Customer Price Sensitivities in Competitive Automobile Insurance Markets

Robert Matthijs Verschuren

ASTIN Online Colloquium

May 19, 2021

Universiteit van Amsterdam
Amsterdam School of Economics

Motivation

- Traditional cost-based pricing in non-life insurance aimed at directly increasing profits
- In reality, also indirect effects due to highly competitive markets:
  (i) *Easy to compare insurers online*
  (ii) *Relatively low switching costs*
  (iii) *Marketing campaigns to attract new customers*
- Interested therefore in a more demand-based method but requires expected response to alternative, counterfactual offers

Causal inference problem

- Confounding in treatment assignment and response mechanism:

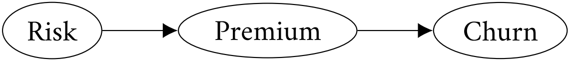

Causal inference problem

- Confounding in treatment assignment and response mechanism:

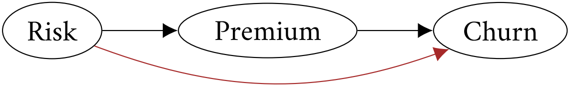

Causal inference problem

- Confounding in treatment assignment and response mechanism:

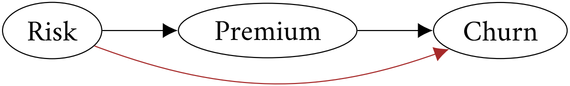

- In addition, premia are not offered at random in practice
  → *So risk characteristics will be insufficiently balanced across treatments*
- Causal inference solution by Guelman and Guillén (2014):
  (i) *Discretize percentage premium changes*
  (ii) *Impute counterfactual responses with propensity score matching*
  (iii) *Optimize next period's profit given predicted responses*

Relevant previous studies

- Causal inference framework:
  (i) *Discrete treatments*
  (Rosenbaum and Rubin, 1983; Rubin, 1997; Morgan and Winship, 2007; Guo and Fraser, 2009; Rosenbaum, 2010; McCaffrey et al., 2013; Guelman and Guillén, 2014; Wager and Athey, 2018)
  (ii) *Continuous treatments*
  (Hirano and Imbens, 2004; Imai and Van Dyk, 2004; Fryges and Wagner, 2008; Guardabascio and Ventura, 2014; Zhu et al., 2015; Kreif et al., 2015; Zhao et al., 2018)
- Applications of continuous treatment framework sparse in non-life insurance

Customer price sensitivity

- Let random variable $Y_i(t) \in \{0, 1\}$ denote policy i's churning response to any potential rate change, or treatment, $t \in \mathcal{T}$
- Actually assigned treatment given by T_i with risk characteristics X_i
- Causal inference relies on two assumptions:
  - (i) *Actual rate changes depend only on the observed risk characteristics (weak unconfoundedness):* $Y_i(t) \perp T_i | X_i \quad \forall t \in \mathcal{T}$
  - (ii) *Each customer has non-zero probability of receiving every rate change (common support):* $0 < \pi(t, X_i) := \mathbb{P}[T_i = t | X_i] < 1 \quad \forall t \in \mathcal{T}$
- Together this allows identification of average treatment effects without bias by controlling for confounders (strong ignorability)

Discrete treatment categories

- Discretize observed treatments T_i in T categories $\{t_1, \ldots, t_T\}$
- Match customers based on similarity:
  (i) *Challenging or even impossible with many risk characteristics*
  (ii) *Propensity score* $\pi(t_s, X_i)$ *one-dimensional alternative, sufficient due to balancing property:* $T_i \perp X_i | \pi(t_s, X_i) \quad \forall s \in \{1, \ldots, T\}$
  (iii) *If strong ignorability holds conditional on* X_i *then also conditional on* π
- Propensity score to explain treatments T_i as accurately as possible
  → *XGBoost of Chen and Guestrin (2016) is appropriate for this* Details
- Impute counterfactual responses from propensity score matches
  → *Multiple imputation to (partially) include response uncertainty* Details
- Form global response model from both observed and imputed data

Continuous treatment doses

- Continuum of potential treatment doses $\mathcal{T} = [\underline{T}, \overline{T}]$
- Generalized propensity score $\pi(T_i, X_i)$ continuous in treatment
- Balancing and strong ignorability property still valid
- Traditional global response model only conditional on $\pi(T_i, X_i)$:

  (i) $$\mathbb{E}\left[Y_i(T_i)|\pi(T_i, X_i)\right] = \alpha_0 + \alpha_1\pi(T_i, X_i) + \alpha_2\pi(T_i, X_i)^2 + \alpha_3 T_i + \alpha_4 T_i^2 + \alpha_5\pi(T_i, X_i)T_i$$

  (ii) $$\widehat{\mathbb{E}\left[Y(t)\right]} = \frac{1}{N}\sum_{i=1}^{N}\big(\hat{\alpha}_0 + \hat{\alpha}_1\hat{\pi}(t, X_i) + \hat{\alpha}_2\hat{\pi}(t, X_i)^2 + \hat{\alpha}_3 t + \hat{\alpha}_4 t^2 + \hat{\alpha}_5\hat{\pi}(t, X_i)t\big)$$

- No direct causal interpretation of global response model

  → *So can use XGBoost for this as well*

  → *Can still use it to predict individual potential responses*

Dutch automobile insurance portfolio

- Applied to individual policy renewals from 2017-2019:
  - (i) 71,522 *policies by* 30,738 *customers with* 20,649 (28.87%) *lapses*
  - (ii) *Churn defined as lapsing before renewal but after receiving rate change*
  - (iii) *Rate change quintile intervals* $[-9.28\%, 1.53\%]$, $(1.53\%, 6.06\%]$, $(6.06\%, 8.58\%]$, $(8.58\%, 12.58\%]$ *and* $(12.58\%, 27.01\%]$
- Includes premia offered by six largest competitors:
  - (i) *Competitiveness* $(B - A)/A$ *of each renewal offer before any rate changes* (A) *relative to current cheapest competing offer* (B)
  - (ii) *Underpricedness* $(D_z - C)^+$ *of renewal offers after autonomous corrections but before competitiveness adjustments* (C) *compared to cheapest* (D_1) *and second-cheapest competitor* (D_2)

Observed churn proportions

- Stable, slowly increasing churn ratios
- Relatively large inflection at small rate changes
  - → *Indication of let sleeping dogs lie effect*

Churn proportions, rate changes (left) - competitiveness (right)

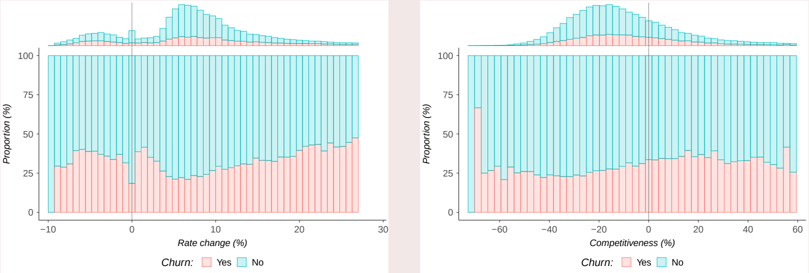

Propensity score matching

- Increase in balance of risk factors, up to 95%
- Discrete approach improves balance considerably more:
  - (i) *Optimizes the rate change interval assignments directly*
  - (ii) *Only has to distinguish between five categories*
- Common support and hence strong ignorability hold

Average (generalized) propensity scores

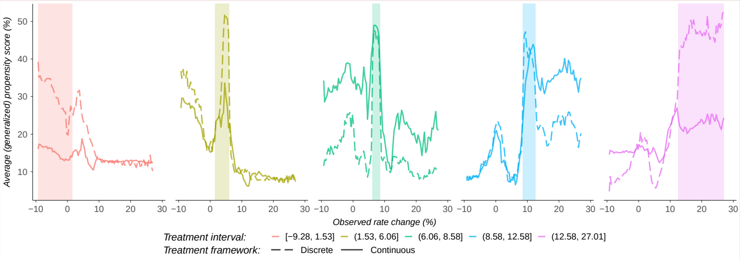

Customer price sensitivities

- Intuition:
  - (i) *More worthwhile to switch at higher or very small rate changes*
  - (ii) *Comparison of insurers more likely for very competitive policies*
  - (iii) *New policies become relatively expensive for very good and bad risks*

Average customer churn, discrete (left) - continuous (right)

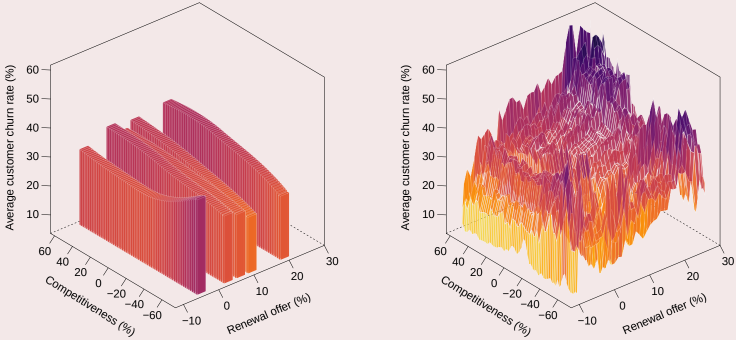

Efficient frontier

- Constrained optimization of next year's expected profit: Details
  - (i) *Trade-off between customer churn and profit potential*
  - (ii) *Only small improvement due to multiple imputation*
  - (iii) *Substantially more profit in continuous approach due to XGBoost and ability to distinguish between rate changes in each interval*

Efficient frontier, discrete (left) - continuous (right)

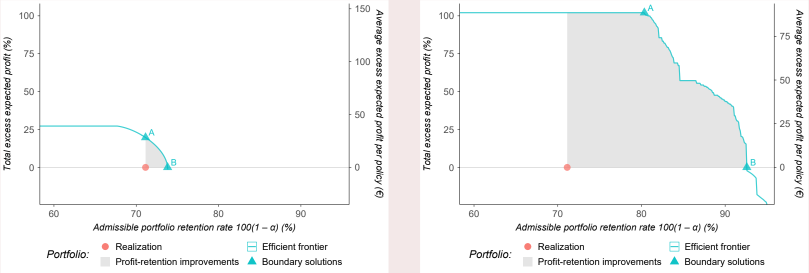

Efficient frontier

- Constrained optimization of next year's expected profit: Details
  - (i) *Trade-off between customer churn and profit potential*
  - (ii) *Only small improvement due to multiple imputation*
  - (iii) *Substantially more profit in continuous approach due to XGBoost and ability to distinguish between rate changes in each interval*

Efficient frontier, discrete one imputation (left) - continuous (right)

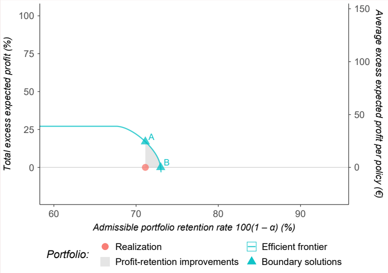

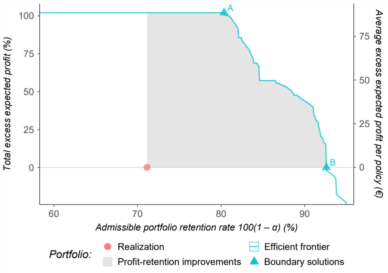

Efficient frontier

- Constrained optimization of next year's expected profit: Details
  - (i) *Trade-off between customer churn and profit potential*
  - (ii) *Only small improvement due to multiple imputation*
  - (iii) *Substantially more profit in continuous approach due to XGBoost and ability to distinguish between rate changes in each interval*

Efficient frontier, discrete XGBoost (left) - continuous (right)

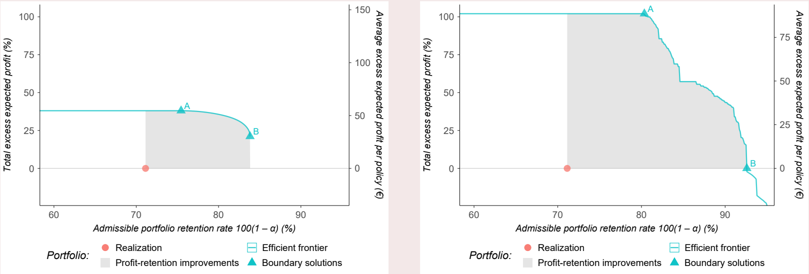

Efficient frontier

- Constrained optimization of next year's expected profit: Details
  (i) *Trade-off between customer churn and profit potential*
  (ii) *Only small improvement due to multiple imputation*
  (iii) *Substantially more profit in continuous approach due to XGBoost and ability to distinguish between rate changes in each interval*

Efficient frontier, discrete XGBoost (left) - continuous restricted (right)

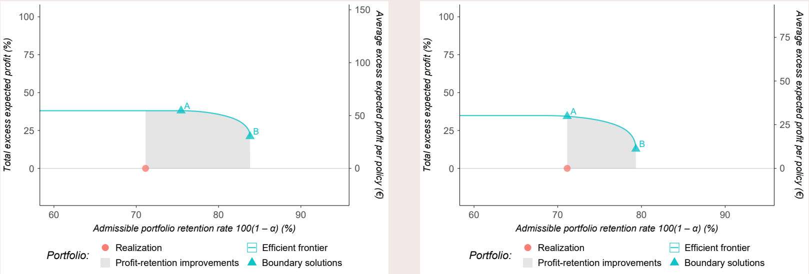

Multi-period renewal optimization

- Constrained optimization of expected profit over τ periods: Details
  (i) *Slightly lower rate changes in first period due to temporal feedback*
  (ii) *Substantially more profit possible, especially in continuous approach*

Expected yearly profit and churn, discrete (left) - continuous (right)

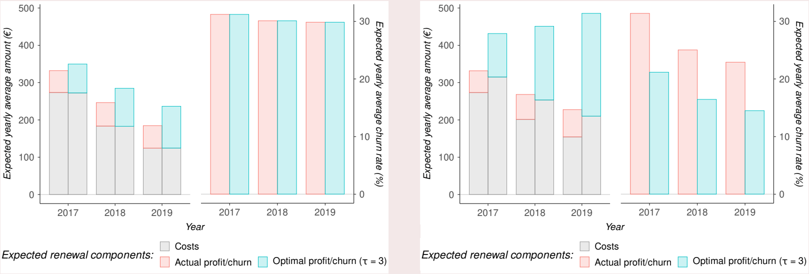

Conclusion

- Shift from cost-based pricing to demand-based pricing
- Causal inference approach required to adjust for confounding
- Application to automobile insurance shows:
  (i) *Policy's competitiveness crucial for price sensitivity*
  (ii) *XGBoost more appropriate than traditional logistic regression*
  (iii) *Substantially more profit can be gained than realized, also already with less churn and in particular using continuous approach*
  (iv) *Temporal feedback of previous rate changes on future demand enabled through competitiveness*

Future research

- Introduce risk characteristics in matching procedure
- Primary focus on logistic GLMs and XGBoost:
  (i) *Compare to alternative machine learning methods, such as (causal) random forests, (deep) neural networks or support vector machines*
  (ii) *Consider ensemble of various (machine learning) models*

Customer Price Sensitivities in Competitive Automobile Insurance Markets

Robert Matthijs Verschuren

ASTIN Online Colloquium

May 19, 2021

Universiteit van Amsterdam
Amsterdam School of Economics

References

CHEN, T. and GUESTRIN, C. (2016). XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 785–794. ACM, New York.

FRYGES, H. and WAGNER, J. (2008). Exports and productivity growth: First evidence from a continuous treatment approach. *Review of World Economics*, **144**(4):695–722.

GUARDABASCIO, B. and VENTURA, M. (2014). Estimating the dose-response function through a generalized linear model approach. *The Stata Journal*, **14**(1):141–158.

GUELMAN, L. and GUILLÉN, M. (2014). A causal inference approach to measure price elasticity in automobile insurance. *Expert Systems with Applications*, **41**(2):387–396.

GUO, S. and FRASER, M. W. (2009). *Propensity Score Analysis: Statistical Methods and Applications*. Sage, Thousand Oaks, CA.

HIRANO, K. and IMBENS, G. W. (2004). The propensity score with continuous treatments. In W.A. SHEWHART, S.S. WILKS, A. G. and MENG, X., editors, *Applied Bayesian Modeling and Causal Inference from Incomplete-Data Perspectives*, volume 2074, pages 73–84. John Wiley & Sons, Ltd, Chichester, UK.

IMAI, K. and VAN DYK, D. A. (2004). Causal inference with general treatment regimes: Generalizing the propensity score. *Journal of the American Statistical Association*, **99**(467):854–866.

KREIF, N., GRIEVE, R., DÍAZ, I., and HARRISON, D. (2015). Evaluation of the effect of a continuous treatment: A machine learning approach with an application to treatment for traumatic brain injury. *Health Economics*, **24**(9):1213–1228.

References

McCaffrey, D. F., Griffin, B. A., Almirall, D., Slaughter, M. E., Ramchand, R., and Burgette, L. F. (2013). A tutorial on propensity score estimation for multiple treatments using generalized boosted models. *Statistics in Medicine*, **32**(19):3388–3414.

Morgan, S. L. and Winship, C. (2007). *Counterfactuals and Causal Inference*. Cambridge University Press, New York.

Rosenbaum, P. R. (2010). *Design of Observational Studies*. Springer, New York.

Rosenbaum, P. R. and Rubin, D. B. (1983). The central role of the propensity score in observational studies for causal effects. *Biometrika*, **70**(1):41–55.

Rubin, D. B. (1997). Estimating causal effects from large data sets using propensity scores. *Annals of Internal Medicine*, **127**(8):757–763.

Wager, S. and Athey, S. (2018). Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, **113**(523):1228–1242.

Zhao, S., Van Dyk, D. A., and Imai, K. (2018). Causal inference in observational studies with non-binary treatments. Working Paper, October.

Zhu, Y., Coffman, D. L., and Ghosh, D. (2015). A boosting algorithm for estimating generalized propensity scores with continuous treatments. *Journal of Causal Inference*, **3**(1):25–40.

XGBoost and multiple imputation

- Gradient Boosting Models for propensity score: Return
  (i) *Combines many weak learners to learn from errors of previous learners*
  (ii) *Flexible non-linear effects of risk factors*
  (iii) *Identification of complex interactions in tree-learning algorithm*
  (iv) *Built-in variable selection procedure*
- XGBoost of Chen and Guestrin (2016) more flexible and faster
- Multiple imputation to (partially) include response uncertainty:
  (i) *Randomly sample M counterfactual responses from I closest matches*
  (ii) *Combine global response estimates $\delta_m = (\beta_m, \gamma_m)$:*
  $\bar{\delta} = \frac{1}{M}\sum_{m=1}^{M}\hat{\delta}_m$ *and* $\text{Var}(\bar{\delta}) = \bar{W} + \left(1 + \frac{1}{M}\right) B$
  (iii) *Within-imputation, or parameter, uncertainty:*
  $\bar{W} = \frac{1}{M}\sum_{m=1}^{M}\hat{W}_m$
  (iv) *Between-imputation, or imputation, uncertainty:*
  $B = \frac{1}{M-1}\sum_{m=1}^{M}\left(\hat{\delta}_m - \bar{\delta}\right)'\left(\hat{\delta}_m - \bar{\delta}\right)$

Constrained renewal optimization

- Constrained optimization of next year's expected profit: Return

$$\max_{\{t_i\}_{i=1}^N \in \mathcal{T}^N} \left\{ \sum_{i=1}^N \left(1 - \hat{Y}_i(t_i)\right) \left(\texttt{Premium}_i - \texttt{Costs}_i\right) \right\} \text{ s.t. } \frac{1}{N} \sum_{i=1}^N \hat{Y}_i(t_i) \leq \alpha$$

- Constrained optimization of expected profit over τ periods: Return

$$\max_{\{t_{i,j}\}_{i=1,j=1}^{N,\tau} \in \mathcal{T}^{N\tau}} \left\{ \sum_{i=1}^N \sum_{j=1}^\tau \left(\prod_{h=1}^j \left[1 - \hat{Y}_i(t_{i,1}, \ldots, t_{i,h})\right] \right) \left(\texttt{Premium}_{i,j} - \texttt{Costs}_{i,j}\right) \right\}$$

$$\text{s.t. } \frac{1}{N} \sum_{i=1}^N \hat{Y}_i(t_{i,1}, \ldots, t_{i,j}) \leq \alpha_j \quad \text{for} \quad j = 1, \ldots, \tau$$

→ *Overall churn rate limited to average churn rate expected for actual renewal offers, or* $\alpha_j = \frac{1}{N} \sum_{i=1}^N \hat{Y}_i(T_{i,1}, \ldots, T_{i,j})$ *for* $j = 1, \ldots, \tau$