



Vector Autoregressive Modeling of Longitudinal Sales Data using Simulated Populations Informed by Conjoint Experiments

Kevin Lattery VP Methodology & Innovation | Sawtooth Software Conference May 2022

Thanks for Allowing Talk

- Not representative of typical topic or directly relevant for Sawtooth Software
- We show a **custom HB model** with **conjoint** component
 - Sawtooth Software uses a different fixed HB model
 - Our model is a choice model with similarities to a conjoint simulator with softmax (or logit rule) $\frac{e^u}{\sum e^u}$
- **Estimation in Stan (open source software for general HB modeling)**
 - *Hamiltonian Monte Carlo* (Stan) more robust than *Gibbs Sampling* (Sawtooth Software)
 - Especially with many parameters
 - Stan faster using Linux + Arm chip + multi-threading (32 cores x 2 chains)
 - Amazon Web Services c6g.8xlarge with SKIM custom built image

Testament to Sawtooth Software's support for HB and larger analytics community, beyond their own software



Objective : Use Longitudinal Sales Data
and Conjoint for Better Predictions

| Conjoint or Real World?

- Our clients (like many) trust their real world data more than survey data.
 - Want us to use real world data to develop predictive model.
- They also believe (along with academics) consumers in **real** world behave somewhat differently than the **experimental** (conjoint survey) world.
 - Adjusting scale factor is one step (tuning conjoint results)

However, there is no reason to expect the alternative specific constants and the scale factor to be the same for stated preference data [conjoint] as for revealed-preference data [sales data]. These parameters reflect the impact of unobserved factors, which are necessarily different in real choice situations than hypothetical survey situations.

Kenneth Train

Discrete Choice Methods with Simulation (March 2002)

Pg 177, section 7.2

Vector Autoregression Often Helps Predictions

Our clients (like many) have their own forecasting models

- Models based on **aggregate** sales data for their products and competitive products over time
 - 2 – 5 years of data (monthly and some weekly)
 - Also have data like price, distribution, etc.
- Client models use previous sales to predict future sales, Vector Autoregression (VAR)
 - Previous sales are a great predictor of future sales
 - Difficult for non-VAR model to perform as well as VAR model
 - Correlations of .9 or higher, even for 12 months
 - Mean Absolute Error for family rolled-up SKUs about .15% - .30%



Vector Autoregression (VAR): Use Shares to Predict Shares

- **Vector**: At any specific time we have sales data for multiple items (100-200 in our case)
- **Autoregression**: use previous sales to forecast new sales

Previous Shares						Future Shares	
SKU	Time 1	Time 2	Time 3	Time 4	Time 5	Future 1	Future 2
1	20%	18%	17%	17%	16%	?	?
2	30%	31%	31%	32%	32%	?	?
3	10%	11%	12%	12%	13%	?	?
4	15%	15%	15%	16%	15%	?	?
5	25%	25%	25%	23%	24%	?	?



Optional Additional Exogenous Variables

Price

Distribution

...

| Our Goal

- Model choices (VAR is linear)
- Framework should be a microlevel model of individual consumer dynamics/choices*
 - Respondent level conjoint is example of microlevel model
 - Model tradeoffs and detailed dynamics

Can we somehow combine conjoint experiments
with real world longitudinal data into a VAR +
Conjoint model?

And can we get even more accurate predictions
than either alone?



Modeling Real World Changes with Simulated Shoppers

Multinomial Probit (MNP): Importance of ϵ_i for Sourcing Dynamics

$$\begin{aligned} U_1 &= \text{Int}_1 + \beta \text{price}_1 * \log(\text{price}_1) + \beta \text{dist}_1 * \log(\text{dist}_1) + \epsilon_1 \\ U_2 &= \text{Int}_2 + \beta \text{price}_2 * \log(\text{price}_2) + \beta \text{dist}_2 * \log(\text{dist}_2) + \epsilon_2 \\ &\dots \\ U_n &= \underbrace{\text{Int}_n + \beta \text{price}_n * \log(\text{price}_n) + \beta \text{dist}_n * \log(\text{dist}_n)}_{\mu_i} + \epsilon_i \end{aligned}$$

Where $\{\epsilon_i\} \sim \text{MVN}(\mathbf{0}, \Sigma)$

Not IID, covariance Σ controls sourcing

$$\{\mu_i + \epsilon_i\} \sim \text{MVN}(\mu_i, \Sigma)$$

The alternative is chosen with
the highest total utility among
 $U_i = \{\mu_i + \epsilon_i\}$

Prob of choosing has
No closed form
**Must simulate to
get probabilities
from (aggregate)
utilities**

MNP with Smoothed Accept-Reject (A-R) Simulator

- 1) Simulate a population with utilities representing unknown error $\{\epsilon_i\} \sim \text{MVN}(\mu=0, \Sigma)$
- 2) For each simulated shopper in 1)
 - a) add the known/global part μ_i to their ϵ_i
 - b) **compute predictions for simulated shopper**
- 3) Average the predictions

Kenneth Train:
Discrete Choice Methods with Simulation

5	Probit	111
5.1	Choice probabilities	111
5.2	Identification	114
5.3	Taste variation	121
5.4	Substitution patterns/non-IIA	123
5.5	Panel data	126
5.6	Simulation of the choice probabilities	130
5.6.1	Accept-reject simulator	131
5.6.2	Smoothed A-R simulators	136
5.6.3	GHK simulator	139

2b) Can compute predictions w/many methods.
First choice is most simple.

Smoothed A-R simulator can use scaled softmax (logit rule) with λ in $(0,1] = 1$:

$\frac{e^{U/\lambda}}{\sum e^{U/\lambda}}$ Requires less sim population than first choice
Avoids small items getting 0 choice

A-R MNP Simulator \approx Conjoint Simulator

In a respondent level conjoint simulator we can think of utilities as a global μ plus respondent level deviations ϵ_i

- Compute the mean vector μ across typical conjoint respondent utilities
- Subtract μ from each respondent's utilities
 - Result is conjoint respondent's deviations from the mean $\epsilon_i \sim \text{MVN}(\mu=0, \Sigma)$

To simulate, add μ to each row of ϵ_i (call this U)

- Prediction for each respondent is $\frac{e^{U_x}}{\sum e^{U_x}}$

Smoothed Accept-Reject simulator behaves like conjoint simulator and enables us to build microlevel models of choice



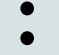

Our simulator differs from conjoint because it varies μ over time based on a specific time period

Multinomial Probit For Any Two Periods

What do Simulated Shoppers buy in this Period (Lag)?

For each row in ϵ_i add

$$\mu_{i,\text{lag}} = \text{int}_i @ \text{lag} + \beta_{\text{price}_i} * \log(\text{price}_i @ \text{lag}) + \beta_{\text{dist}_i} * \log(\text{dist}_i @ \text{lag})$$

	SKU1	SKU2	...	SKU _n
				
				
				
				

Simulated Shopper Utilities
 $\epsilon_i \sim \text{MVN}(\mu=0, \Sigma)$

What do Simulated Shoppers Buy in this Period (Forecast)?

For each row in ϵ_i add

$$\mu_{i,\text{new}} = \text{int}_i @ \text{new}_i + \beta_{\text{price}_i} * \log(\text{price}_i @ \text{new}) + \beta_{\text{dist}_i} * \log(\text{dist}_i @ \text{new})$$

Reminder:

Exact same errors ϵ_i (simulated shoppers)

Lag and forecast have different mean vectors μ_i that we add to each simulated shopper in ϵ_i

Difficulty: Single **intercept** term for each sku does not work well, as it likely varies over time. We will model intercept autoregressively





Use Known Share of Lag Period to Drop Intercept and Simplify

What do Simulated Shoppers buy in Lag Period?

$$\begin{aligned}\mu_{i,\text{lag}} &= \text{int}_i @ \text{lag} + \\ &\beta_{\text{price}_i} * \log(\text{price}_i @ \text{lag}) + \\ &\beta_{\text{dist}_i} * \log(\text{dist}_i @ \text{lag}) \\ &\approx \log(\text{lag_share}_i)\end{aligned}$$

Implies

$$\begin{aligned}\text{int}_i @ \text{lag} &\approx \log(\text{lag_share}_i) \\ &- \beta_{\text{price}_i} * \log(\text{price}_i @ \text{lag}) \\ &- \beta_{\text{dist}_i} * \log(\text{dist}_i @ \text{lag})\end{aligned}$$

	SKU1	SKU2	...	SKU _n
				
	Simulated Shopper Utilities $\epsilon_i \sim \text{MVN}(\mu=0, \Sigma)$			
				
				

What do Simulated Shoppers Buy in Forecast Period?

$$\begin{aligned}\mu_{i,\text{new}} &= \text{int}_i @ \text{lag} + \beta_{\text{trend}_i} * [\text{time diff}] + \\ &\beta_{\text{price}_i} * \log(\text{price}_i @ \text{new}) + \\ &\beta_{\text{dist}_i} * \log(\text{dist}_i @ \text{new})\end{aligned}$$

Sub for $\text{int}_i @ \text{lag}$





$$\begin{aligned}\mu_{i,\text{new}} &= \log(\text{lag_share}_i) \\ &- \beta_{\text{price}_i} * \log(\text{price}_i @ \text{lag}) \\ &- \beta_{\text{dist}_i} * \log(\text{dist}_i @ \text{lag}) + \\ &\beta_{\text{trend}_i} * [\text{gap periods}] + \\ &\beta_{\text{price}_i} * \log(\text{price}_i @ \text{new}) + \\ &\beta_{\text{dist}_i} * \log(\text{dist}_i @ \text{new})\end{aligned}$$

Final Model: Paired MNP and VAR of Differences

What do Simulated Shoppers buy in Lag Period?

1

For each row in ϵ_i add
 $\mu[\text{sku}_i]_{\text{lag}} =$
 $\log(\text{share}[\text{sku}_i]_{\text{lag}})$

	SKU1	SKU2	...	SKU _n
				
				
				
				

Simulated Shopper Utilities
 $\epsilon_i \sim \text{MVN}(\mu=0, \Sigma)$

What do Simulated Shoppers Buy in Forecast Period?

For each row in ϵ_i add

$\mu[\text{sku}_i]_{\text{new}} =$
 $\log(\text{share}[\text{sku}_i]_{\text{lag}}) + \beta_{\text{trend}_i} * [\text{time diff}] +$
 $\beta_{\text{price}_i} * [\log(\text{price}_i @ \text{new}) - \log(\text{price}_i @ \text{lag})] +$
 $\beta_{\text{dist}_i} * [\log(\text{dist}_i @ \text{new}) - \log(\text{dist}_i @ \text{lag})]$

2

Fit difference in log(shares) for simulated vs observed:

$$\log(\text{Sim}_{\text{Fore}}) - \log(\text{Sim}_{\text{Lag}}) \sim \log(\text{Obs}_{\text{Fore}}) - \log(\text{Obs}_{\text{Lag}})$$

$$\log(\text{Sim}_{\text{Fore}}) - \log(\text{Sim}_{\text{Lag}}) + \log(\text{Obs}_{\text{Lag}}) \sim \log(\text{Obs}_{\text{Fore}})$$

$$\text{Softmax}(\log(\text{Sim}_{\text{Fore}}) - \log(\text{Sim}_{\text{Lag}}) + \log(\text{Obs}_{\text{Lag}})) \sim \text{Obs}_{\text{Fore}}$$



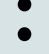

Final Model: Conceptual Explanation



What do Simulated Shoppers buy in Lag Period?

1

Calibrate Shoppers Approximately to Shares in Lag Period

	SKU1	SKU2	...	SKU _n
				
				
				
				

Simulated Shopper Utilities
 $\epsilon_i \sim \text{MVN}(\mu=0, \Sigma)$

What do Simulated Shoppers Buy in Forecast Period?




Lag Calibrated Shoppers + Changes from Lag Period

2

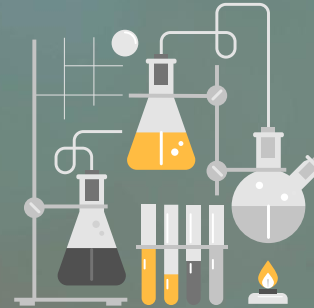
Adjust Forecast Simulations for Error in Calibration to Lag Period

$$\log(\text{Sim}_{\text{Fore}}) + \{\log(\text{Obs}_{\text{Lag}}) - \log(\text{Sim}_{\text{Lag}})\} \sim \log(\text{Obs}_{\text{Fore}})$$



	SKU1	SKU2	...	SKU _n
				
	Simulated Shopper Utilities $\epsilon_i \sim \text{MVN} (\mu=0, \Sigma)$			
•				
•				
				

ϵ_i as Microlevel of Consumer Dynamics: Fusion with Conjoint Experiments



| Why Conjoint can Complement Real World Data

Stated-preference data [conjoint] complement revealed-preference data [sales data].

...

The advantage of stated-preference data is that the experiments can be designed to contain as much variation in each attribute as the researcher thinks is appropriate. While there might be little price variation over suppliers in the real world, the suppliers that are described in the experiments can be given sufficiently different prices to allow precise estimation. Attributes can be varied over respondents and over experiments for each respondent. This degree of variation contrasts with market data, where often the same products are available to all customers, such that there is no variation over customers in the attributes of products. Importantly, for products that have never been offered before, or for new attributes of old products, stated-preference data allow estimation of choice models when revealed-preference data do not exist.

Kenneth Train



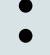

Discrete Choice Methods with Simulation

Pg 175, section 7.2

** Text in this color added (Kevin's interpretation)*

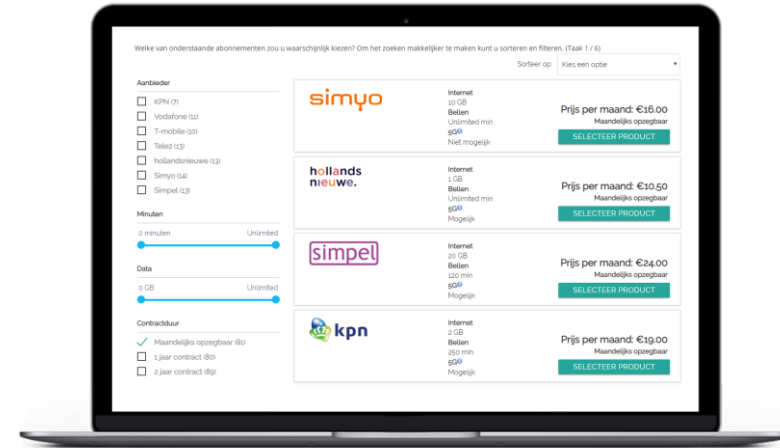
| How to Estimate Σ ?

- Our microlevel choice model depends upon generating simulated shoppers from **MVN** ($\mu=0, \Sigma$)
- With our specific data and hundreds of SKUs, estimating Σ **with just sales data** was problematic
 - Sourcing results were close to IIA (points or draws)
 - The posterior of correlation implied by Σ varies around 0 (with a range that depends on prior)
 - With many products, estimating Σ is one of the known difficulties for MNP

	SKU1	SKU2	...	SKU _n
				
	Simulated Shopper Utilities			
				
				

Conjoint Experiments Can Inform Σ



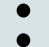

- Conjoint Experiments allow us to observe changes at the respondent level to different stimuli in a designed experiment
- Respondents move from one SKU to another because of changes in price or product availability.
- This **respondent level sourcing** information allows us to estimate the covariance of utilities Σ for our conjoint population.



- We conducted 5 different conjoint studies with overlapping SKUs in the same category
 - Total of 14,341 respondents

Linking Simulated Shoppers and Conjoint Respondents

MVN (α, Σ^*)

	SKU1	SKU2	...	SKU _n
				
				
				
				

Conjoint Respondent Utilities

Simulated shopper utilities in real world are similar to our conjoint respondents






Some Possible Relationships

- 1) $\Sigma = k \Sigma^*$
- 2) $\Sigma \sim \Sigma^*$
- 3) $R = R^*$ Where R is the correlation matrix
- 4) $R \sim R^*$ corresponding to covariance $\Sigma = DRD'$

Academic Note:

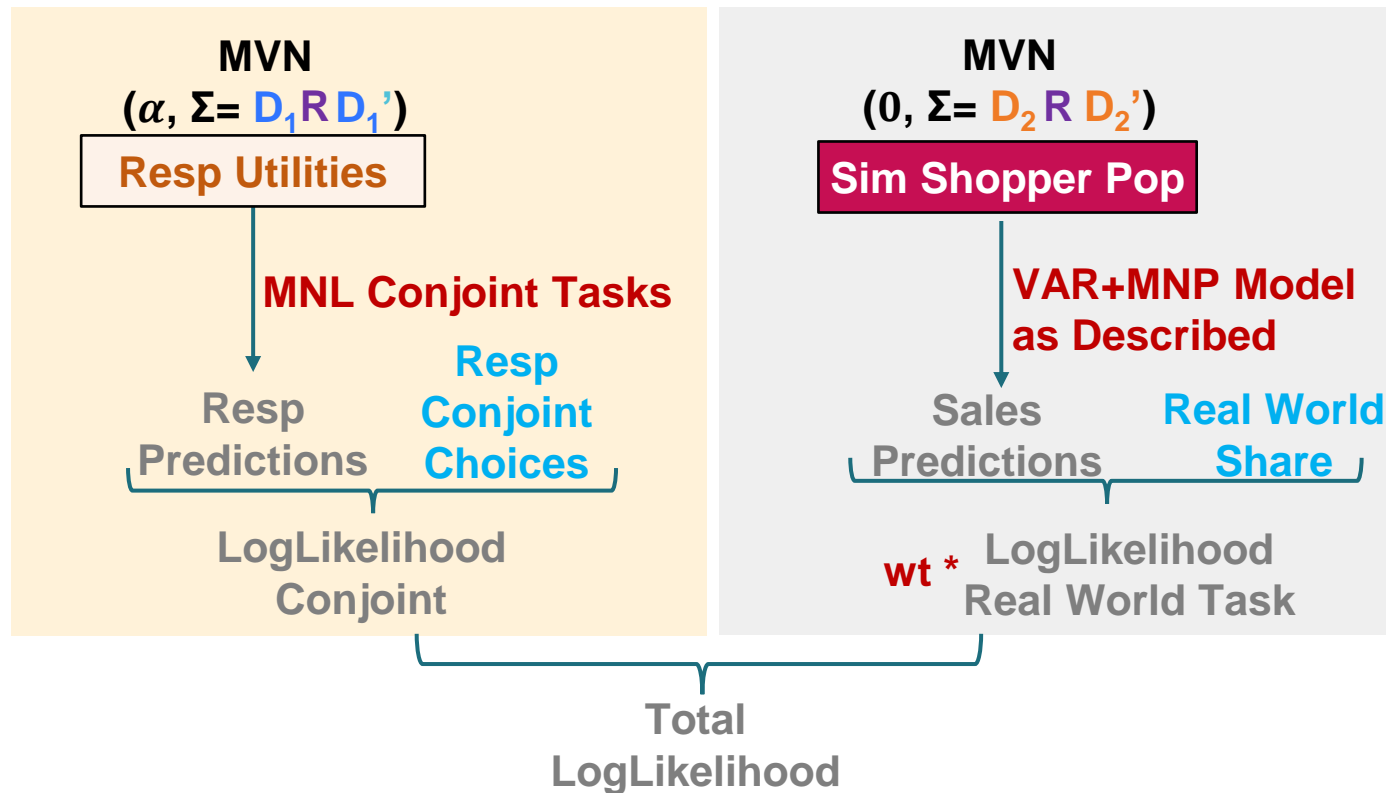
This is an extension of the idea of building conjoint simulators that “*simulate from the upper level model*”.

$\epsilon_i \sim \text{MVN} (\mu=0, \Sigma)$

	SKU1	SKU2	...	SKU _n
				
				
				
				
				

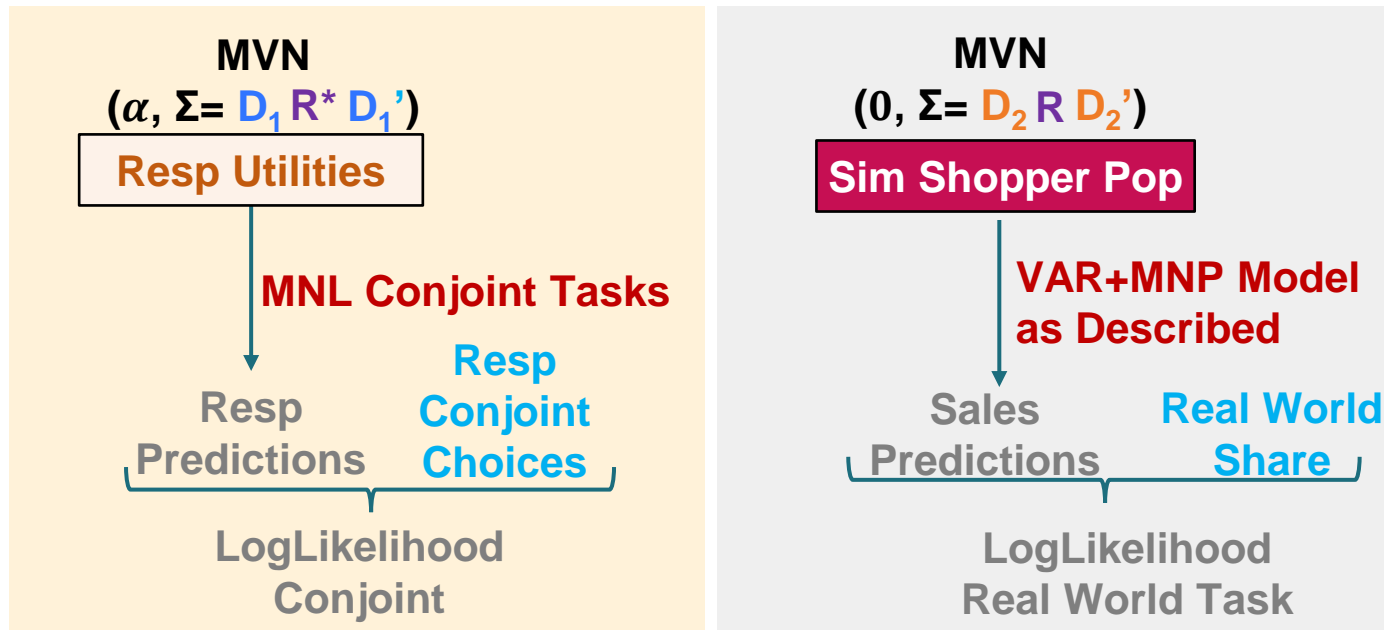
Simulated Shopper Utilities

Simultaneous Estimation of Sales and Conjoint with Shared Correlation R



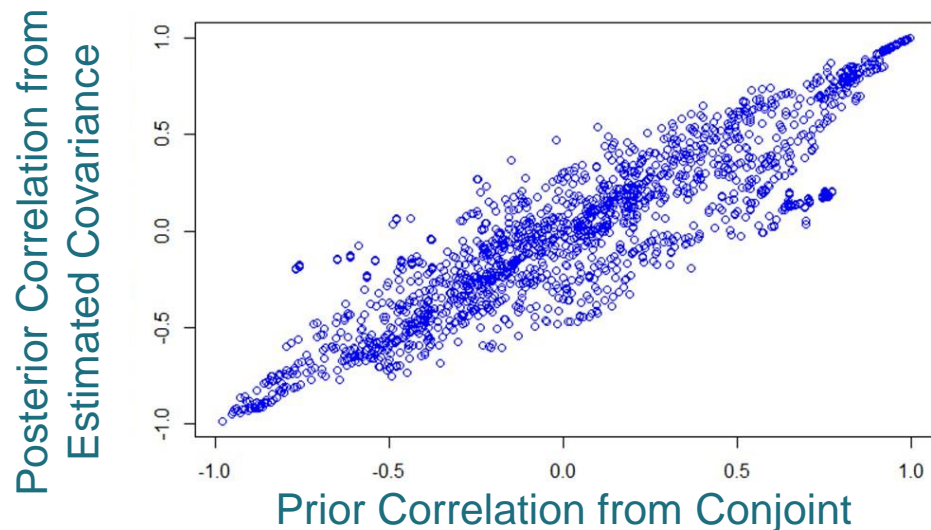
Use Correlation of Conjoint Utilities as Prior

$$R \sim R^*$$



No more combined loglikelihood or weighting
Much faster to estimate than both together

Correlation implied by Posterior Σ is Similar to Prior, Even when Conjoint is Vague Prior



Estimating Σ without conjoint prior (= 0 correlation) gives posterior correlation near 0

$\Sigma \sim \text{Wishart}(\text{DF} = P+2, \text{Scale} = \sigma_{\text{diag}} * \text{prior_corr} * t(\sigma_{\text{diag}})/(P+2))$, where
prior_corr is fixed prior correlation from conjoint,
P is number of columns (or rows) in prior_corr,
 σ_{diag} is diagonal of standard deviations (estimated parameter)
Note: in Stan we sample from Wishart using Bartlett Decomposition



Results & Future Steps

History of Main Methods and Overall Success

Holdout MAE
SKUs Rolled to Family
20 = .2%

Method	Covariance Σ	Sourcing	1Mo	3Mo	6mo
MNP Only	Estimated from Sales Data	IIA sourcing	18.3	26.1	40.4
MNP + Conjoint	Used Conjoint Respondents	Too much sourcing	18.0	27.8	42.2
MNP + Conjoint Joint	Sales and Conjoint Jointly	Best, but long time	18.2	23.9	31.9



Much better predictions with VAR

Simple VAR	None	Volume to Share IIA	12.7	19.9	25.5
MNP + VAR	Estimated from Sales Data	IIA sourcing	12.5	15.7	19.2
MNP + Conjoint + VAR	Used Conjoint Respondents	Too much sourcing	12.5	17.2	19.6
MNP + Conjoint Joint + VAR	Sales and Conjoint Jointly	Best, but long time	12.3	13.9	14.8
MNP + Conjoint Prior + VAR	Conjoint Prior	Similar to Above, Less Time	12.5	13.8	15.0






VAR model more accurate than Non-VAR model

Infer intercept utility of sku from $\log(\text{share}[\text{sku}_i]_{\text{lag}})$ and model differences from lag to forecast periods

Difference of Obs ~ Difference of Simulated

What do Simulated Shoppers buy in Lag Period?

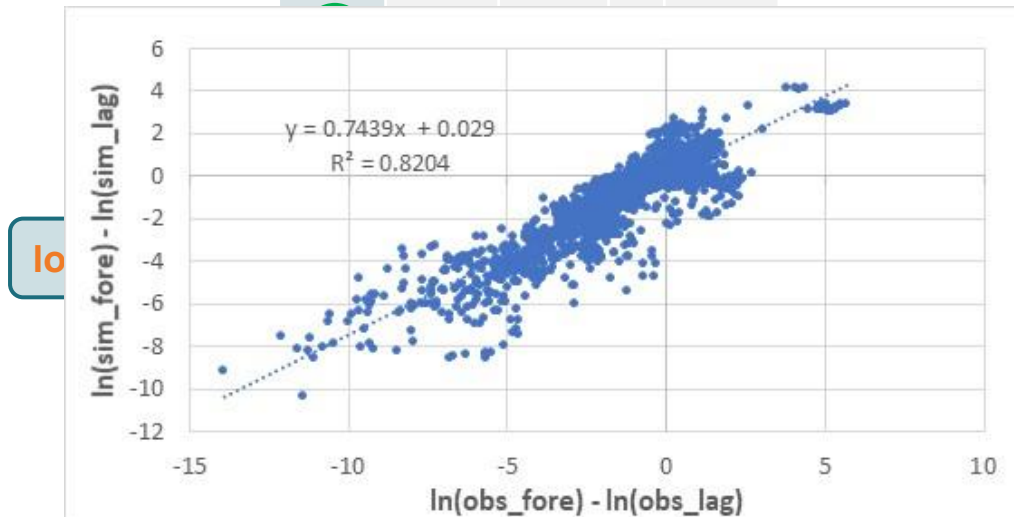
Calibrate Shoppers Approximately to Shares in Lag Period

	SKU1	SKU2	...	SKU _n
				
				
				
				
				

Simulated Shopper Utilities
 $\epsilon_i \sim \text{MVN}(\mu=0, \Sigma)$

What do Simulated Shoppers Buy in Forecast Period?






Lag Calibrated Shoppers + Changes from Lag Period



Difference of Obs ~ Difference of Simulated

What do Simulated Shoppers buy in Lag Period?

Calibrate Shoppers Approximately to Shares in Lag Period

	SKU1	SKU2	...	SKU _n
				
	Simulated Shopper Utilities $\epsilon_i \sim \text{MVN}(\mu=0, \Sigma)$			
				
				
				

What do Simulated Shoppers Buy in Forecast Period?

Lag Calibrated Shoppers + Changes from Lag Period

Adjust Forecast Simulations for Error in Calibration to Lag Period

$$\log(\text{Sim}_{\text{Fore}}) + \{\log(\text{Obs}_{\text{Lag}}) - \log(\text{Sim}_{\text{Lag}})\} \sim \log(\text{Obs}_{\text{Fore}})$$

What if Sim_{Lag} is calibrated exactly to Obs_{Lag} ?

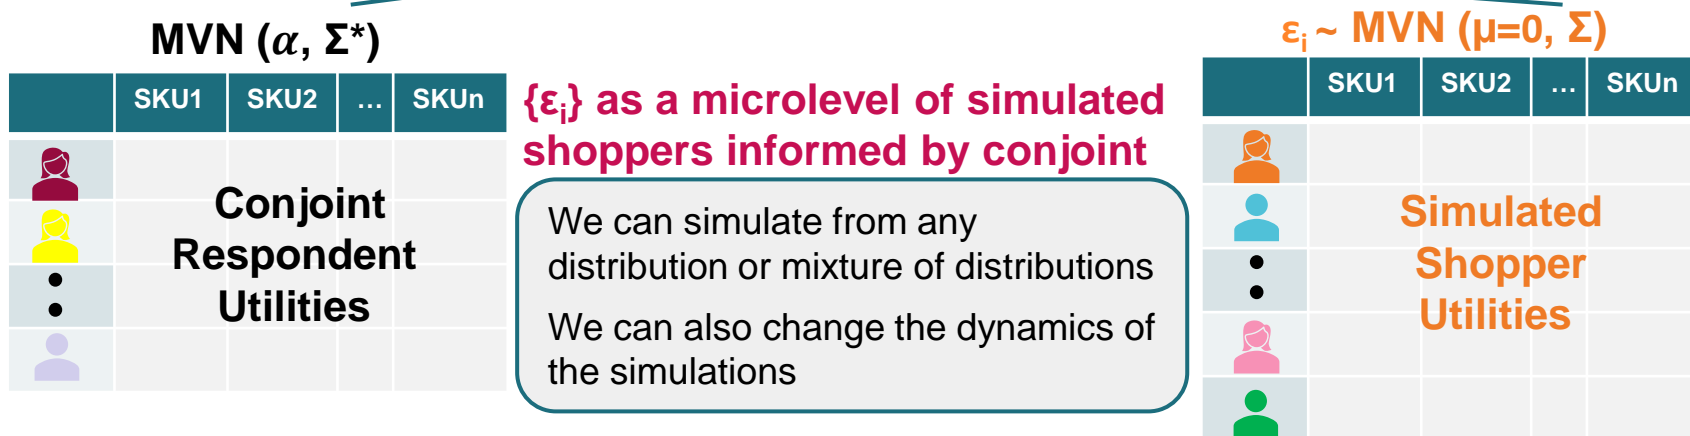
Post-Hoc Calibrate/Tune Intercept of Non-VAR

Holdout MAE
SKUs Rolled to Family
20 = .2%

Method	Covariance Σ	Sourcing	1Mo	3Mo	6mo
MNP Only	Estimated from Sales Data	IIA sourcing	18.3	26.1	40.4
MNP + Conjoint	Used Conjoint Respondents	Too much sourcing	18.0	27.8	42.2
MNP + Conjoint Joint	Sales and Conjoint Jointly	Best, but long time	18.2	23.9	31.9
MNP Only + Tune	Estimated from Sales Data	IIA sourcing	12.7	18.9	22.2
MNP + Conjoint + Tune	Used Conjoint Respondents	Too much sourcing	12.9	19.2	23.1
MNP + Conjoint Joint + Tune	Sales and Conjoint Jointly	Best, but long time	12.6	17.5	19.9
Simple VAR	None	Volume to Share IIA	12.7	19.9	25.5
MNP + VAR	Estimated from Sales Data	IIA sourcing	12.5	15.7	19.2
MNP + Conjoint + VAR	Used Conjoint Respondents	Too much sourcing	12.5	17.2	19.6
MNP + Conjoint Joint + VAR	Sales and Conjoint Jointly	Best, but long time	12.3	13.9	14.8
MNP + Conjoint Prior + VAR	Conjoint Prior	Similar to Above, Less Time	12.5	13.8	15.0

Post-Hoc
Tune Intercept

Data Fusion with Conjoint Utilities: an Extendable Framework

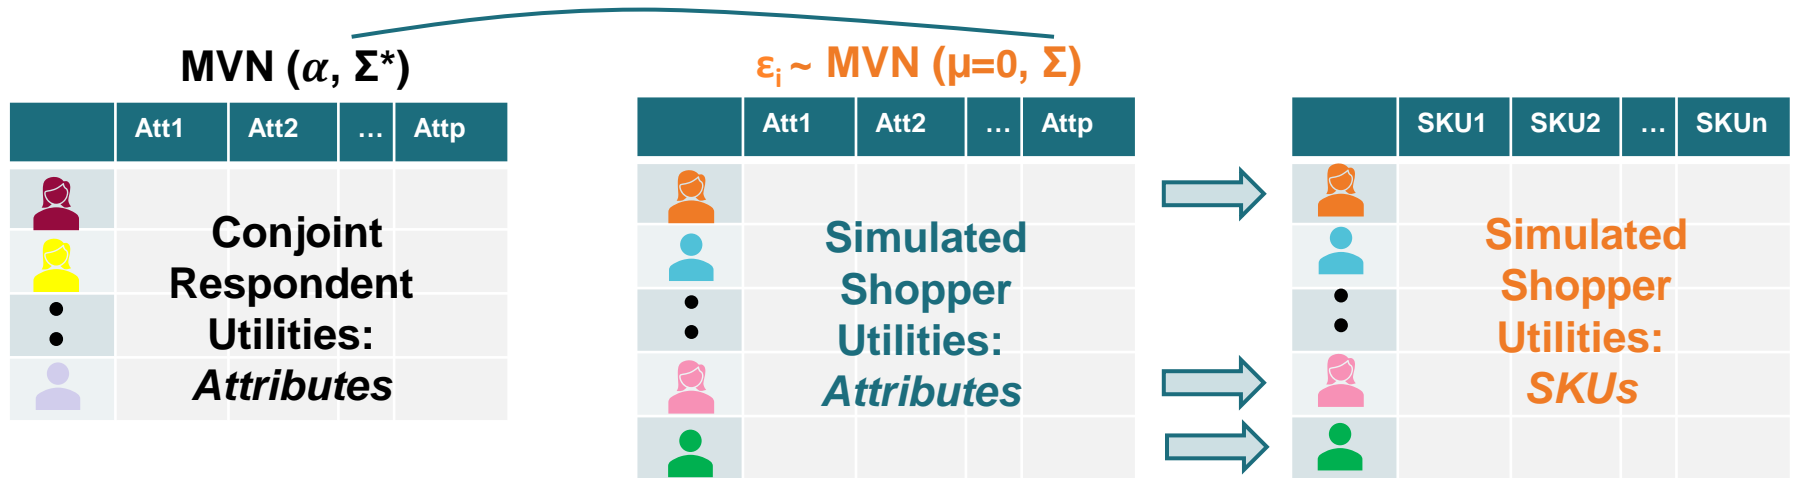


- **Conjoint provides information about our simulated shoppers**
- **Two-stage estimation using conjoint correlation as prior worked about as well as simultaneous estimation (in our studies)**

Benefits of using conjoint as prior:

- No need to weight sales data vs conjoint
- Much quicker to estimate and update parameters when we get more data

Actual Extension Used Attributes and Compute SKUs



From each simulated shopper's attribute utilities $\epsilon_{i,\text{Att}}$ we compute the corresponding utility of each SKU (based on its attribute definition), resulting in $\epsilon_{i,\text{SKU}}$

Conclusion: Experiments + Real World Data = Better Predictions

- **Conjoint Experiments can inform microlevel of simulated shoppers that help predict aggregate shares**



- **We described a framework that can be further expanded and modified**
- **Welcome more discussion on data fusion linking different data sources**

Questions?



Kevin Lattery

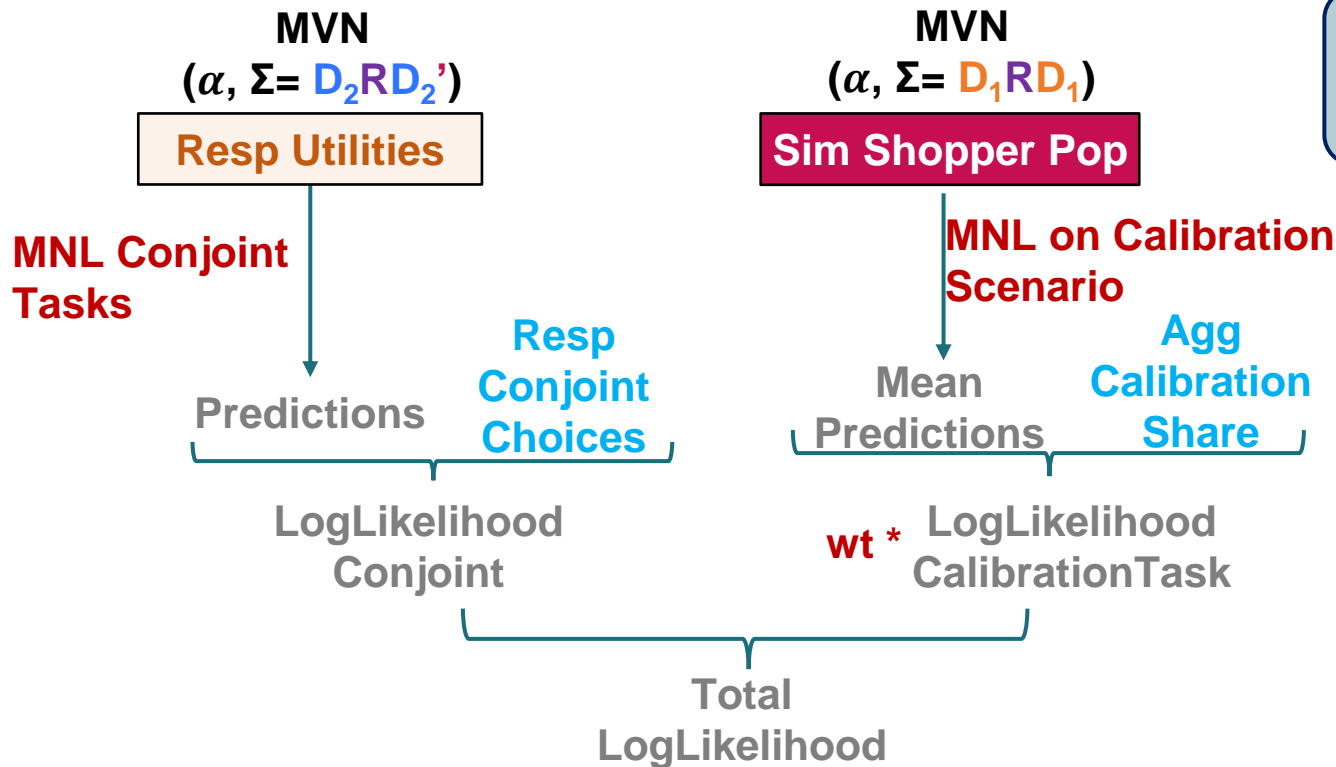
email: k.lattery@skimgroup.com
www.skimgroup.com
[@skimgroup](https://twitter.com/skimgroup)

- 1) How many simulated respondents do you need?
- 2) What if you have real world SKUs without corresponding conjoint priors?
- 3) What are some of the other microlevel consumer dynamics you modeled?
- 4) Could we use something similar (but simpler) than this to fuse conjoint with sales data for just one time period (calibrate conjoint data to sales)?



Appendix

Simultaneous Estimation Similar for “Static Calibration Data”



We can also set:
Sim Shopper Pop =
Respondent Utilities