Unbiased Estimation of the Economic Value of Pricing Strategies

Dimitri Semenovich and James Petterson

IAG

Introduction I

Tactical pricing of insurance products can often be effectively carried out adopting the so called "semi-myopic" customer model. Under this model customers have private willingness-to-pay, drawn from a distribution potentially dependent on their observed characteristics, and are taken to arrive at random. If customers' willingness-to-pay exceeds the proposed premium, they purchase the policy.

A key assumption made in most real-world pricing systems is that the willingness-to-pay distributions (or, equivalently, demand functions), as well as the cost of providing cover, are known exactly for each customer. While this makes the problem more tractable it also introduces substantial statistical difficulties as we will show in this talk.

Introduction II

The prevalent approach (Murphy et al., 2000; Krikler et al., 2004) follows along these lines:

1. specify a demand model and a cost of cover model,
2. estimate their parameters using sales, exposure and claims cost data,
3. set up an optimal pricing problem using the above two models, with individual contract prices as decision variables,
4. use the objective/constraints values corresponding to the solution as estimates of the economic value of the resulting pricing strategy to the firm.

We demonstrate that this framework is only adequate if both demand and risk cost model estimates are unbiased and have minimal prediction uncertainty. Once realistic assumptions are adopted, however, the economic value is overstated to a considerable degree. Inflation factors between 1.2 and 5 are consistent with our experience.

Traditional tests for goodness of fit, predictive accuracy and calibration used to validate risk cost and demand models, are ultimately neither necessary nor sufficient to ensure correct estimation of the economic value. We propose a new family of unbiased evaluation metrics for pricing procedures, inspired by work in uplift modeling and reinforcement learning.

Motivating Example I

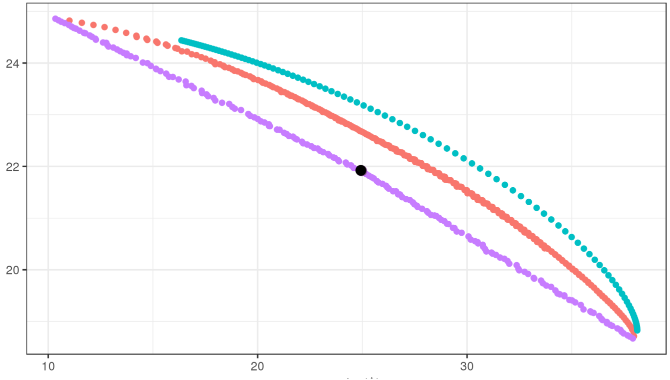

Figure 1: Expected effects on conversions (x-axis) and margin (y-axis) as a result to $\pm 10\%$ per quote premium change, evaluated using a demand model. **Black** dot denotes the current portfolio position. Purple frontier represents operating points achievable with a base rate change. Red frontier indicates the effect of moving premiums towards a target loss ratio. Green frontier is a biased estimate derived using the traditional optimisation procedure.

Motivating Example II

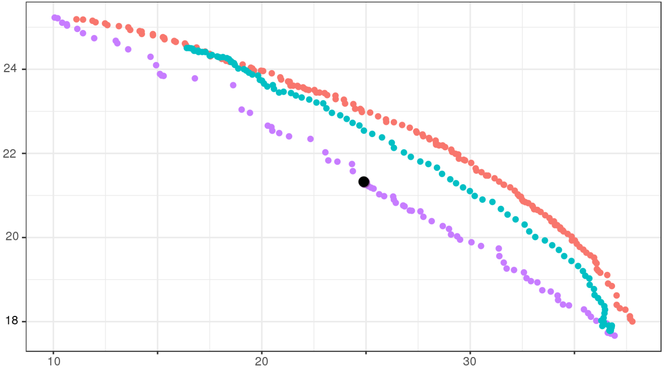

Figure 2: Same premium changes as in the previous plot but evaluated using the proposed unbiased estimator. Note that the order of the frontiers is reversed, the simpler profitability based adjustment is now expected to outperform the "optimisation".

Single period optimal pricing problem I

In a single period optimal pricing problem we seek to maximise the total expected profit objective for a cohort of n policies subject to a constraint on the minimum retention level q.

For the i-th policy with risk characteristics $\mathbf{x}_i$ the proposed premium is denoted p_i, the demand is a random variable with mean $d(p_i, \mathbf{x}_i)$, indexed by premium, cost of claims is a random variable with mean $c(\mathbf{x}_i)$ and expected profit is given by $r_i(p_i) = \big(p_i - c(\mathbf{x}_i)\big) d(p_i, \mathbf{x}_i)$, which yields:

$$
\begin{aligned}
\underset{p_1,\ldots,p_n}{\text{maximise}} \quad & \sum_{i=1}^{n} \big(p_i - c(\mathbf{x}_i)\big) d(p_i, \mathbf{x}_i) = \sum_{i=1}^{n} r(p_i, \mathbf{x}_i) = r(\mathbf{p}) \\
\text{subject to} \quad & \sum_{i=1}^{n} d_i(p_i, \mathbf{x}_i) = q.
\end{aligned}
\tag{1}
$$

We will refer to the solution of this problem as $\mathbf{p}^*$, with optimal underwriting profit given by $r(\mathbf{p}^*)$.

Single period optimal pricing problem II

In practice, we do not have access to the parametrised expectations of demand and cost random variables and instead we are working with their estimates $\hat{d}(p_i, \mathbf{x}_i)$ and $\hat{c}(\mathbf{x}_i)$ respectively. It is common practice to still use the optimisation problem of the same form as (1):

$$\begin{aligned} \underset{p_1,\ldots,p_n}{\text{maximise}} \quad & \sum_{i=1}^{n} \big(p_i - \hat{c}(\mathbf{x}_i)\big)\hat{d}(p_i, \mathbf{x}_i) = \sum_{i=1}^{n} \hat{r}(p_i, \mathbf{x}_i) = \hat{r}(\mathbf{p}) \\ \text{subject to} \quad & \sum_{i=1}^{n} \hat{d}_i(p_i, \mathbf{x}_i) = q. \end{aligned} \tag{2}$$

The solution to this surrogate problem is denoted $\hat{\mathbf{p}}^*$ and the objective value, which is often taken as the estimate of $r(\mathbf{p}^*)$ is $\hat{r}(\hat{\mathbf{p}}^*)$. We will later show that under realistic assumptions on model error the following obtains:

$$r(\hat{\mathbf{p}}^*) < r(\mathbf{p}^*) < \hat{r}(\hat{\mathbf{p}}^*). \tag{3}$$

Single period optimal pricing problem III

Before examining the properties of the naive estimate of the objective value $\hat{r}(\hat{\mathbf{p}}^*)$, we observe that the problem (1) can be rewritten using policy demand as the decision variable, assuming one-to-one correspondence between premium and demand $p(d_i, \mathbf{x}_i) = d^{-1}(d_i, \mathbf{x}_i)$:

$$\begin{aligned} \underset{d_1,\dots,d_n}{\text{maximise}} \quad & \sum_{i=1}^{n} \big(p(d_i, \mathbf{x}_i) - c(\mathbf{x}_i)\big) d_i = r(\mathbf{d}) \\ \text{subject to} \quad & \sum_{i=1}^{n} d_i \ = q. \end{aligned} \tag{4}$$

We can then formulate the Lagrangian:

$$L(d_1, \dots, d_n, \lambda) = \sum_{i=1}^{n} \big(p(d_i, \mathbf{x}_i) - c(\mathbf{x}_i)\big) d_i + \lambda \Big(\sum_{i=1}^{n} d_i - q \Big)$$

Single period optimal pricing problem IV

and write the optimality conditions as:

$$\frac{\partial L}{\partial d_i} = 0, \quad 1 \leq i \leq n,$$
$$\frac{\partial L}{\partial \lambda} = 0.$$

Observe that $\frac{\partial L}{\partial d_i} = \frac{\partial r}{\partial d_i} + \lambda$ and that therefore if the portfolio is priced optimally, marginal profit with respect to demand for each policy is constant: $\frac{\partial R}{\partial d_i} = -\lambda$.

This condition is intuitive – should $\frac{\partial R}{\partial d_i} \neq \frac{\partial R}{\partial d_j}$ for some i and j, we can reallocate demand between contracts i and j in such a way as to increase total profit.

Effects of model uncertainty I

We now demonstate that the surrogate optimisation problem (2) is subject to a facet of the phenomenon that often causes overparametrised statistical models to "overfit" in sample.

The effect of model uncertainty can be studied more easily if instead of (2) we consider a *local linearisation* (i.e. first order Taylor expansion) of the demand parametrised problem (4) around demand vector $\mathbf{d}^{(0)}$ instead:

$$\begin{aligned} \underset{w_1,\dots,w_n}{\text{maximise}} \quad & \sum_{i=1}^{n} \left(r\left(d_i^{(0)}, \mathbf{x}_i\right) + \frac{\partial r}{\partial d_i} w_i \right) = r\left(\mathbf{d}^{(0)}\right) + r(\mathbf{w}) \\ \text{subject to} \quad & \sum_{i=1}^{n} \left(d_i^{(0)} + w_i\right) = q, \\ & -1 \leq w_i \leq 1. \end{aligned} \tag{5}$$

Omitting the constant term $r(\mathbf{d}_o)$ from the objective and observing that $\sum_{i=1}^{n} d_i^{(0)} = q$, we can simplify the above as:

Effects of model uncertainty II

$$\begin{aligned} \underset{w_1,\ldots,w_n}{\text{maximise}} \quad & \sum_{i=1}^{n} \frac{\partial r}{\partial d_i} w_i = r(\mathbf{w}) \\ \text{subject to} \quad & \sum_{i=1}^{n} w_i = 0, \\ & -1 \leq w_i \leq 1. \end{aligned} \tag{6}$$

It is intuitive that the solution $\mathbf{w}^*$ is attained if we set $w_i^* = 1$ for those policies i where $\frac{\partial r}{\partial d_i}$ is larger than M, the median entry of $\left(\frac{\partial r}{\partial d_1}, \ldots, \frac{\partial r}{\partial d_n}\right)$, and $w_i^* = -1$ where it is smaller.

The objective value corresponsing to $\mathbf{w}^*$ is then given by $\sum_{i=1}^{n} \left| \frac{\partial r}{\partial d_i} - M \right|$. It represents improvement to profit r attainable by perturbing demand by no more than one unit for each contract relative to the initial demand vector $\mathbf{d}$.

Notice that if we substitute a noisy estimate of marginal profit $\frac{\hat{\partial r}}{\partial d} = \frac{\partial r}{\partial d} + \epsilon$, our view of expected profit improvements can generally only go up. This means that any model uncertainty will result in statistically biased estimates of expected profit.

Effects of model uncertainty III

Now we attempt to quantify this bias. This will require further assumptions:

$$\begin{aligned} \epsilon &\sim \mathcal{N}(0, \sigma_a), \\ \frac{\partial R}{\partial d} &\sim \mathcal{N}(0, \sigma_b). \end{aligned}$$

Note the slight abuse of notation, profit function r has become a random variable R. What is the degradation in true performance and how over-optimistic do we become as the noise parameter σ_a is increased?

Effects of model uncertainty IV

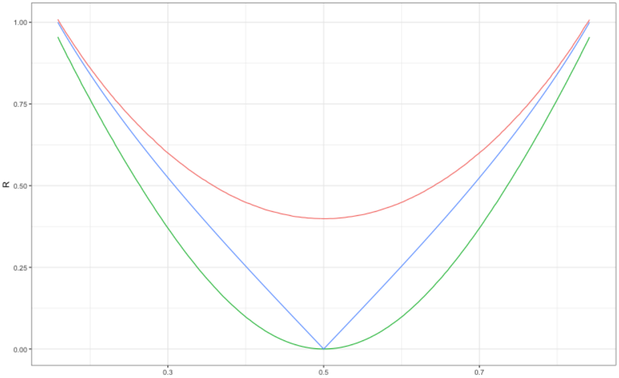

Figure 3: A numerical example showing the bias inherent in the traditional "optimal" pricing procedures. The x axis corresponds to the quantiles of the true marginal profit of a policy and the y axis to the profit either achieved or estimated. The area under the blue line represents the total profit improvement realisable if the true marginal profit with respect to demand is known. The area under the green line shows the profit attained if the noisy estimate of marginal profit is used to guide pricing decisions. Finally the area under the red line is the *biased estimate* of profit that would be achieved. The gap between red and green lines represents total bias in traditional optimal pricing.

The three profit estimates I

We can now examine the three cases (see paper for details):

1. Decision and performance estimate are based on the true marginal profit $\frac{\partial R}{\partial d}$:

$$\mathbb{E}R(w^*) = \frac{2\sigma_b}{\sqrt{2\pi}}. \tag{7}$$

2. Decision is based on a noisy estimator $\frac{\partial R}{\partial d} + \epsilon$, but we measure the profit using the true metric:

$$\mathbb{E}R(\hat{w}^*) = \frac{2}{\sqrt{2\pi}} \frac{{\sigma_b}^2}{\sqrt{{\sigma_a}^2 + {\sigma_b}^2}}. \tag{8}$$

3. Decision and profit estimates are both based on the noisy estimator $\frac{\partial R}{\partial d} + \epsilon$.

$$\mathbb{E}\hat{R}(\hat{w}^*) = \frac{2}{\sqrt{2\pi}} \sqrt{{\sigma_a}^2 + {\sigma_b}^2}. \tag{9}$$

The three profit estimates II

This provides the following decomposition:

$$\mathbb{E}\hat{R}(\hat{w}^*) = \frac{2}{\sqrt{2\pi}} \frac{\overbrace{{\sigma_a}^2}^{\text{contribution of noise on metric estimation}} + {\sigma_b}^2}{\sqrt{\underbrace{{\sigma_a}^2}_{\text{contribution noise on decision}} + {\sigma_b}^2}}. \tag{10}$$

We observe that when $\sigma_a = 0$ we recover (7), adding noise to the decision criterion reduces the expected value of profit R and adding noise to the evaluation metric increases it, yielding:

$$\mathbb{E}R(\hat{w}^*) \leq \mathbb{E}R(w^*) \leq \mathbb{E}\hat{R}(\hat{w}^*).$$

Unbiased estimation I

We can construct an unbiased estimator of expected profit if we conduct validation "out of sample".

Assume we have history of sales and claims data in the form $S = \{(\mathbf{x}_i, p_i, d_i, c_i, \psi_i)\}_{i=1}^{N}$, where ψ_i is the propensity estimate of charging premium p_i for risk $\mathbf{x}_i$. In the ideal scenario these propensities are based on active randomisation with known probabilities.

This history has not been used directly to parametrise either demand or claims cost models (and so we can assume individual realisations to be independent of prediction error).

We construct a vector $\hat{\mathbf{p}}^*$ of proposed prices for each policy using a procedure such as (2). An unbiased estimate of profit can be obtained by the so called inverse probability weighted estimator (Horvitz and Thompson, 1952; Dudik et al., 2014):

$$\hat{r}_{\mathrm{IPWE}}(\hat{\mathbf{p}}^*) = \frac{1}{N} \sum_{i=1}^{n} (p_i - c_i) d_i \frac{\mathbb{I}(p_i = \hat{p}_i^*)}{\psi_i}.$$

IPWE example

v_i	p_i	$\hat{p}_i^*$	ψ_i	d_i	c_i	$\hat{r}_i(\hat{p}_i^*)$
250	10	0	0.25	1	200	—
375	0	0	0.50	0	310	$\frac{(375+0-310) \times 0}{0.50}$
500	-10	-10	0.25	1	370	$\frac{(500-10-370) \times 1}{0.25}$
150	-10	-10	0.25	1	120	$\frac{(150-10-120) \times 1}{0.25}$
230	0	10	0.5	1	200	—

Table 1: Example calculation of $\hat{r}_{\text{IPWE}}$ for a premium adjustment $a \in \{-10, 0, 10\}$ where d_i is the flag indicating purchase. Profit is given by $r_i = d_i(v_i + a_i - c_i)$.

Practical considerations

The IPW estimate can be somewhat noisy on small samples. The variance is magnified by the ratio of at least $\frac{1}{\underset{i}{\operatorname{argmax}}\ \psi_i}$:

$$\operatorname{Var}[\hat{r}_{\mathrm{IPWE}}(\hat{p}_i^*)] = \frac{1}{n}\sum_{i=1}^{n}\left[\frac{\mathbb{E}\left[\left((\hat{\mathbf{p}}^* - c_i)d_i\right)^2\right]}{\psi_i} - \mathbb{E}\left[(\hat{p}_i^* - c_i)d_i\right]^2\right].$$

To be used successfully, it is essential that the randomisation of p_i is carried out over the same small set of values in relation to some reference price p_i^0 as that used in the optimisation procedure to derive $\hat{p}_i^*$. In some cases it may also be necessary to substitute c_i with model based value $\hat{c}(\mathbf{x}_i)$.

We note that replacing $\mathbb{I}(p_i = \hat{p}_i^*)$ with a kernel $\kappa(p_i, \hat{p}_i^*)$ satisfying certain properties may substantially reduce this variance while the resulting estimator remains unbiased under only mild assumptions. This will be explored in future work.

Conclusion

In this presentation we have:

- demonstrated that the traditional approach to optimal pricing can significantly overstate benefits,
- derived correction terms in a simplified model, and
- proposed an unbiased validation procedure, equivalent to out-of-sample testing of predictive models.

References

Dudik M., Erhan D., Langford J., and Li, L., Doubly Robust Policy Evaluation and Optimization, *Statistical Science*, 2014, 29, pp. 485– 511.

Horvitz, D. G., Thompson, D. J., A generalization of sampling without replacement from a finite universe, *Journal of the American Statistical Association*, 1952, 47, pp. 663–685.

Krikler S., Dolberger D., and Eckel J., Method and tools for insurance price and revenue optimisation, *Journal of Financial Services Marketing*, 2004, 9, pp. 68–79.

Murphy K. P., Brockman M. J., and Lee P. K. W., Using generalized linear models to build dynamic pricing systems for personal lines insurance, *Casualty Actuarial Society Forum*, Winter 2000, pp. 107–140.