fitelson)

- Modeling well is essential – leverage **prior knowledge & representative data**
  - The stronger your Bayesian belief network (BBN) model, the stronger your inferences & insights into interventions → i.e., new policies.

- ***BayesiaLab strengthens its position as the leading Bayesian Machine Learning environment available!***