Dynamics Recovery of Individual-level Fusion Trajectories from Surveys

Samuel Levy[1] Longxiu Tian[2]

[1]University of Virginia, Darden School of Business
[2]University of North Carolina, Kenan-Flagler Business School

Marketing Dynamics Conference 2026
ESSEC Business School
May 26–28, 2026

Customer Surveys: A Fundamental Tension

Customer surveys are ubiquitous:

- Gauge satisfaction, identify churn drivers
- Understand competitive positioning
- Foundation of CRM strategy

But practical realities create friction:

- Most surveys are **Repeated Cross-Sections (RCS)**
- Different respondents each wave
- Cannot track individual evolution

The Problem

"Pseudo-longitudinal" data

Snapshots of distinct customers at different moments

$\Downarrow$

Cannot separate heterogeneity from dynamics

Cannot forecast individual trajectories

The Data Structure Challenge

Repeated Cross-Sections (RCS)

- \+ Broad population coverage
- \+ Large sample sizes
- \+ Low respondent burden
- – No individual continuity
- – Cannot separate heterogeneity from dynamics

Panel Surveys

- \+ Track individual trajectories
- \+ Identify within-person variation
- \+ Separate static vs. dynamic effects
- – Expensive, high attrition
- – Limited coverage

Key Insight: Can we *fuse* complementary data sources to get the best of both worlds?

An Analogy: Incomplete Information

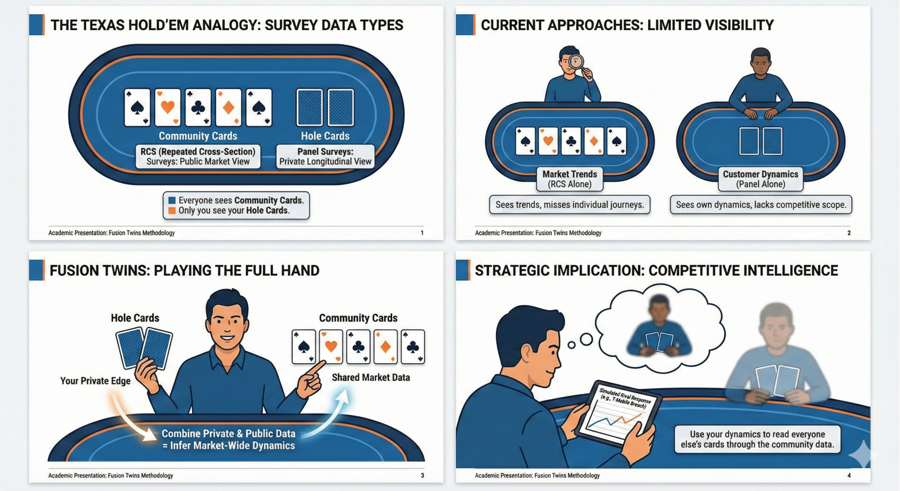

Research Questions

RQ1 Can marketers **augment pseudo-longitudinal RCS surveys** with parsimonious panel surveys to enable inference on customer dynamics and heterogeneity?

RQ2 How well can a fusion of sparse panel and dense RCS surveys **predict individual-level attitudes** out-of-sample, relative to state-of-the-art benchmarks (BART, cGAN, CVAE, GP, hierarchical cohort)?

RQ3 What are the **data requirements** for such a fusion exercise — how thin can the panel be, and how few RCS waves are needed, to recover individual-level trajectories?

Related Literature

Pseudo-Panel Econometrics

- Deaton (1985); Verbeek & Nijman (1992); Moffitt (1993)
- Enabled dynamics from RCS
- – Pools over individual heterogeneity
- – Attenuation bias from sampling noise

Marketing RCS Methods

- Du & Kamakura (2015): Dynamic factor model
- Denoises RCS, extracts common trends
- – No causal identification

Digital Twins (Engineering)

- Tsialiamanis et al. (2021); Kapteyn et al. (2021)
- Virtual replicas for scenario testing
- – Not adapted to multi-firm CRM
- – No causal/strategic framework

Data Fusion

- Gilula et al. (2006)
- Fuse datasets from same population
- ✓ Foundation for our approach

Our Approach: DRIFTS

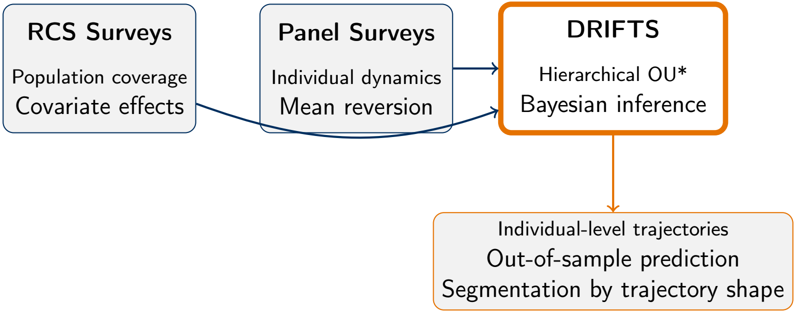

Important Point

Respondent pools *need not share individuals*—only drawn from same population

*OU = Ornstein-Uhlenbeck process: a mean-reverting stochastic differential equation

Contributions

Methodological

- Extend pseudo-panel techniques
- Principled RCS + panel fusion
- Dynamically coherent framework
- Managerially interpretable

Substantive

- Short-term shocks dissipate
- Strategic improvements persist
- Reframes tactical vs. strategic resource allocation

Managerial

- Surveys $\rightarrow$ forecasting engines
- Individual-level forecasts
- Cross-carrier monitoring without rival panel data

Bottom Line

Transform customer surveys from *descriptive dashboards* into *prescriptive decision engines*

What We Recover: Individual-Level Latent Trajectories

The Deliverable:

- Per-individual posterior over the latent OU state $x_i(t)$
- Full trajectory with credible bands at every t
- Recovered jointly from **sparse panel** + **dense RCS**

No shared respondents required — only that panel and RCS sample from the same population.

What it enables (prediction-first):

- **Forecast** next-period satisfaction / intent for each respondent
- **Segment** by trajectory *shape*, not snapshot value
- **Early-warning** detection on equilibrium drift

Key Capability

Each respondent — even one observed in a single RCS wave — gets a full latent trajectory, posterior-distributed.

Empirical Setting: U.S. Wireless Telecom

Industry-wide RCS Surveys

- AT&T, T-Mobile, Verizon
- Q3 2020 – Q2 2022 (10 quarters)
- ~6,756 respondents/carrier/quarter
- 220,690 total respondents

Partnering Carrier's Panel Survey

- Concurrent, identical questions
- 53,535 unique respondents
- Customer IDs tracked via email links
- .61% response rate (industry standard)

Key Outcomes

- **Satisfaction** (NPS, 0–10)
- **Intent to Switch** (1–5)

Natural Experiment

August 2021
T-Mobile Data Breach
76.6M individuals exposed

Data Validation: Sample Selection

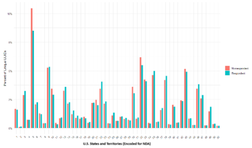

(a) State

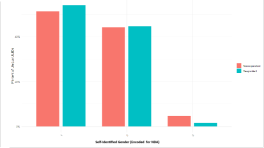

(b) Gender

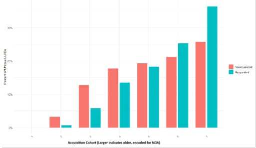

(c) Acquisition

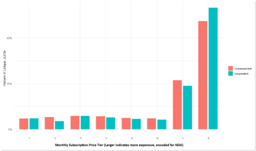

(d) Subscription

Takeaway

- Category frequencies rank-wise identical between respondents & non-respondents
- Large-scale surveys mitigate selection concerns (Huang & Sudhir, 2021)

Data Validation: Survey-CRM Covariance

Key Findings:

- 19.7% responded to $\geq$2 waves
- Only 64 (.001%) responded to all quarters
- $<$.01% returned after gap

Implication: "One-and-done" respondents
$\Rightarrow$ RCS unlikely affected by repeat obs.

5.82% churn rate in panel
(survey window + 12 months forward)

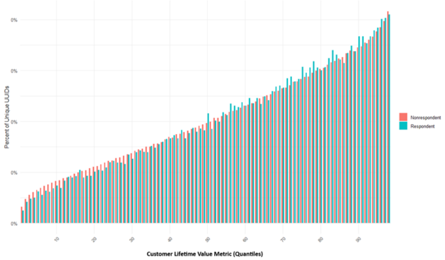

Figure: CLV: Respondents vs. Non-respondents

Modeling Philosophy: Three Sources of Variation

1. **Stable Individual Differences**
   Some customers inherently more satisfied (observed & unobserved)

2. **Dynamic Fluctuations**
   Deviations from equilibrium due to random experiences, service encounters

3. **Systematic Responses**
   Time-varying conditions (promotions, disruptions) create predictable drift

Goal

Panel data $\rightarrow$ identify trajectories, separate heterogeneity from fluctuations
RCS data $\rightarrow$ population coverage, covariate space

The Ornstein-Uhlenbeck (OU) Process

Stochastic Differential Equation

$$dx_i(t) = \underbrace{\left[-\boldsymbol{\Gamma}(x_i(t) - \boldsymbol{\mu}_i)\right]}_{\text{Mean Reversion}} dt + \underbrace{\boldsymbol{\beta} Z_i(t) dt}_{\text{Drift}} + \underbrace{\boldsymbol{\Sigma} dW_i(t)}_{\text{Innovation}}$$

Parameters:

- $x_i(t)$: latent state (satisfaction, intent)
- $\boldsymbol{\Gamma}$: reversion rate (diagonal)
- $\boldsymbol{\mu}_i$: individual equilibrium
- $\boldsymbol{\beta}$: time-varying covariate effects
- $\boldsymbol{\Sigma}$: innovation matrix

▸ Technical Details

Intuition:

- When $x_i(t) > \mu_i$: pulls down
- When $x_i(t) < \mu_i$: pulls up
- $Z_i(t)$: market forces "fight" reversion
- Half-life: $t_{1/2} = \ln(2)/\gamma$

OU Process: Visual Intuition

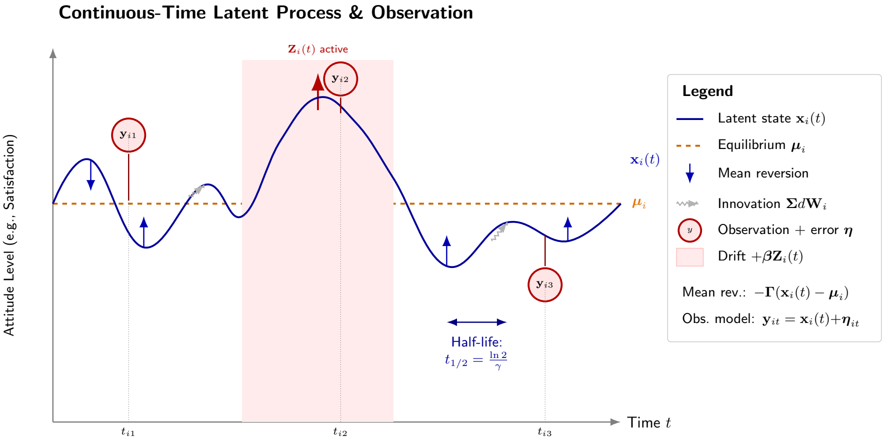

Latent state $x_i(t)$ fluctuates around individual equilibrium $\boldsymbol{\mu}_i$
Mean reversion pulls back; drift $\beta Z_i(t)$ pushes systematically

Hierarchical Equilibria

Individual Equilibrium Structure

$$\boldsymbol{\mu}_i = \underbrace{\boldsymbol{\mu}}_{\text{Population mean}} + \underbrace{\mathrm{Bw}_i}_{\text{Static covariates}} + \underbrace{\boldsymbol{\epsilon}_i}_{\text{Unobserved heterogeneity}}$$

Example: High-income customer in their 60s has equilibrium satisfaction:

$$\mu_1 + B_{1,\text{income}} \cdot w_{i,\text{income}} + B_{1,\text{age}} \cdot w_{i,\text{age}} + \epsilon_{i1}$$

w_i (Static)

Demographics, contract type, tenure, carrier

$\boldsymbol{\epsilon}_i \sim \mathcal{N}(0, \boldsymbol{\Omega})$

Unobserved loyalty, personality, past experiences

Ordinal Outcome Model

Survey responses are **ordinal categories** (e.g., satisfaction 1–10):

Ordered Logit Link

$$\Pr(y_{it,j} = k \mid \mathrm{x}_i(t)) = \Lambda(\kappa_{j,k} - x_{i,j}(t)) - \Lambda(\kappa_{j,k-1} - x_{i,j}(t))$$

where $\Lambda(z) = (1 + e^{-z})^{-1}$ and $\kappa_{j,k}$ are ordered cutpoints.

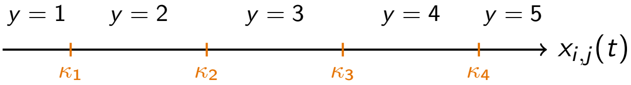

Transition Dynamics

Exact Gaussian Transitions

Given state at time s, distribution at $t > s$:

$$\mathrm{x}_i(t) \mid \mathrm{x}_i(s), \epsilon_i \sim \mathcal{N}(\mathrm{m}_{i,s,t}, \mathrm{V}_{s,t})$$

Conditional Mean:

$$\mathrm{m}_{i,s,t} = \underbrace{e^{-\mathbf{\Gamma}(t-s)}\mathrm{x}_i(s)}_{\text{Decayed starting state}} + \underbrace{(\mathrm{I} - e^{-\mathbf{\Gamma}(t-s)})\boldsymbol{\mu}_i}_{\text{Pull toward equilibrium}} + \underbrace{\int_s^t e^{-\mathbf{\Gamma}(t-\tau)}\beta \mathrm{Z}_i(\tau) d\tau}_{\text{Accumulated covariate effects}}$$

Key Property: As $t - s \to \infty$, state approaches equilibrium $\boldsymbol{\mu}_i$

▸ Transition Covariance Details

Data Fusion Likelihood

Panel Likelihood

$$\mathcal{L}_i^{\text{panel}} = \int p(\epsilon_i) \left[\prod_j p(\mathrm{x}_{ij} \mid \mathrm{x}_{i,j-1}) p(\mathrm{y}_{ij} \mid \mathrm{x}_{ij}) \right] d\epsilon_i$$

- Multiple obs. per individual
- Identifies $\boldsymbol{\Gamma}$, $\boldsymbol{\Sigma}$, $\boldsymbol{\Omega}$
- Within-person variation

RCS Likelihood

$$\mathcal{L}_{i'}^{\text{RCS}} = \int p(\mathrm{y}_{i'} \mid \mathrm{x}_{i'}) \mathcal{N}(\mathrm{x}_{i'}; \mathrm{m}^{\text{RCS}}, \mathrm{C}^{\text{total}}) d\mathrm{x}_{i'}$$

- Single observation
- Marginalizes over $\epsilon_{i'}$
- Population coverage

Joint Fusion Likelihood

$$\mathcal{L}(\Theta) = \prod_{i=1}^{n_{\textbf{panel}}} \mathcal{L}_i^{\text{panel}} \times \prod_{i'=1}^{n_{\textbf{RCS}}} \mathcal{L}_{i'}^{\text{RCS}}$$

Estimation: Fully Bayesian via HMC (NUTS)

Joint posterior sampled by the No-U-Turn Sampler:

$$p\left(\mathrm{B},\, \boldsymbol{\beta},\, \boldsymbol{\mu},\, \boldsymbol{\Gamma},\, \boldsymbol{\Sigma},\, \boldsymbol{\Omega},\, \text{cutpoints} \,\middle|\, \mathrm{y}^{\text{panel}},\, \mathrm{y}^{\text{RCS}}\right)$$

Latent states are integrated out, not sampled:

- The OU state $\mathrm{x}_i(t)$ at each panel observation is marginalized inside $\mathcal{L}_i^{\text{panel}}$
- Keeps the posterior geometry tractable for HMC — avoids gradient pathologies on a high-dimensional latent path

Implementation: NumPyro + JAX on GPU; multi-chain $\hat{R} < 1.05$ on the canonical run.

Evaluating the Panel Likelihood: Quadrature Kalman Filter

NUTS evaluates the marginal likelihood via the prediction-error decomposition:

$$\log \mathcal{L}_i^{\text{panel}}(\theta) = \sum_t \log p(y_{i,t} \mid y_{i,1:t-1},\, \theta)$$

Each factor is supplied by a **forward Gauss–Hermite Kalman filter**; latent states are never sampled.

Predict (exact — OU dynamics are linear-Gaussian):

$$\hat{x}_{k|k-1} = e^{-\Gamma\Delta t}\,\hat{x}_{k-1|k-1} + (\mathrm{I} - e^{-\Gamma\Delta t})\,\mu_i + \text{covariate drift}$$

Update (Gauss–Hermite quadrature integrates the ordered-logit at nodes $\{\xi_j, w_j\}$):

$$p(y_{i,k} \mid y_{i,1:k-1}) \approx \sum_j w_j\, p\left(y_{i,k} \,\middle|\, \hat{x}_{k|k-1} + \sqrt{P_{k|k-1}}\,\xi_j\right)$$

Trajectory recovery (after MCMC, per posterior draw $\theta^{(s)}$): a backward RTS pass reconstructs $\hat{x}_{k|T}$ — **not** in the NUTS gradient path.

What Each Data Source Identifies

Parameter	Panel	RCS
Mean reversion $\boldsymbol{\Gamma}$	✓	
Innovation $\boldsymbol{\Sigma}$	✓	
Individual heterogeneity $\boldsymbol{\Omega}$	✓	
Individual effects $\{\boldsymbol{\epsilon}_i\}$	✓	
Static effects B		✓
Drift coefficients $\boldsymbol{\beta}$		✓
Population means $\boldsymbol{\mu}$		✓
Joint estimation	*Pools information*	

Intuition

Panel: *how* individuals evolve (dynamics); RCS: *who* differs from whom (population structure)

Why Mean Reversion Requires Panel Data

From RCS alone:

- Observe cross-sectional variance
- Cannot distinguish:
  - High $\boldsymbol{\Omega}$ (heterogeneous equilibria)
  - High $\boldsymbol{\Sigma}$ (volatile dynamics)
  - Low $\boldsymbol{\Gamma}$ (slow reversion)

With panel fusion:

- Within-person autocorrelation $\to \boldsymbol{\Gamma}$
- Residual variance $\to \boldsymbol{\Sigma}$
- Between-person variance $\to \boldsymbol{\Omega}$

Key Insight

Mean reversion (γ) governs shock decay half-life

$$t_{1/2} = \frac{\ln(2)}{\gamma}$$

Unidentifiable from RCS!
Critical for forecasting decay

University of Virginia | Darden School of Business

Simulation Design

Data Generating Process:

- Randomly initialized DRIFTS
- Two ordinal outcomes per period
- 50 time periods per individual
- Panel: contiguous windows (avg. 5 periods)
- RCS: single snapshots

Configurations:

- Panel-to-RCS ratio: 0%–10%
- $N_{\text{RCS}} = 5{,}000$ (held constant)

Comparison Models

- **CVAE** (Sohn et al., 2015)
- **cGAN** "Digital Twins" (Tsialiamanis et al., 2021)
- **Cohort Model** (Deaton, 1985)
- **Gaussian Processes**
- **BART** (Chipman et al., 2012)

RMSE: DRIFTS Dominates

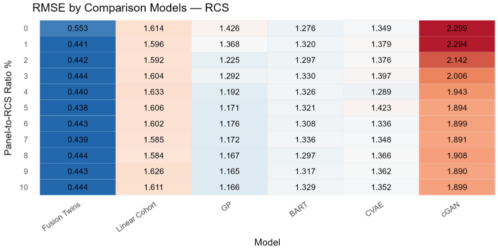

Figure: RCS Test Data

Figure: Panel Test Data

Key Finding

At 1% panel-to-RCS ratio: **25% RMSE improvement** on RCS test data
DRIFTS *continuously improves* with more panel data

Posterior Coverage: Uncertainty Quantification

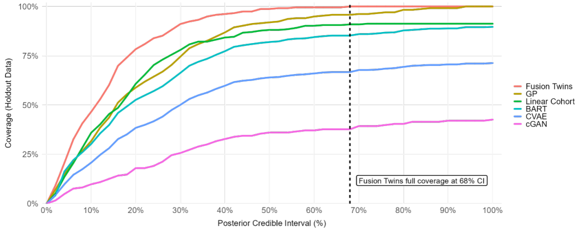

Figure: Coverage of Ground-Truth Outcomes

DRIFTS:
100% coverage at 68th percentile

Deep Generative Models:
CVAE: 73% at 99.9th
cGAN: $<$50% at 99.9th

Warning: Neural approaches severely underestimate uncertainty

Extrapolation: Stable Forecasting

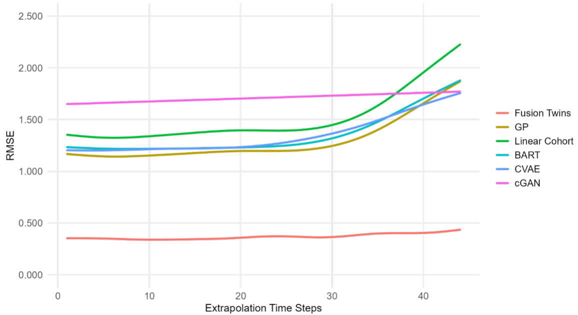

Figure: Forecast Error from Last Observed Point

DRIFTS:

Stable error across horizons (captures mean reversion)

Comparison Models:

Rising error after 30 steps

Cannot extrapolate OU dynamics

Implication

Long-horizon forecasting requires explicit equilibrium-seeking specification

Out-of-Sample Hit Rates

Model	Split	Satisfaction	Intent
Random Guess	—	10.0%	20.0%
GP	Time-contiguous	9.3	20.3
Cohort	Time-contiguous	48.5	46.6
CVAE	Time-contiguous	48.9	53.1
BART	Time-contiguous	32.7	39.2
cGAN	Time-contiguous	42.5	42.3
DRIFTS	**Time-contiguous**	**53.9**	**53.5**
DRIFTS	**Random**	**53.8**	**53.6**
DRIFTS	**Individual**	**54.1**	**53.4**

Key Finding

DRIFTS: **10.2% absolute improvement** over CVAE (next best); **remarkably consistent** across all splits (±0.6pp)

Mean Reversion Rates by Cohort

Cohort	Dimension	Posterior Mean	95% CI	Half-Life
2012–Earlier	Satisfaction	1.845	[1.251, 2.700]	0.38 qtr
	Intent	2.376	[1.518, 3.368]	0.29 qtr
2014–2016	Satisfaction	2.100	[1.416, 3.047]	0.33 qtr
	Intent	2.651	[1.643, 3.903]	0.26 qtr
2017–2019	Satisfaction	1.778	[1.195, 2.581]	0.39 qtr
	Intent	2.111	[1.330, 2.993]	0.33 qtr
2020–2022	Satisfaction	1.895	[1.296, 2.770]	0.37 qtr
	Intent	2.427	[1.530, 3.455]	0.29 qtr

Intent reverts faster (2–3 months) than satisfaction (3–3.5 months)

$\Rightarrow$ Tactical offers decay quickly; satisfaction investments persist

Time-Varying Covariate Effects

Variable	Satisf.	Intent
Network reliability	1.009***	−.692***
Value for price	1.007***	−.866***
Brand is for me	1.130***	−.760***
Rewards (high-tier)	1.708***	−.902***
Rewards (mid-tier)	1.419***	−.562***
Network speed	.335***	.082
Billing accuracy	.027	.101

*** 95% CI excludes zero

Dominant Drivers:

- Network reliability
- Perceived value
- Brand fit

Rewards programs matter substantially!

Speed $\neq$ Reliability:

Speed $\rightarrow$ satisfaction only

Reliability $\rightarrow$ both

▸ Full Parameter Tables

Static Effects on Equilibrium

Variable	Satisf.	Intent
Carrier (vs. T-Mobile)		
Verizon	2.052***	−.146
AT&T	1.669***	.145
Demographics		
Female	3.201***	−.408
Income (vs. <$25k)		
$100k–$150k	2.658***	.126
$75k–$100k	.879***	−.386***

Carrier Effects:
Verizon > AT&T > T-Mobile
(satisfaction only, not intent!)

Income Puzzle:
$75k–$100k = "sweet spot"
High satisfaction + low churn risk

Gender:
Non-male: high satisfaction
but same churn risk

Cross-Carrier Forecasting: Predicting Rival Customers

What this enables:

1. Recover a posterior trajectory for *any* RCS respondent — including those on rival carriers
2. Forecast their next-period attitudes from the latent OU dynamics
3. Diagnose where a rival's customer base is trending — without rival panel data

Why it works:

Fusion transfers dynamics from *your* panel to anyone drawn from the same population.

Precision Gains

Panel fusion delivers:

5.0% narrower CIs for satisfaction
4.4% narrower CIs for intent

$\Rightarrow$ More confident cross-carrier forecasts

Example

Forecast AT&T customers' satisfaction over the next two quarters using only your panel + industry RCS.

Summary

1. **The Recovery:** Every survey respondent gets a *full posterior latent trajectory*
2. **The Fusion:** RCS (broad coverage) + Panel (individual dynamics) → joint inference on attitudes
3. **The Unlock:** Mean reversion (γ) identified—critical for forecasting *when* shocks fade
4. **The Performance:** 25% RMSE improvement, 28% narrower credible intervals vs. benchmarks
5. **The Application:** Cross-carrier forecasting using only *your* panel + industry RCS

Fusion turns survey snapshots into individual-level trajectory forecasts

Managerial Implications: What Can You Do With Trajectories?

1. Forecast & Triage

- Forecast each customer's trajectory
- Identify at-risk segments *before* they churn
- Prioritize retention spend by predicted risk

2. Time Your Response

- Intent reverts in 2–3 months
- Satisfaction in 3–3.5 months
- Some shocks self-correct—don't overreact

3. Target Precisely

- Low-income: crisis-sensitive
- Long-tenure: churn-prone
- Recent cohort: satisfaction-reactive

The Bottom Line

Transform "What happened?" surveys into "What will happen next?" forecasting engines

Extensions

1. **Longer Forecast Horizons:** Validate year-ahead trajectory forecasts and churn-window predictions
2. **Broader Domains:** Brand tracking, B2B audits, platform trust metrics—anywhere RCS + limited panel coexist
3. **Real-Time Integration:** Fuse transactional/engagement data with survey anchors
4. **Computational:** Variational inference and state-space approximations for scale

Bottom Line

Extract more value from data *you already collect*
Turn descriptive dashboards → prescriptive engines

Thank You
Merci beaucoup

Samuel Levy
levys@darden.virginia.edu
University of Virginia, Darden School of Business

Appendix

Appendix: T-Mobile Data Breach (Aug 2021)

The Event:

- 76.6 million individuals exposed
- SSNs, driver's licenses, addresses
- $350M class action settlement
- $31.5M FCC settlement

Why useful as a real shock:

- Plausibly exogenous
- Unexpected, widely publicized
- Unambiguously negative
- Orthogonal to firm actions

Validation Goals

1. **Internal:** Does γ predict decay?
2. **External:** Can we forecast *competitor* customers' trajectories?

Cross-Carrier Forecast

Using *only* a partnering carrier's panel + industry RCS, recover trajectories for T-Mobile customers.

Back

Appendix: Average Treatment Effects

Outcome	DiD (Month 0)	Peak Shock	Peak Month	Half-life	γ
Satisfaction	$+.19$	$-.83$	Sep 2021	2.5 mo	.27
Intent to Switch	$-.17$	$+.41$	Dec 2021	6.1 mo	.11

Satisfaction:

- Peak: -0.83 (8% of scale)
- Half-life: 2.5 months
- Fast recovery

Intent to Switch:

- Peak: $+0.41$ (10% of scale)
- Half-life: 6.1 months
- Slower decay

Novel Insight

γ and half-life **only identifiable** via panel fusion!

Appendix: Treatment Effect Dynamics — Satisfaction

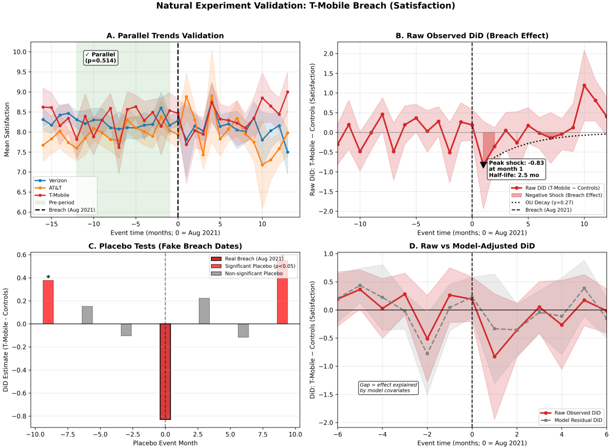

Panel A: Parallel trends pre-breach ($p = .739$) **Panel B:** Raw DiD + OU decay overlay

Panel C: Placebo tests **Panel D:** Raw vs. residual DiD

Back

Appendix: Treatment Effect Dynamics — Intent to Switch

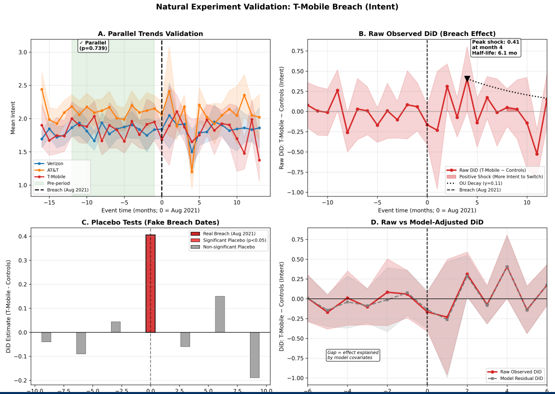

Managerial Interpretation

Breach = **transitory level-shifter**, not trajectory shifter; T-Mobile's long-run CRM strategy remained effective

Appendix: Heterogeneity — Income

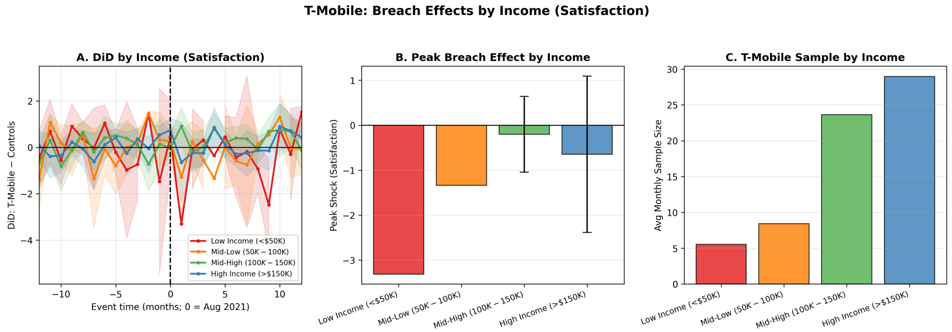

Low-income customers: Largest satisfaction decline (−3.2 points)—most economically vulnerable to identity theft ramifications

◂ Back

Appendix: Heterogeneity — Acquisition Cohort

Cohort	Satisfaction		Intent	
	RCS-only	Fusion	RCS-only	Fusion
Long-tenure (≤2013)	−.09	−.12	+.18*	+.18*
Mid-tenure (2014–16)	−.20	−.08	−.00	−.07
Recent (2017–19)	−.30*	−.30*	+.07	+.03
New (2020–22)	−.34	−.17	+.19	+.15

Satisfaction:
Recent cohort (2017–19) largest decline

Intent:
Long-tenure customers most likely to consider switching!

Targeting Implication

Prioritize: low-income & recent for satisfaction; long-tenure for retention

Appendix: OU Process Technical Details

Stochastic Differential Equation:

$$d\mathrm{x}_i(t) = [-\boldsymbol{\Gamma}(\mathrm{x}_i(t) - \boldsymbol{\mu}_i) + \boldsymbol{\beta}\mathrm{Z}_i(t)]\, dt + \boldsymbol{\Sigma} d\mathrm{W}_i(t)$$

Parameter Interpretations:

- $\boldsymbol{\Gamma} = \mathrm{diag}(\gamma_1, \ldots, \gamma_j)$: Mean reversion rates
- Half-life of shock: $t_{1/2} = \ln(2)/\gamma_j$
- If $\gamma_1 = 1$: shock decays to half in ≈ 0.69 time units
- $\boldsymbol{\beta}_{jk}$: Drift of β_{jk} units per period when covariate k active
- Steady-state under sustained disruption: $\boldsymbol{\mu}_i - \beta/\gamma$

Back to Model

Appendix: Transition Covariance

Transition Covariance:

$$\mathrm{V}_{s,t} = \mathrm{V}_\infty - e^{-\boldsymbol{\Gamma}(t-s)}\mathrm{V}_\infty e^{-\boldsymbol{\Gamma}^\top(t-s)}$$

where V_∞ satisfies Lyapunov equation: $\boldsymbol{\Gamma}\mathrm{V}_\infty + \mathrm{V}_\infty\boldsymbol{\Gamma}^\top = \boldsymbol{\Sigma}\boldsymbol{\Sigma}^\top$

For diagonal Γ:

$$[\mathrm{V}_\infty]_{jk} = \frac{[\boldsymbol{\Sigma}\boldsymbol{\Sigma}^\top]_{jk}}{\gamma_j + \gamma_k}$$

Intuition:

- As $t - s \to \infty$: $\mathrm{V}_{s,t} \to \mathrm{V}_\infty$
- Faster mean reversion $(\gamma) \Rightarrow$ smaller long-run variance

Back to Model

Appendix: Full Covariate Effects (Time-Varying)

Variable	Satisfaction	Intent
Network speed	.335***	.082
Network reliability	1.009***	−.692***
Data plans meet needs	.235***	−.372***
Value for price paid	1.007***	−.866***
Service cost	.845***	−.441***
Billing accuracy	.027	.101
Rewards/recognition	−.073	−.059
Easy to do business	.880***	−.368***
First-call problem solving	.200**	−.160
Brand is for me	1.130***	−.760***
Device selection	−.109	.209**
Plan understanding	.934***	−.209**
Unlimited data (high-tier)	.601**	.595
Rewards program (high-tier)	1.708***	−.902***

◂ Back to Results

Appendix: Full Covariate Effects (Static)

Variable	Satisfaction	Intent
Carrier (vs. T-Mobile)		
Verizon	2.052***	−.146
AT&T	1.669***	.145
Female	3.201***	−.408
Income (vs. <$25k)		
$25k–$35k	.785***	.263***
$35k–$50k	1.182***	.771***
$50k–$75k	1.301***	.883***
$75k–$100k	.879***	−.386***
$100k–$150k	2.658***	.126
$150k–$200k	2.780***	.112
$200k+	2.773***	.332***
Cohort (vs. 2012–Earlier)		
2014–2016	2.715***	.334***
2017–2019	2.724***	.432***
2020–2022	2.921***	.267

Back to Results

Appendix: Model Architecture

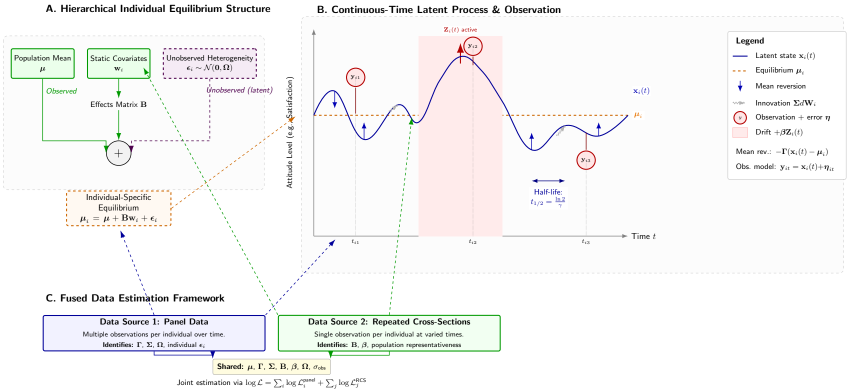

Figure: DRIFTS: Hierarchical OU Framework

◂ Back to Model

Appendix: Simulation Comparison Details

Comparison Model Mechanisms:

- **CVAE:** Encoder-decoder denoises but doesn't transport info across domains
- **cGAN:** Lacks reversion operator; generator noise $\neq$ posterior uncertainty
- **GP:** Smoothness assumptions interpolate *through* rather than *to* equilibria
- **BART:** Tree partitions fail to generalize beyond observed customers
- **Cohort:** Hierarchical shrinkage dilutes longitudinal signal at long horizons

Key Insight: Not all uncertainty quantification is created equal; explicit OU specification captures true dynamics.

Back to Simulations

Appendix: Natural Experiment Robustness

Parallel Trends Test:
$p = .739$ for satisfaction; $p > .05$ for intent

Placebo Tests:
Fake breach dates show no comparable shocks

Residual DiD:
Raw $\approx$ Model-adjusted $\Rightarrow$ effect not explained by observables

Gains from Fusion:
Panel fusion achieves 5% improvement in precision of credible intervals

Interpretation:
Can be viewed as generative extension of DiD with multiple time periods (Callaway & Sant'Anna, 2021)

◂ Back to Natural Experiment