Mortality Effects and Choice Across Private Health Insurance Plans

Jason Abaluck
Yale University

Mauricio Caceres Bravo
Brown University

Peter Hull
University of Chicago

Amanda Starc
Northwestern University

Preview of Our Findings

- Huge differences in causal mortality effects across Medicare Advantage insurance plans

- Consumers respond weakly to these differences

This is a problem: why should insurers invest in making people healthier?

Three Empirical Challenges

To measure plan mortality effects:

1) Consumers of different health sort non-randomly into plans

2) RCTs and individual quasi-experiments are underpowered to estimate mortality effects for each plan

To measure consumer responsiveness:

3) Hard to estimate demand for an unobservable (plan quality)

Our Solutions

We address **selection** and **power** by combining observational estimates with many quasi-experiments (plan terminations)

- We formalize how IV can be used to relate causal effects to observational measures, allowing for selection bias
- We extend this approach to measure other correlates of plan mortality effects (e.g. star ratings, spending…)

We show how consumer willingness to pay can be estimated by combining our IV framework with discrete choice modeling

Literature Focuses on Financial Aspects

- **Within networks** (Handel and Kolstad 2016, Abaluck and Gruber 2016, Bhargava, Loewenstein and Sydnor 2017,…)
- **Argue outcome proxies (e.g. hospital quality) don't change** (Gruber and McKnight 2016)
- **GE impact on premiums or coverage generosity** (Handel 2013, Ho & Scott Morton 2017, Starc and Town 2019, …)

We'll show in MA that outcome differences are first-order

Revealed preference (e.g. Ho and Lee 2019, Gruber, Handel, Kina and Kolstad 2020) is better than nothing, but far from perfect

Our Setting: Medicare Advantage (MA)

Beneficiaries choose from subsidized private-managed care plans (around 30 per county, on average)

- Highly differentiated: premiums, benefits (including prescription drug coverage), spending (medical loss ratios), networks (doctors and hospitals)…
- Ranked by CMS star ratings (no mortality information)

We study the universe of Medicare beneficiaries

- Demographics, enrollment, and mortality (no claims)

Plan for Today

1. Constructing observational mortality
2. Identifying quasi-experimental variation
3. Econometric framework
4. Main results
5. Mechanisms, willingness to pay, and simulations

Constructing Observational Mortality

We estimate mortality rates of plans within markets, among observably similar beneficiaries, accounting for statistical noise

$$Y_{it} = \sum_{j=1}^{J} \mu_j D_{ijt} + X_{it}'\gamma + \nu_{it},$$

Y_{it}: one-year mortality of individual i in year t,

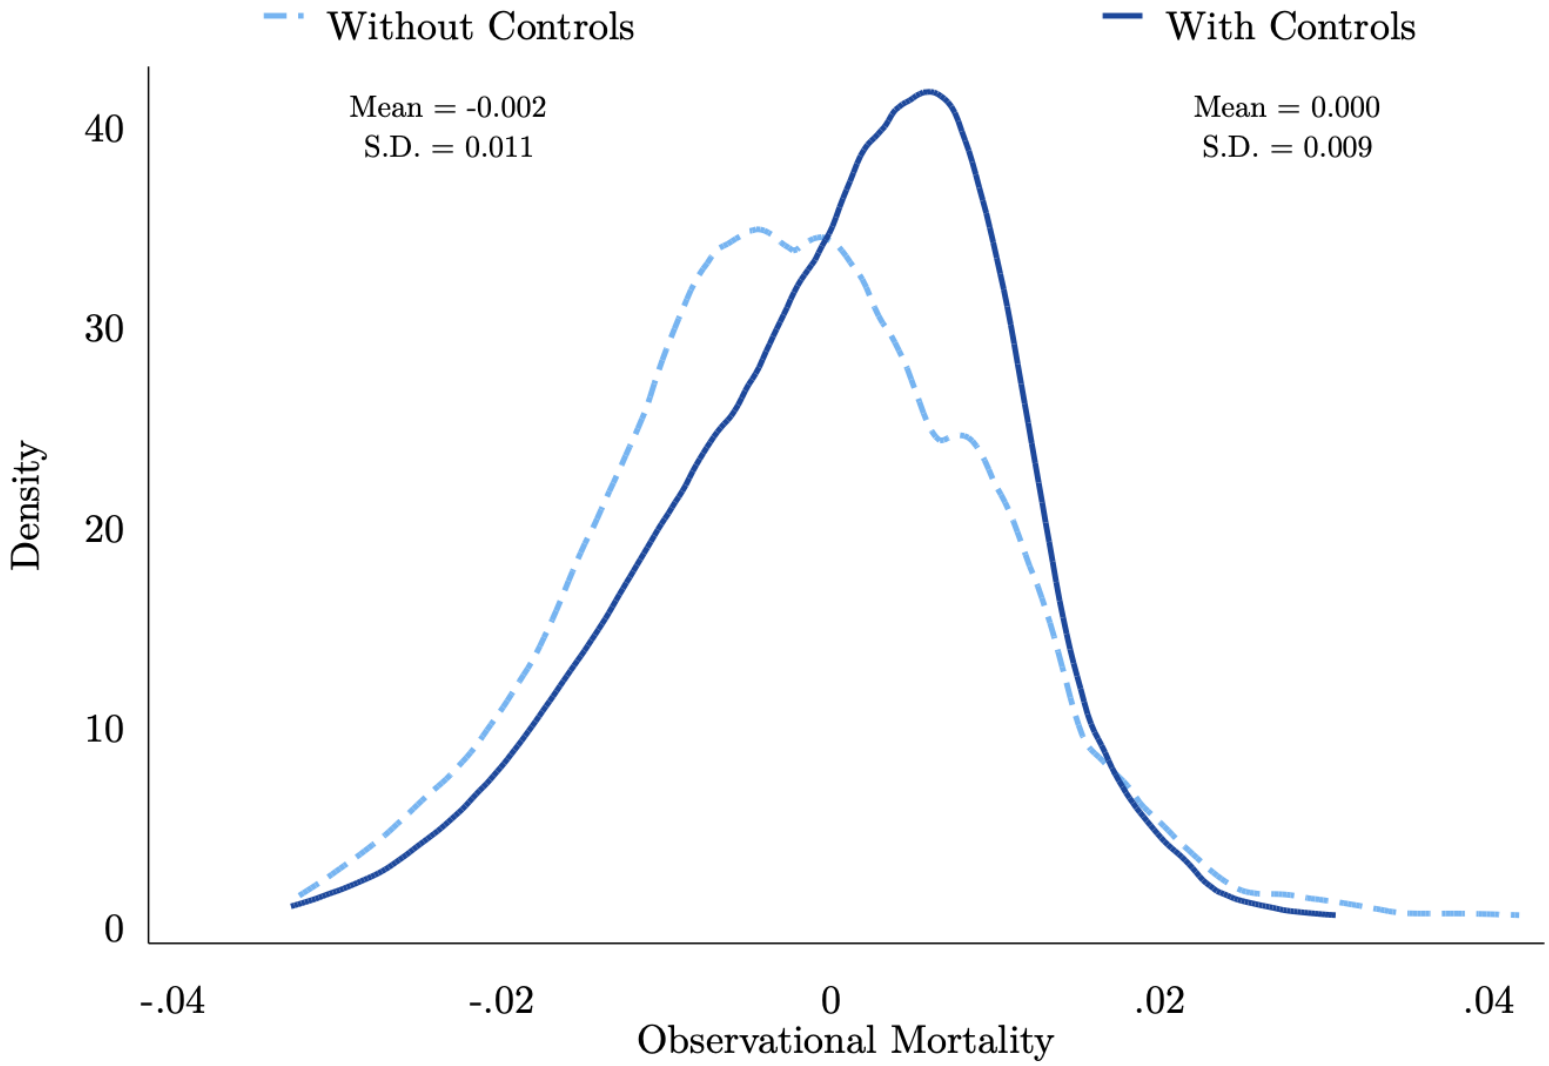
D_{ijt} = 1 if i starts year t enrolled in plan j,

X_{it}: vector of basic individual controls (age, race, sex, dual-eligible status) and county + year fixed effects

ν_{it}: regression residual

μ_j: **observational mortality** (shrunk by empirical Bayes)

Variation in Plan Observational Mortality



Average One-Year Mortality Rate: 4.7%

Quasi-Experimental Choice Set Variation from MA Plan Terminations

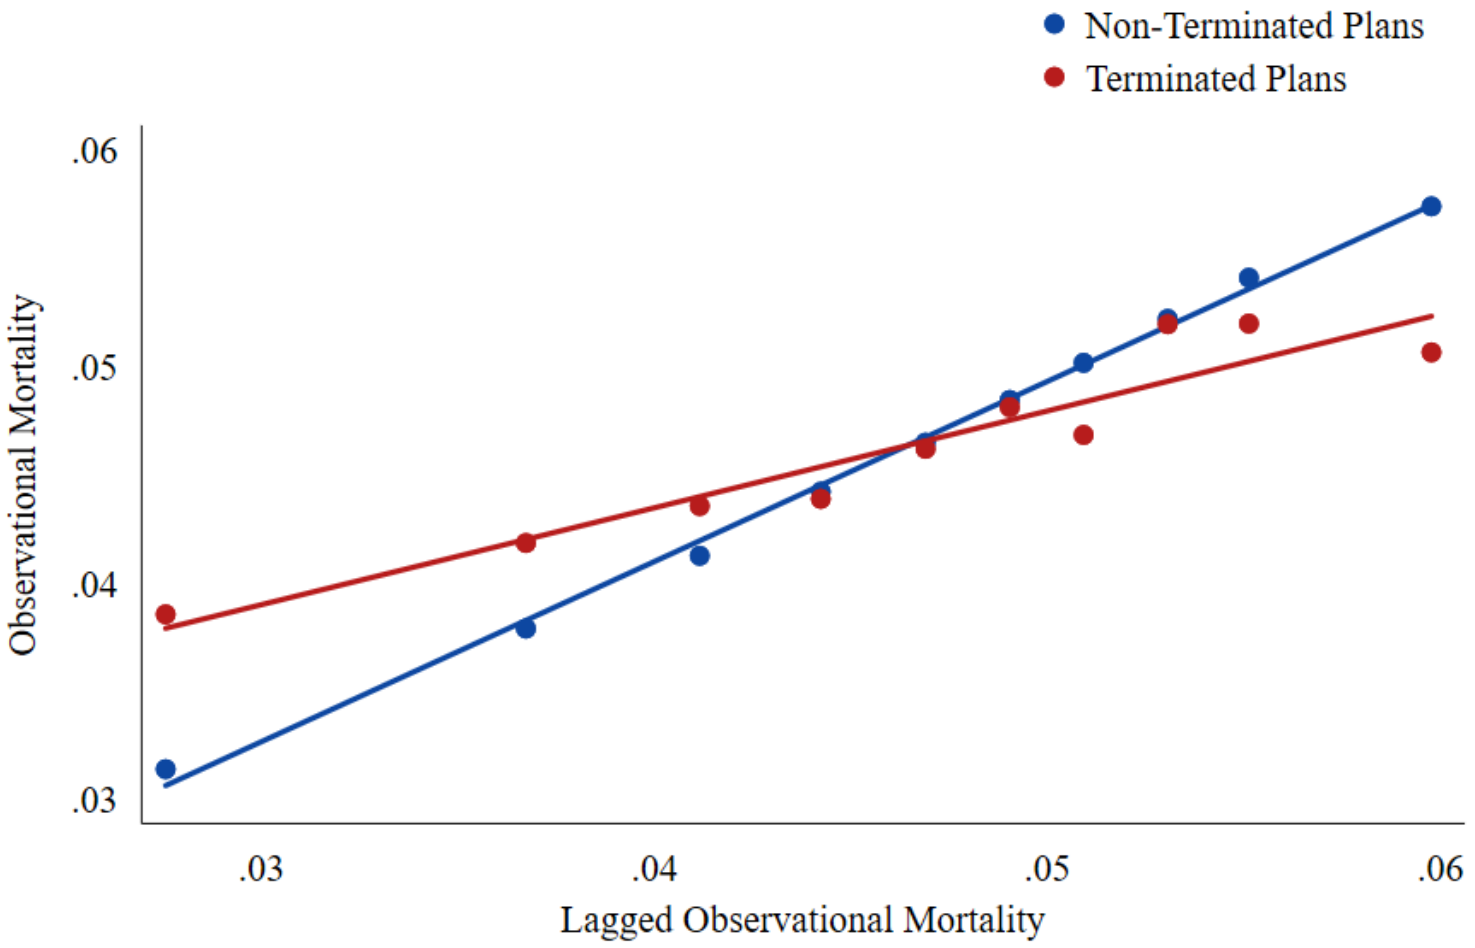
2008 law led to the unexpected exit of private-fee-for-service plans from 2009-2011 (Pelech 2018, Duggan, Gruber and Vabson 2018)

Map of Terminations

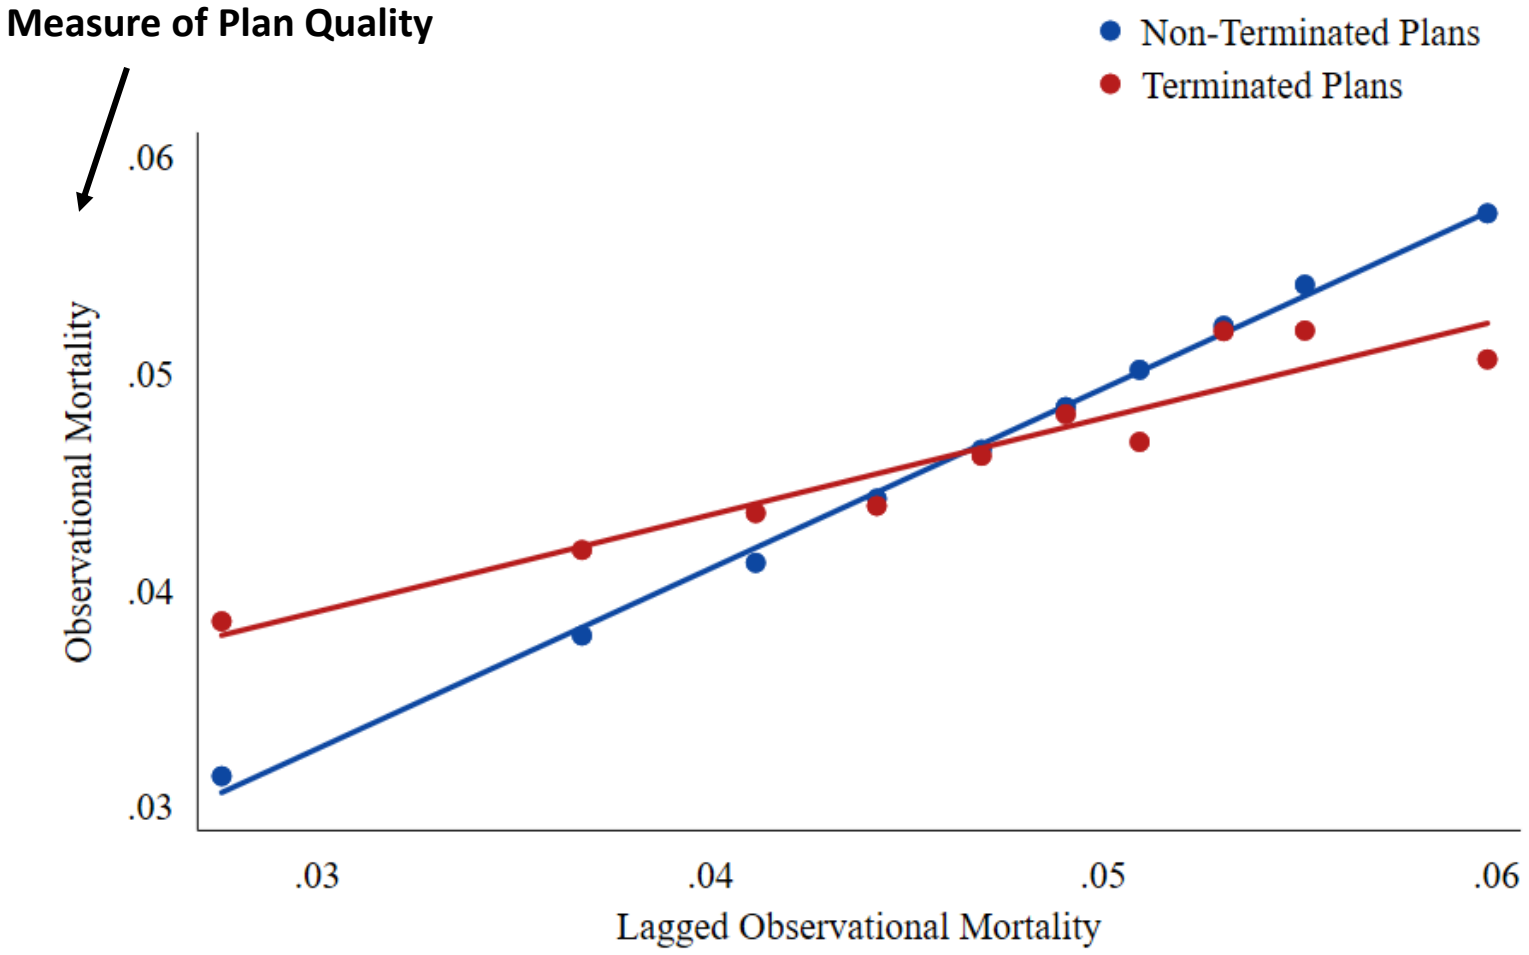
We leverage the interaction of all plan **terminations** and **inertia**

- Absent terminations, beneficiaries enrolled in high- and low-mortality plans tend to stay there
- After terminations, consumers tend to regress to the mean

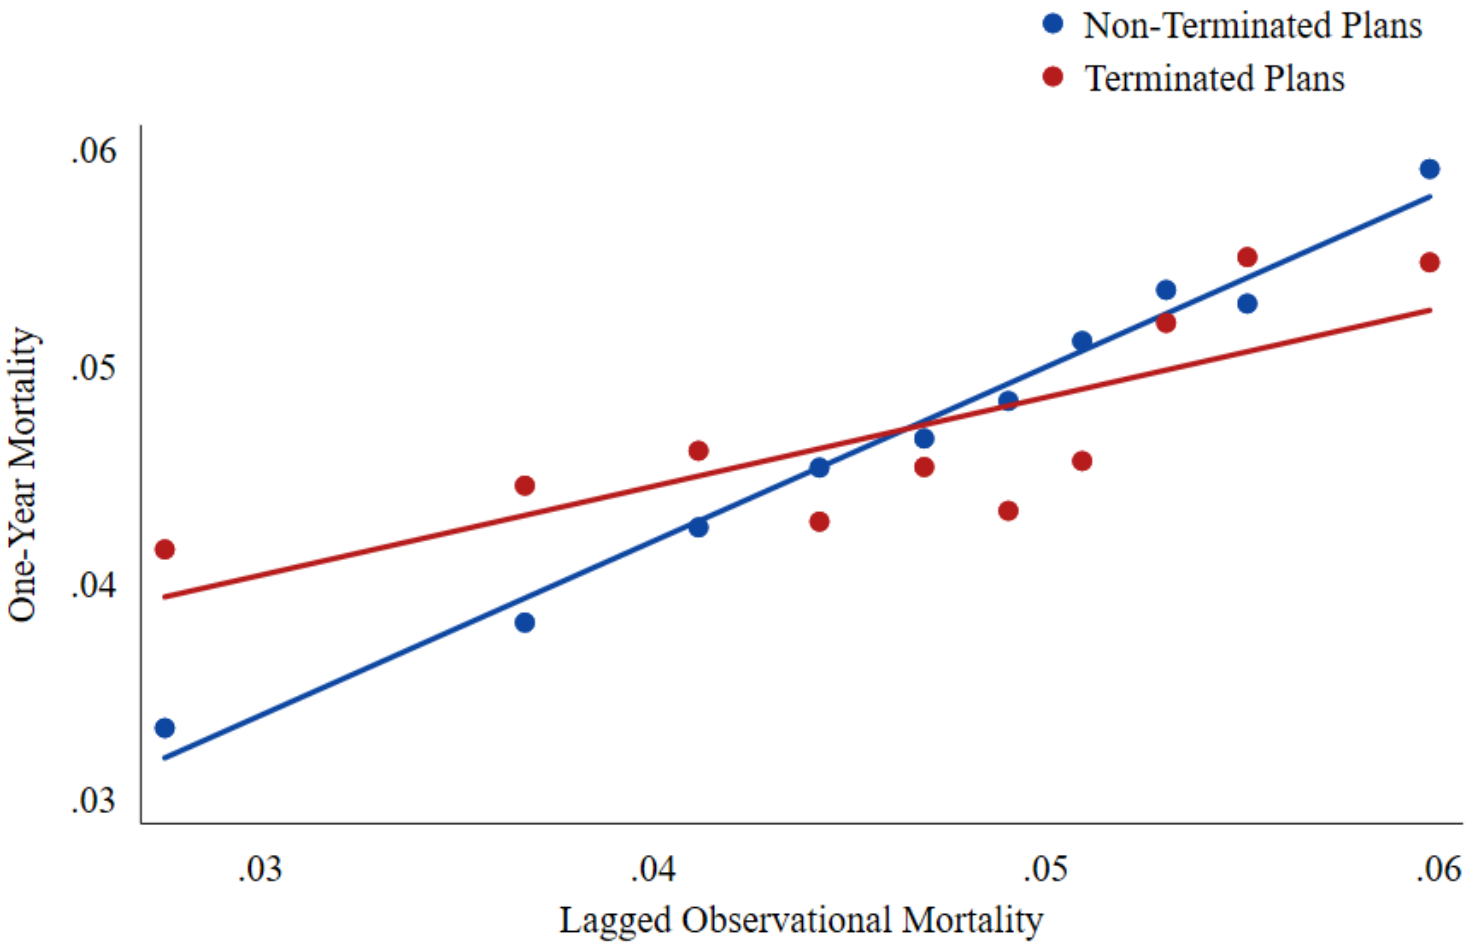
Observational Mortality of Chosen Plans: Terminated and Non-Terminated Plans



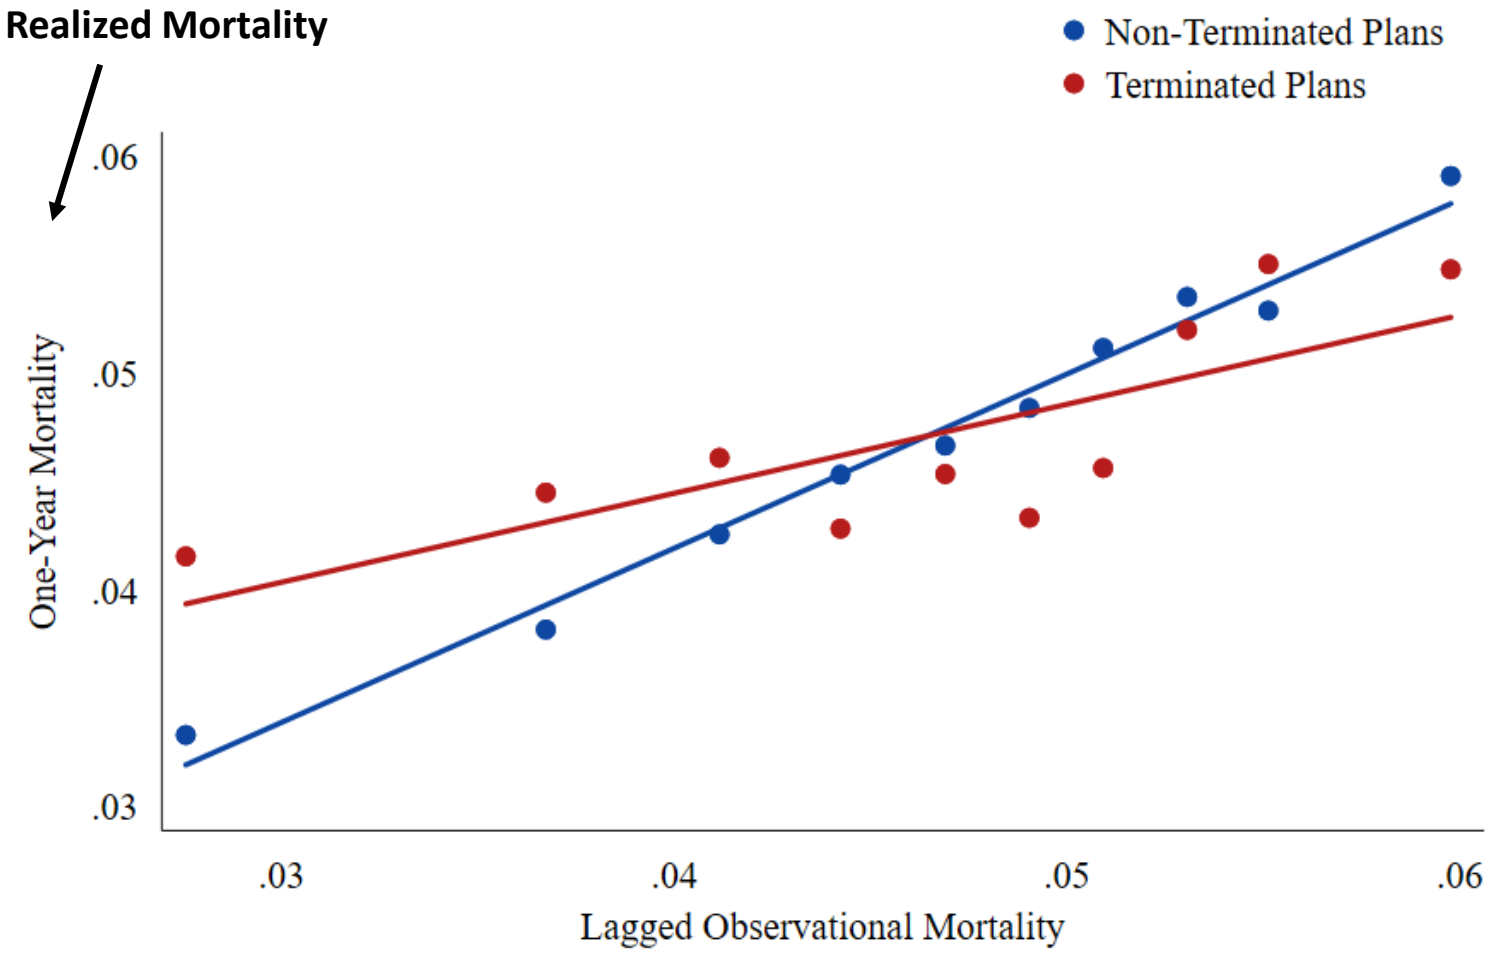
Observational Mortality of Chosen Plans: Terminated and Non-Terminated Plans



Actual Mortality of Plan Beneficiaries: Terminated and Non-Terminated Plans



Actual Mortality of Plan Beneficiaries: Terminated and Non-Terminated Plans



Diff-in-Diff Version

Goal: MA Plan Forecast Regressions

Y_{ij}: potential mortality outcome for individual *i* if in plan *j*

If we randomized people to plans, we could estimate quality:

$$\beta_j = E[Y_{ij}].$$

Imagine regressing (unknown) β_j on (observed) μ_j:

$$\beta_j = \lambda \mu_j + \eta_j$$

Goal: recover λ ("forecast coefficient") from an IV regression.

IV Framework

$$\beta_j = \lambda \mu_j + \eta_j$$

Realized Mortality →

$$Y_{it} = \beta_{j(it)} + X_{it}\gamma + \varepsilon_{it}$$

Second Stage:

New error term! ↙

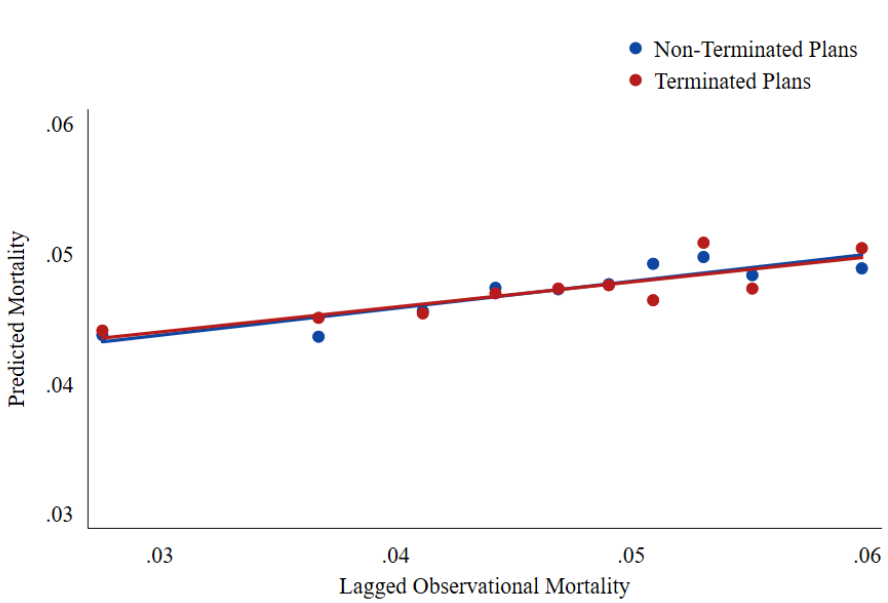
$$Y_{it} = \lambda \mu_{j(it)} + X_{it}\gamma + \varepsilon_{it} + \eta_{j(it)}$$

Instrument enrolled plan observational mortality $\mu_{j(it)}$ with the interaction $Z_{it} = \mu_{j(i,t-1)} \times T_{i,t-1}$

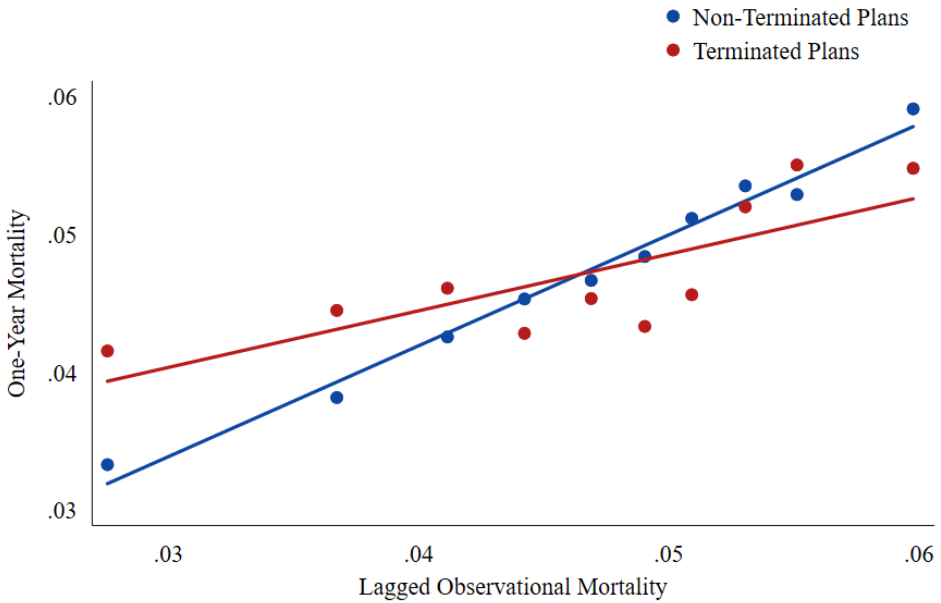
- $T_{i,t-1} = 1$ if i is enrolled in a plan that terminates in $t-1$
- Also control for $\mu_{j(i,t-1)}$ and $T_{i,t-1}$ main effects

Exclusion: Similar Beneficiaries in Terminated and Non-Terminated Plans

Balance: Predicted Mortality



Reduced Form



Risk Score Balance

A New "Fallback Condition"

Random assignment of terminations is not sufficient to get λ

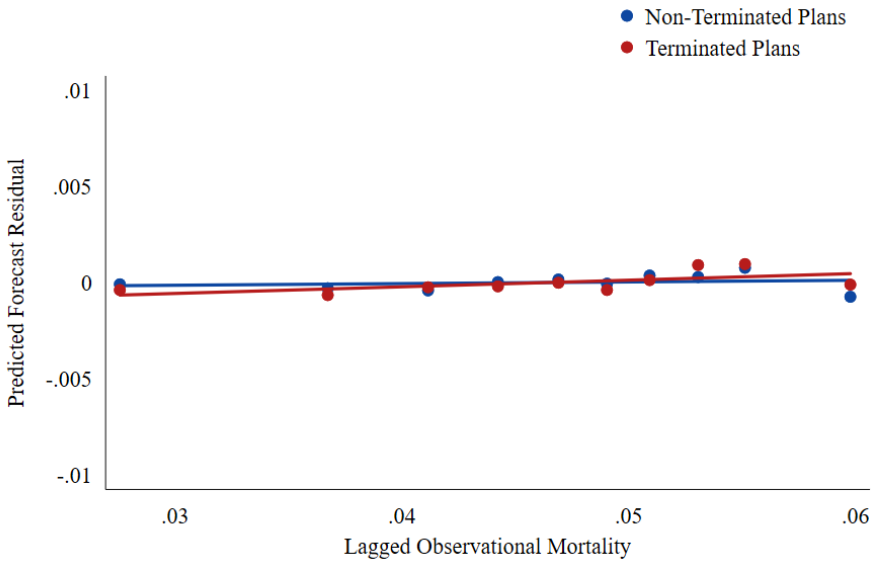
- Need "fallback" plans following termination to be "typical," given observational mortality

In the paper, we give a general class of discrete choice models where this holds and develop balance tests

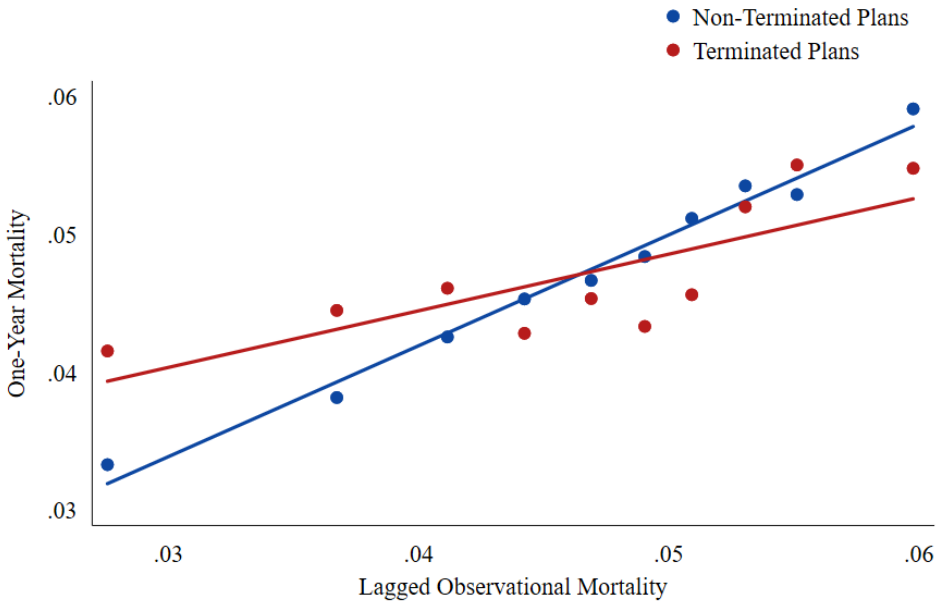
This fallback condition is generally needed for IV validation of "value-added" models (of, e.g., teachers, schools, hospitals…)

Fallback: Beneficiaries in Terminated Plans Have "Typical" Second-Choice Plans

Balance: Forecast Residual



Reduced Form



The IV Forecast Coefficient Is Near one

	(1)	(2)	(3)	(4)
Dep. Var.: Observational Mortality		A. First Stage		
Instrument	−0.349	−0.0055	−0.349	−0.0055
	(0.037)	(0.0011)	(0.037)	(0.0011)
F Statistic	89.6	24.3	89.4	24.3
Dep. Var.: One-Year Mortality		B. Reduced Form		
Instrument	−0.344	−0.0068	−0.386	−0.0065
	(0.099)	(0.0024)	(0.088)	(0.0020)
Dep. Var.: One-Year Mortality		C. Second Stage (Forecast Coefficient)		
Observational Mortality	0.986	1.230	1.107	1.183
	(0.230)	(0.353)	(0.187)	(0.310)
Specification	Linear	Median	Linear	Median
Demographic Controls	No	No	Yes	Yes
N Beneficiary-Years		15,012,189		

Notes: County-clustered standard errors in parentheses

Robustness

What Does This Mean?

Observational mortality **reliably predicts** true mortality effects (doesn't mean there's no selection bias!)

Variation in true effects β_j is **at least as large** as in observed μ_j

- Magnitude: 1 sd = a 19% reduction in one-year mortality

Compare to:

- Extensive margin: 15-20% mortality from insurance (Card, Dobkin and Maestas 2008; Miller, Altekruse, Johnson and Wherry 2019; Sommers, Gawande and Baicker 2017)
- Place-based: moving from 10th-90th perc -> 30% mortality reduction (Finkelstein, Gentzkow and Williams 2019, Deryugina and Molitor 2018)
- Hospital effects: 1 sd = 20% lower mortality (Hull (2020) for emergency room patients; we replicate with all inpatients)

Mechanisms: Other Plan Attributes & Causal Mortality Effects

Star Rating	0.0006				
	(0.0014)				
Premium		−0.0051			−0.0044
		(0.0020)			(0.0026)
Has Donut Hole Coverage			−0.0041		−0.0001
			(0.0016)		(0.0023)
Medical Loss Ratio				−0.0214	−0.0223
				(0.0044)	(0.0044)
First-Stage F Statistic	2,860.3	2,085.8	1,437.8	1,644.6	372.9
Maximum Forecast R^2	0.0005	0.0218	0.0214	0.0095	0.0298
N Beneficiary-Years			15,012,189		

Notes: County-clustered standard errors in parentheses

OLS Estimates

WTP for Plan Quality is Positive, But Low

Premium Elasticity	Premium Coefficient (α)	Forecast Coefficient (κ)	Minimum Quality Coefficient (τ)	Maximum WTP: $\tau/(100 \times \alpha)$
(1)	(2)	(3)	(4)	(5)
-10	−0.0229	−0.0003	−403.95	176.38
		(0.0001)	(129.74)	(56.65)
-7	−0.0160	−0.0004	−284.07	177.19
		(0.0001)	(91.30)	(56.95)
-3.5	−0.0080	−0.0007	−144.81	180.66
		(0.0002)	(47.45)	(59.20)
-1	−0.0023	−0.0017	−46.43	202.75
		(0.0008)	(22.23)	(97.05)
-0.5	−0.0011	−0.0015	−25.23	220.30
		(0.0013)	(21.68)	(189.30)

Notes: County-clustered standard errors in parentheses

Calculation Details

Large Potential Mortality Reductions from Better Aligning Quality and Plan Choice

	Change Among Reassigned (1)	% of Mean Mortality (2)	Unconditional Change (3)	% of Mean Mortality (4)
Random Assignment to Plans	0.0027	5.7	0.0027	5.7
Assignment to Minimum-Mortality Plans	−0.0192	-40.8	−0.0192	-40.8
Assignment from Top- to Bottom-Quartile Plans	−0.0077	-16.3	−0.0019	-4.1
Random Assignment from Top 5% of Plans	−0.0108	-23.0	−0.0005	-1.1

Note: partial equilibrium analysis, so many caveats may apply (e.g. capacity constraints, insurer exit/entry, market disruption effects…)

Simulation Details

Conclusions

There is significant variation in the mortality effects of MA plans, but consumers are mostly insensitive to this variation

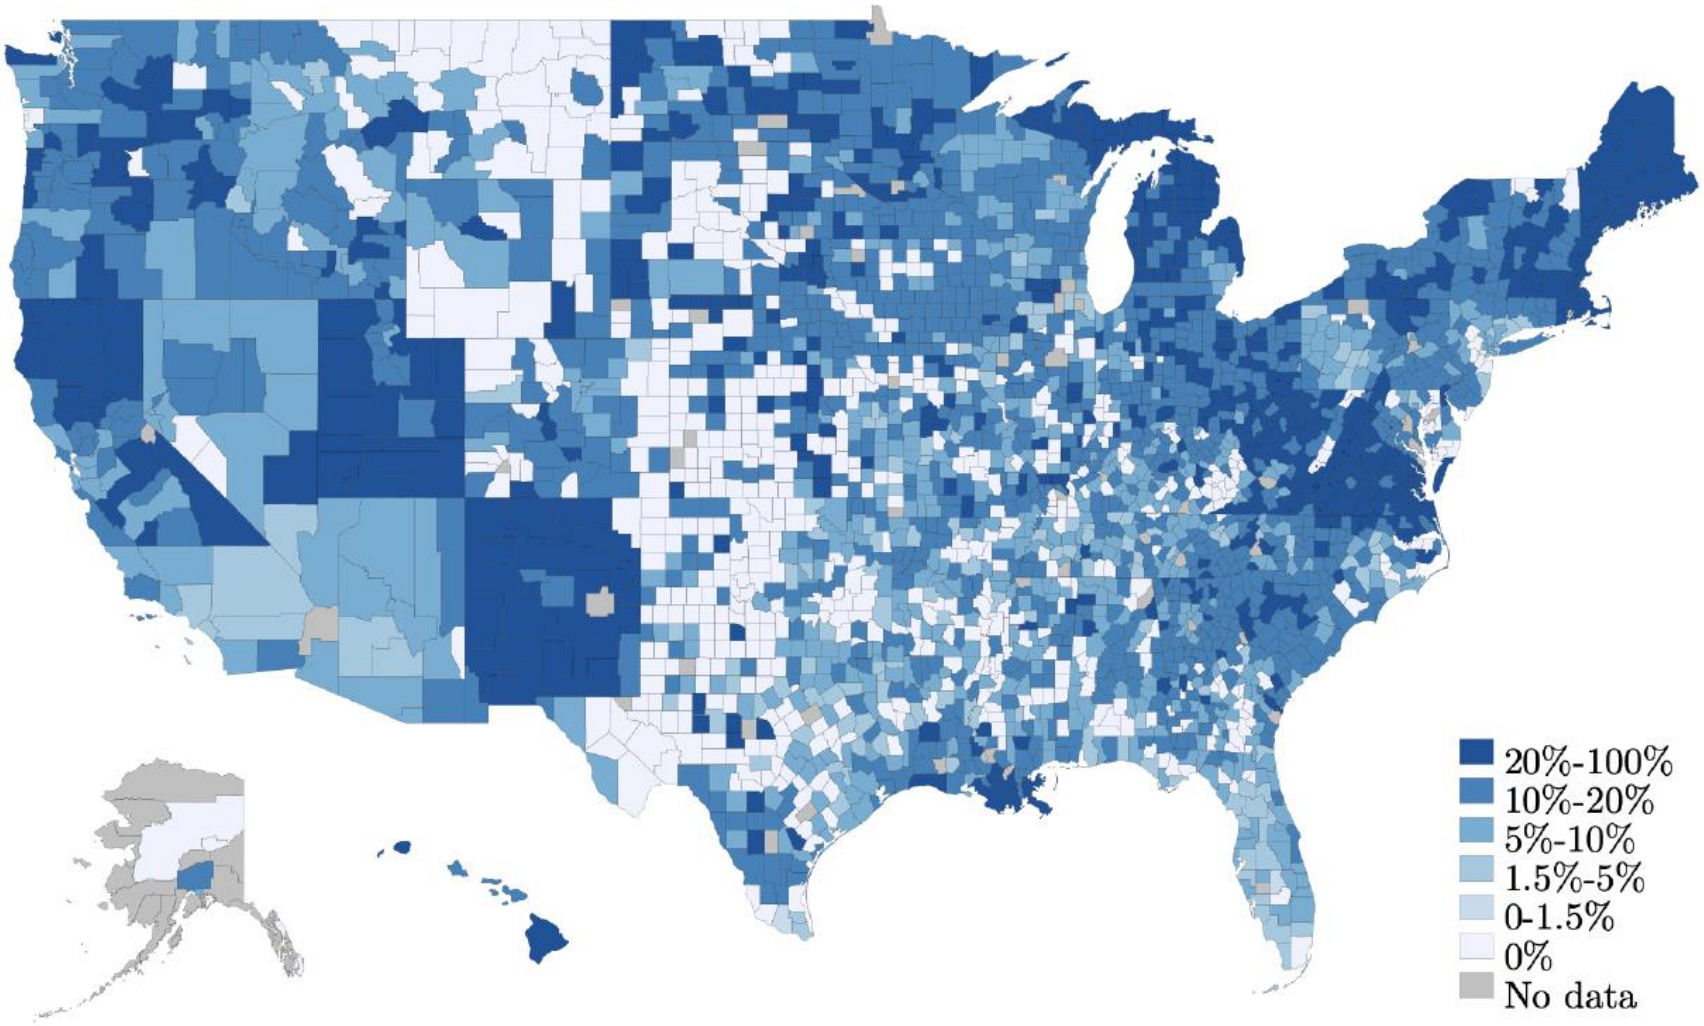
Potential policy responses (my co-authors each hate one):

- Release risk-adjusted mortality rates, like with hospitals
- Subsidize risk-adjusted mortality instead of star ratings
- Audit bad plans and possibly remove them
- Integrate health insurance and life insurance

Frontier: linking plan outcomes to provider outcomes, GE responses, consequences for innovation and tech adoption

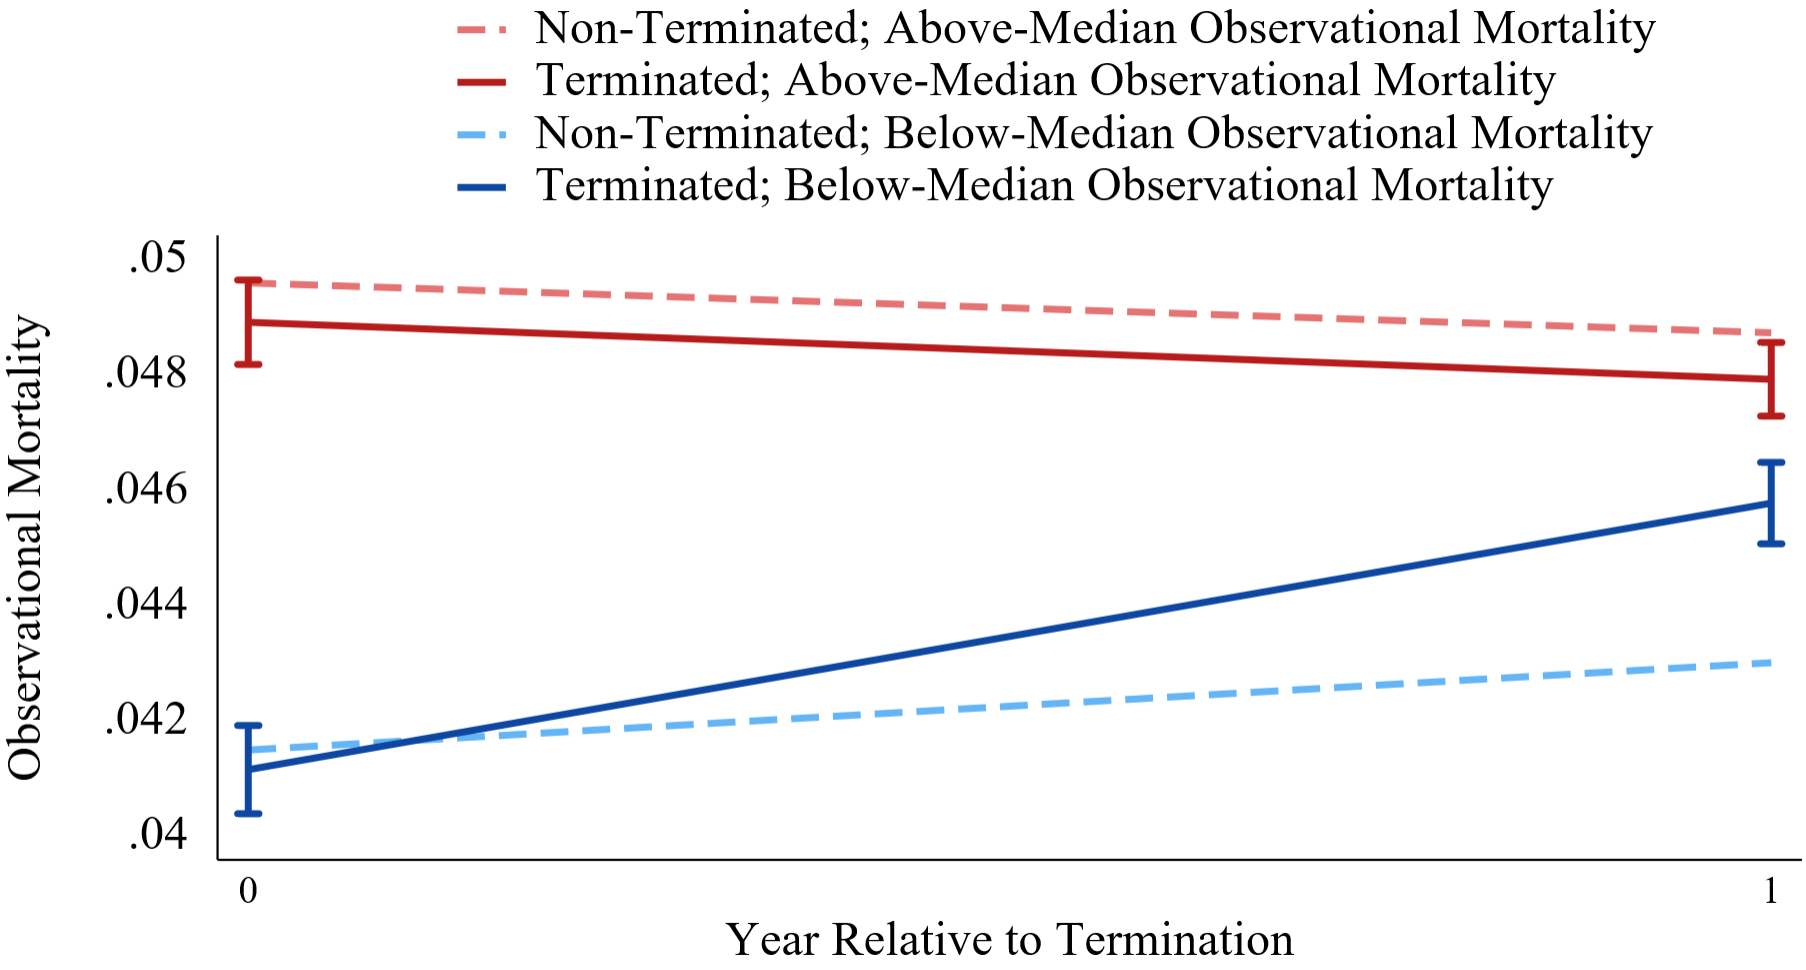
Backup Slides

Most Counties See an MA Termination

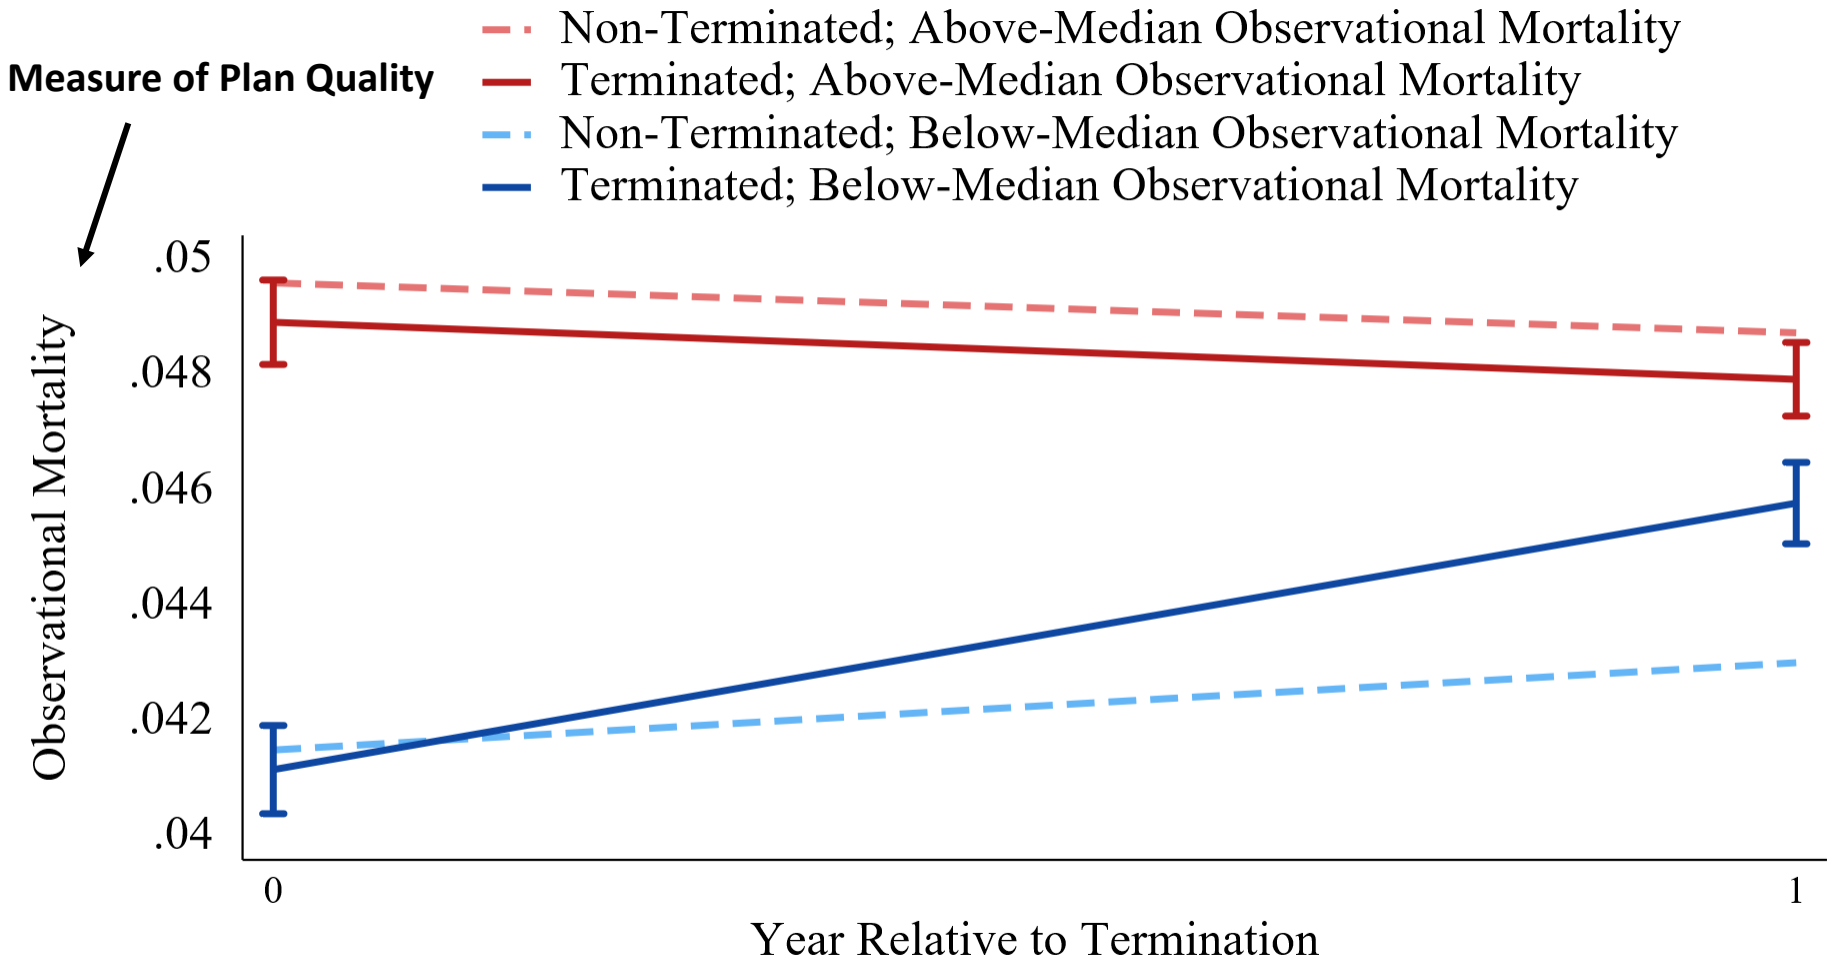


BACK

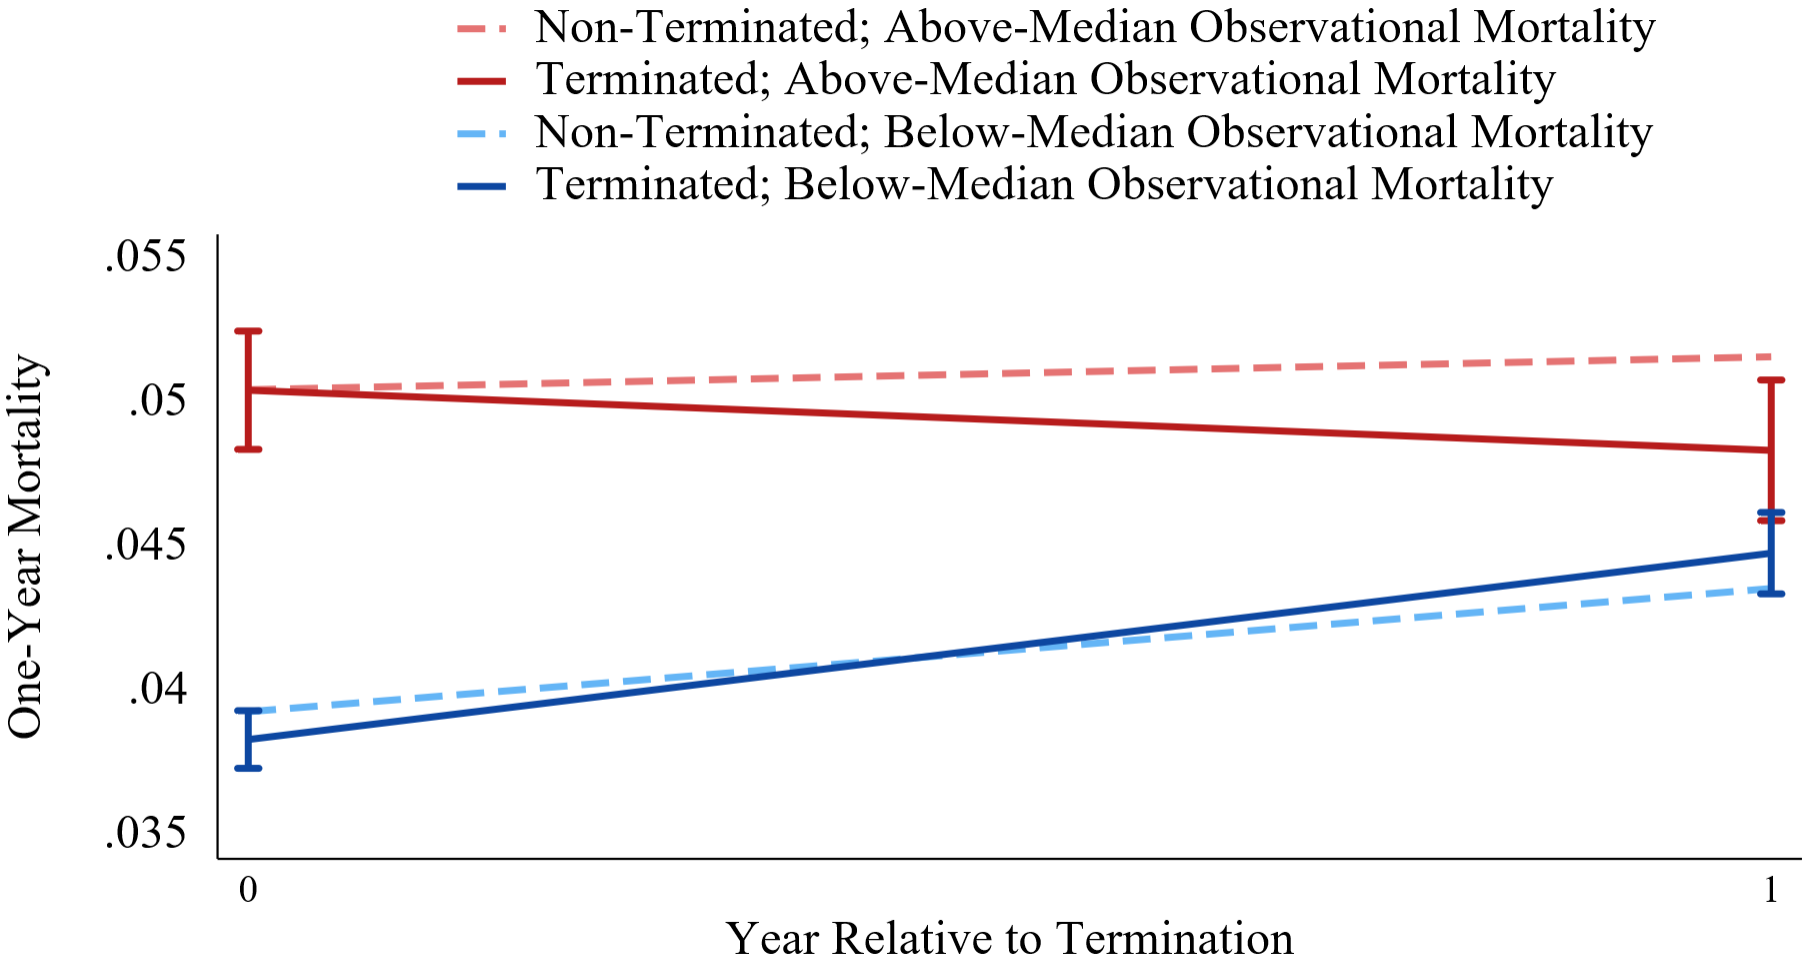
Observational Mortality of Chosen Plans: Pre- and Post-Termination



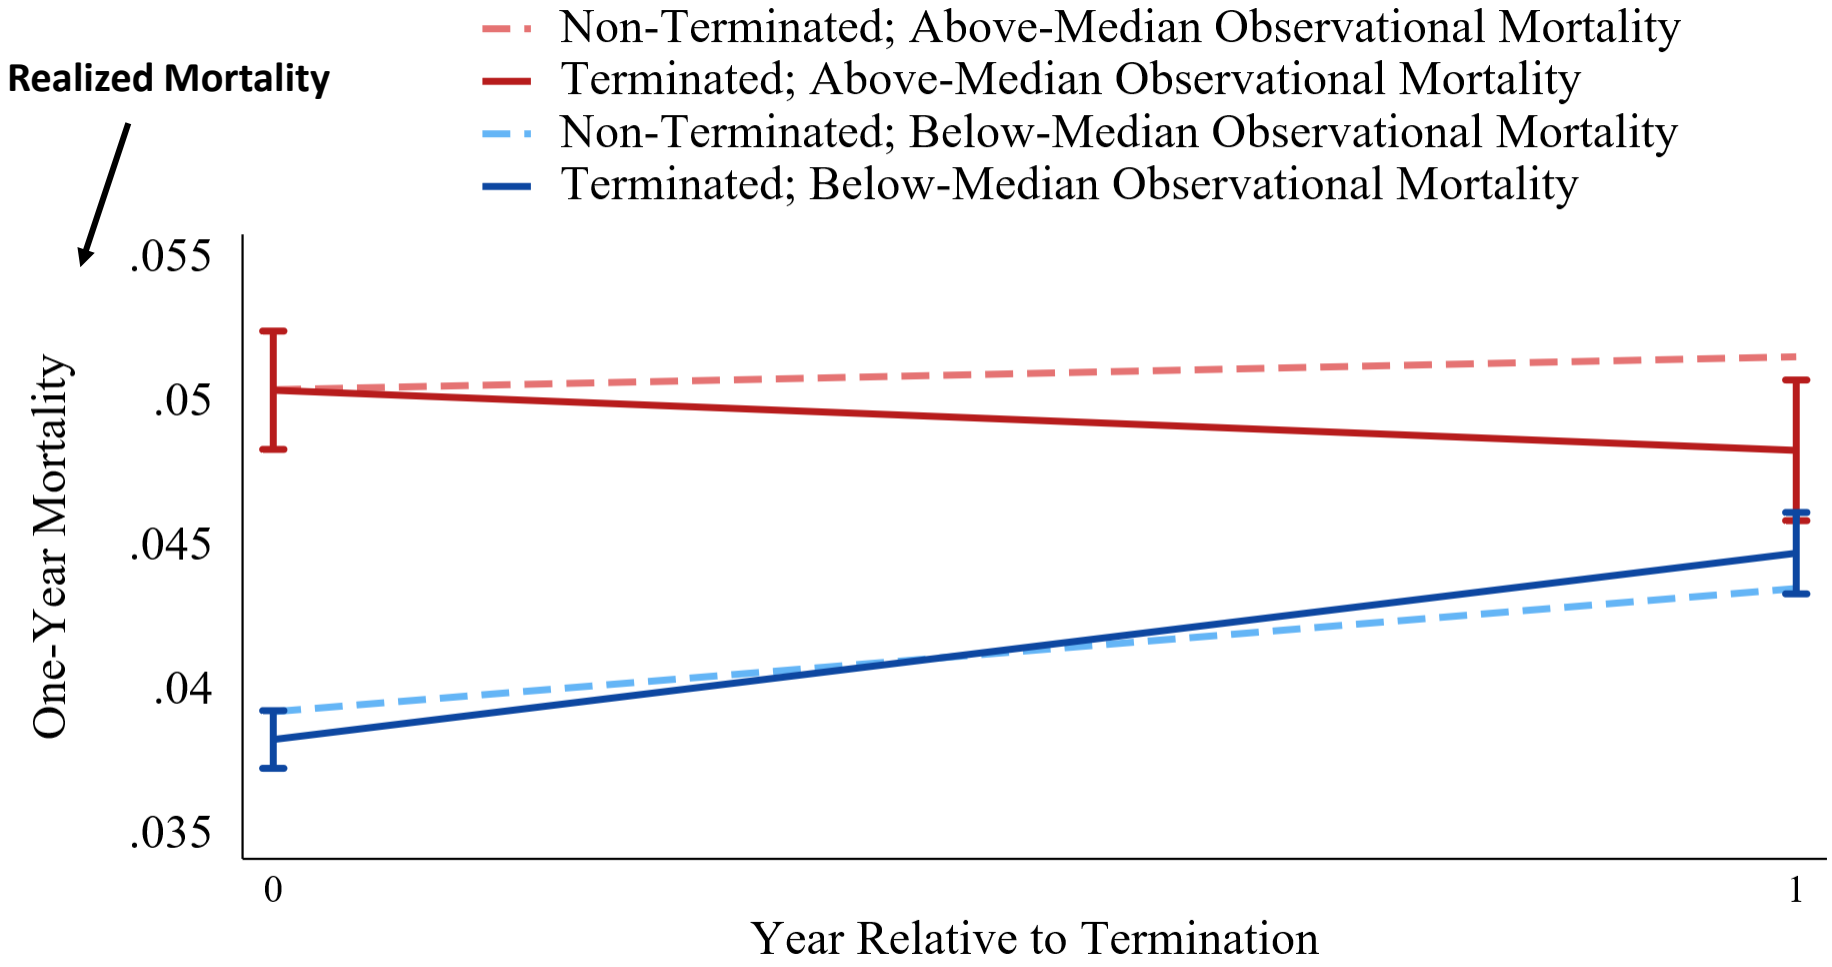
Observational Mortality of Chosen Plans: Pre- and Post-Termination



Actual Mortality of Plan Beneficiaries: Pre- and Post-Termination

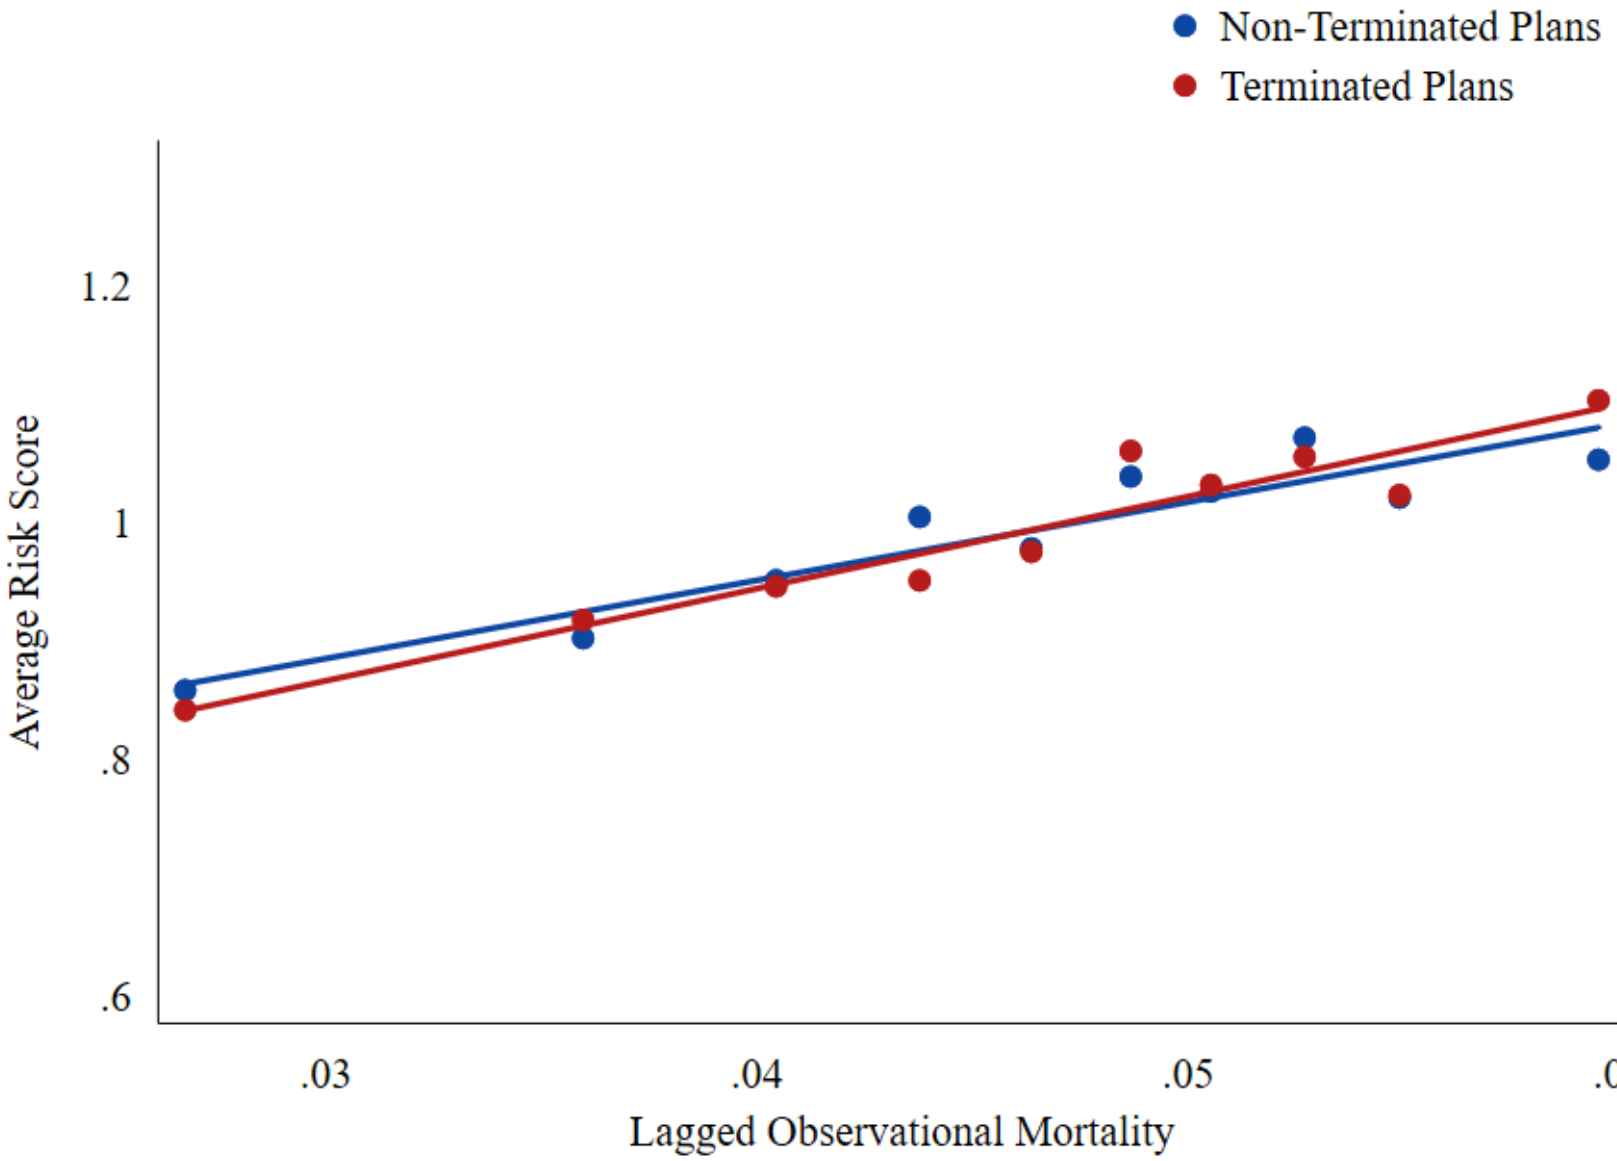


Actual Mortality of Plan Beneficiaries: Pre- and Post-Termination



BACK

Exclusion: Similar Risk Scores in Terminated and Non-Terminated Plans



BACK

Robustness Checks

- Exclude counties w/o terminations
- Exclude beneficiaries in TM
- Limit to PFFS terminations
  - E.g. Pelech (2018)
- Exclude Dual Eligibles
- Allow for heterogeneity by age

BACK

	(1)	(2)
	A. Counties With Terminations	
Observational Mortality	1.085	1.150
	(0.189)	(0.309)
First-Stage F Statistic	89.6	24.4
N Beneficiary-Years	14,644,200	
	B. No TM Enrollments	
Observational Mortality	1.380	1.325
	(0.219)	(0.289)
First-Stage F Statistic	122.0	32.9
N Beneficiary-Years	14,166,119	
	C. PFFS Terminations	
Observational Mortality	1.154	1.987
	(0.369)	(0.778)
First-Stage F Statistic	54.1	7.2
N Beneficiary-Years	14,904,951	
	D. No Dual-Eligibles	
Observational Mortality	1.132	1.169
	(0.207)	(0.313)
First-Stage F Statistic	107.1	30.5
N Beneficiary-Years	13,151,504	
	E. Age-Specific Effects	
Observational Mortality	1.146	1.135
	(0.088)	(0.140)
First-Stage F Statistic	829.2	231.1
N Beneficiary-Years	15,012,189	
Specification	Linear	Median
Demographic Controls	Yes	Yes

Mechanisms: Other Plan Attributes & Observational Mortality

	(1)	(2)	(3)	(4)	(5)
	Panel A: OLS (Observational Mortality)				
Star Rating	0.0042				
	(0.0003)				
Premium		0.0048			0.0051
		(0.0005)			(0.0005)
Has Donut Hole Coverage			−0.0021		−0.0024
			(0.0003)		(0.0003)
Medical Loss Ratio				0.0142	0.0087
				(0.0035)	(0.0033)

BACK

Estimating Demand for Mortality Effects

In a standard discrete choice model, can invert market shares to recover mean utility: $\ln(s_j) - \ln(s_0) = \delta_j = \alpha p_j + \tau \beta_j + \psi_j$

- Challenge: β_j is unobserved and maybe correlated with p_j

Given a price elasticity (and implied α), we can our forecast IV framework to implicitly regress β_j on $\delta_j - \alpha p_j = \tau \beta_j + \psi_j$

- Reversing this regression with our lower bound on $Var(\beta_j)$ allow us to bound the utility coefficient τ
- Together with α, we obtain the implicit willingness to pay for a one percentage point decrease in plan mortality effects

BACK

Simulations

What happens if we assign everyone to the best plan in their choice set?

- No capacity constraints, partial equilibrium

Be careful, might naively think we can compute:

$$\lambda(\mu_{curr} - \min_j \mu_j)$$

This is wrong:

$$E(\beta_j|\mu_j) = \lambda_j \mu_j$$

$$E(\mu_j|\min_j \mu_j) \neq \lambda_j \min_j \mu_j$$

Simulations

We need to account for the fact that even the shrunken observational measure is noisy

1. Simulate noisy $\hat{\mu}_j = \mu_j + e_j$ and empirical Bayes posterior u_j
2. Identify the smallest $u_{j\min}$ in each market
3. Compute $\lambda(\mu_{curr} - \mu_{jmin(i)})$, the predicted change in mortality from redirecting consumers

BACK