GLMs – the Good, the Bad, and the Ugly

Casualty Actuaries of the Southeast

30 September 2009

Christopher Cooksey, FCAS, MAAA

EagleEye Analytics

Agenda

1. A Brief History of GLMs
2. The Good – what GLMs do well
3. The Bad – what GLMs don't do well
4. The Ugly – what GLMs can't do
5. Solutions

Section 1

GLM History

A Brief History of GLMs

- Formulated by Nelder and Wedderburn in 1972.
- First edition of McCullagh/Nelder book on GLMs in 1983.
- One of the first examples of use in insurance was "Statistical Motor Rating: making effective use of your data" by Brockman and Wright in 1992.
- "Practitioner's Guide to Generalized Linear Models" written in 2007.

Section 2

The Good – what GLMs do well

The Good – what GLMs do well

- There is an established and understood literature.
- There is increasing DOI acceptance.
- There are readily available software solutions.
- GLMs extrapolate over predictor levels with little or no data.
- GLMs provide easily calculated relativities to use as a classification plan.
- GLMs clearly find significant signal in insurance data.

The Good – what GLMs do well

- GLMs are parametric and come with all the advantages of parametric approaches.
  - By assuming you know the form of the "noise" you can do statistical inference to evaluate predictors.
  - You can also provide confidence intervals to communicate the inherent uncertainty in the output.
  - Parametric approaches are very accurate when the assumptions hold.

Section 3

The Bad – what GLMs don’t do well

The Bad – what GLMs don't do well

- The assumptions underlying GLMs may not hold.
- Investigating this issue takes time, as do corrections to the basic assumptions (if necessary).
- Issues include…
  - Independence of the data
  - Appropriateness of the link function
  - Appropriateness of the error function
  - Predictiveness of the model

The Bad – what GLMs don't do well

One assumption is that the data is independent.

- Normally not a bad assumption, at least for frequency.
- With severity, size of loss can group around values.
  - Limits can lead to distortions in the size of loss
  - Claims adjusters tend to settle for round numbers.
- The solution to this problem is…?
- This is usually counted as a minor distortion.

The Bad – what GLMs don't do well

Another assumption is that the log link works well for insurance data.

- This can be tested with a Box Cox Transformation (an example of this can be found in the "Practitioner's Guide").
- Use the following link function.

$$g(x) = (x^{\lambda} - 1)/\lambda \qquad \text{when } \lambda \neq 0$$

$$g(x) = \ln(x) \qquad \text{when } \lambda = 0$$

GLMs – the Good, the Bad, and the Ugly

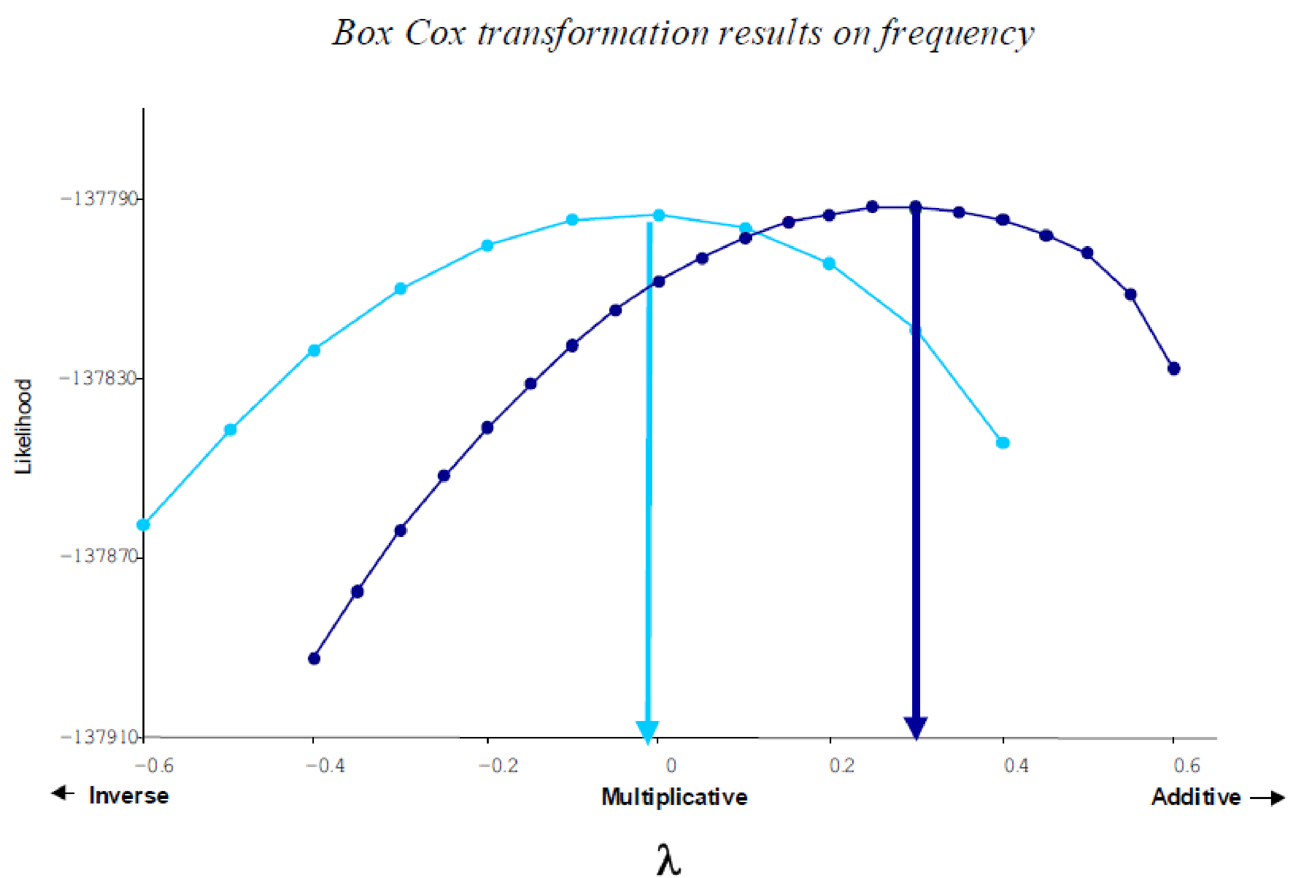

Taken from "A Practitioner's Guide to Generalized Linear Models", Third Edition, page 59.

GLMs – the Good, the Bad, and the Ugly

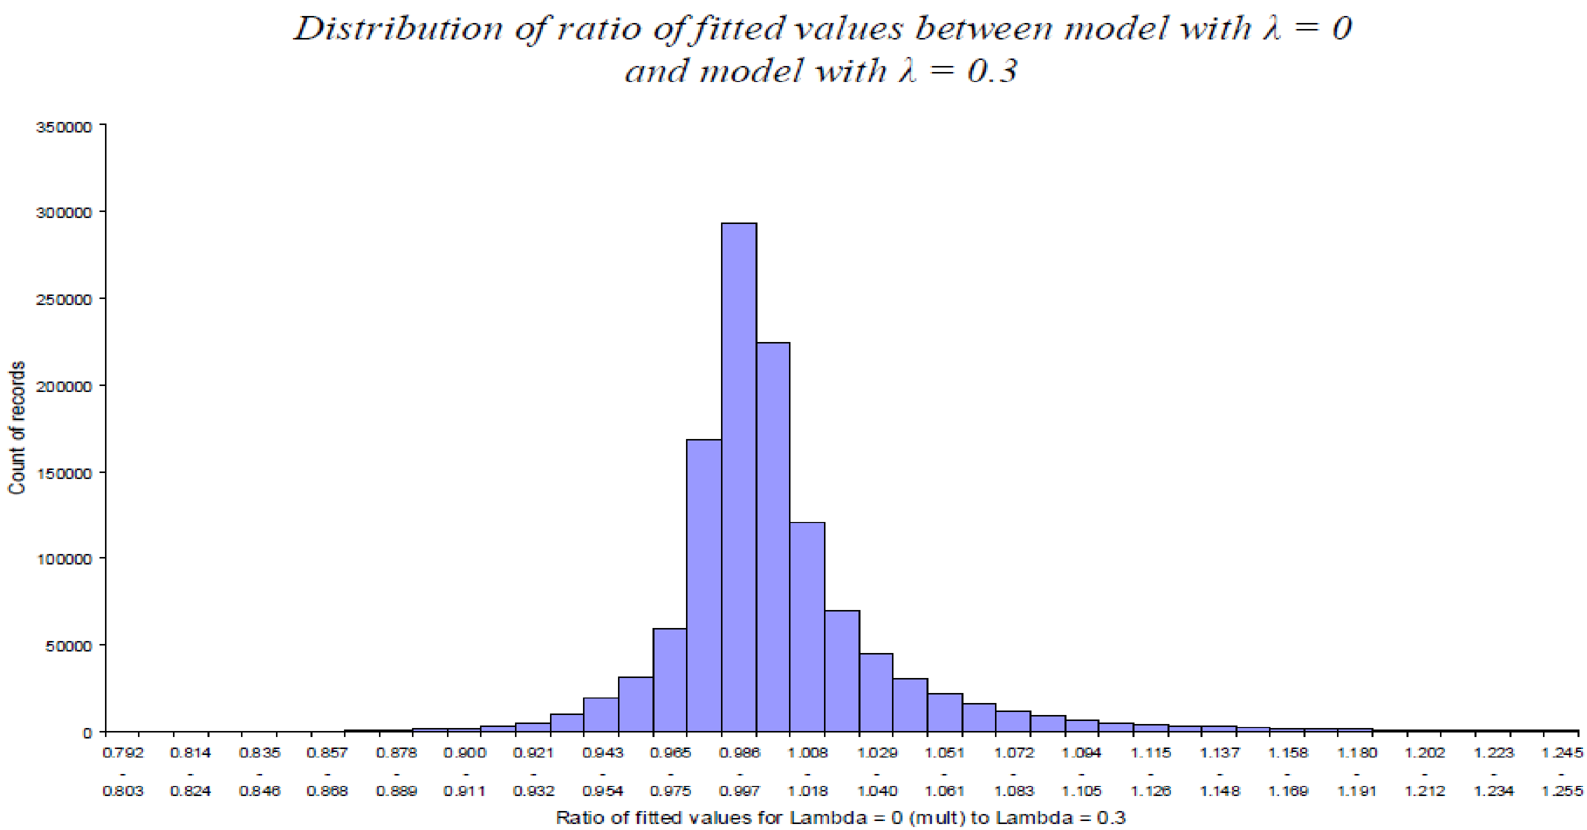

Taken from "A Practitioner's Guide to Generalized Linear Models", Third Edition, page 60.

The Bad – what GLMs don't do well

Another assumption is that the log link works well for insurance data.

- Rarely, if ever, does this test show that the most appropriate model is strictly multiplicative. Usually it shows it to be mostly multiplicative.
- Consequently, multiplicative models are used. This is usually counted as a minor distortion.

The Bad – what GLMs don't do well

A third assumption is that the typical error functions (Poisson and gamma) work well for insurance data.

- This can be tested by looking at the residuals.
- Many things can be done to correct for patterns in residuals, but you rarely, if ever, have perfectly homogeneous residuals.
- Sometimes you can correct for known distortions (zero-inflated Poisson, for example).
- These issues are usually counted as minor distortions.

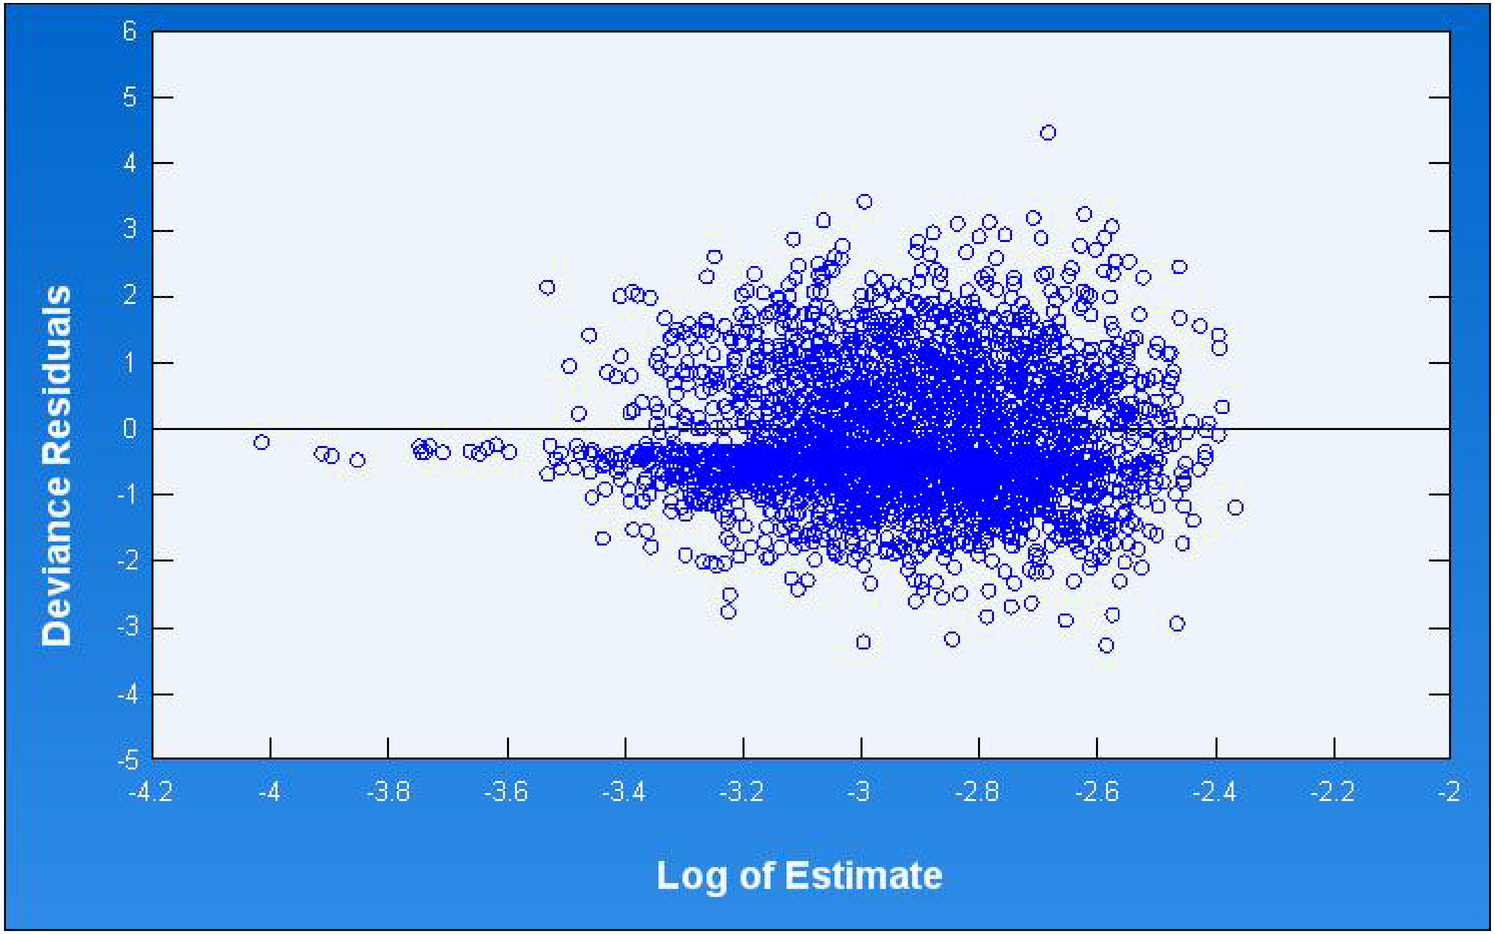
Deviance Residuals
6
5
4
3
2
1
0
-1
-2
-3
-4
-5
-4.2
-4
-3.8
-3.6
-3.4
-3.2
-3
-2.8
-2.6
-2.4
-2.2
-2
Log of Estimate

The Bad – what GLMs don't do well

The predictiveness of the model is an additional assumption that usually isn't considered.

- Certainly people should look at how their final model performs on holdout data.
- One way to do this is to fit the model to the holdout data. Solve for new fitted values.
- Are the new fitted values within the confidence intervals identified by the training data?
- Significance testing tends to overfit models.

The Bad – what GLMs don't do well

The final category of issues with GLMs revolves around the time and effort involved in doing them well.

- GLMs are technically sophisticated, with multiple assumptions and an extensive modeling process.
- Knowledgeable practitioners are required, but supply and demand makes them costly resources.
- Learning from scratch is an alternative, but it too takes an investment of time and money.

The Bad – what GLMs don't do well

The final category of issues with GLMs revolves around the time and effort involved in doing them well.

- Mitigating the model risk posed by GLMs' assumptions also requires time and expertise.
- The trial and error process of determining the design matrix in each case requires significant time.
- Modeling is done separately for each coverage, and likely for both frequency and severity. This multiplies the effort described in the two points above.

Section 4

The Ugly – what GLMs can't do

The Ugly – what GLMs can't do

- GLM model risk can be mitigated but not removed.
- GLMs are linear models. They can only incorporate nonlinear effects through the explicit inclusion of interactions. But GLMs simply do not provide a system for finding all of the relevant interactions. One must know them in advance.
- GLMs are not formulated to find local interactions.
- Combining frequency and severity models leads to an inevitable loss of signal.

The Ugly – what GLMs can't do

GLM model risk can be mitigated but not removed.

- There is no theoretical reason that any given error function should fit precisely.
- Testing shows that insurance data is only "mostly" multiplicative.
- Insurance data is mostly independent.
- There is always some risk that the imperfections of the model assumptions will substantively impact results.

The Ugly – what GLMs can't do

GLMs simply do not provide a system for finding all of the relevant interactions. One must know them in advance.

- It is not practically possible to test through trial and error all possible combinations of two-way interactions, let alone interactions involving three, four, five or more predictors.
- Many people therefore assume there is no such thing as relevant interactions involving more than two or three predictors.

The Ugly – what GLMs can't do

Another problem with interactions is that GLMs are not formulated to find local interactions.

- GLMs use global interactions – the interaction between all levels of two predictors.
- Once this interaction is included, it is possible to note relevant portions and to smooth over irrelevant portions, thus creating local interactions between only certain levels of each predictor.
- This process is only practical for simple interactions.

The Ugly – what GLMs can't do

A final issue is that combining frequency and severity models leads to an inevitable loss of signal.

- After creating models predicting frequency and severity, the models must be combined to find relativities.
- This is usually done by multiplying the predicted frequency and severity of each record into a predicted pure premium, and then regressing relativities onto this.
- This regression is another layer of approximation on top of the already approximate frequency & severity models.

Section 5

Solutions

Solutions

Keeping in mind a realistic view of GLMs, there are at least three possible responses.

1. Continue to rely solely on GLMs
2. Abandon GLMs for some other alternative
3. Find some supplement to cover for GLMs' weaknesses

Solutions

If you stick with GLMs, remember the difficulties…

1. GLMs are parametric. Model assumptions impact the results.
   - Make sure you test the assumptions and consider alternatives to the typical Poisson/frequency and gamma/severity combinations.
2. GLMs provide no good way to explore the universe of possible interactions.
   - Make sure you set aside time to find these. Use intuition and scan your competitors for options. Also look for where your model is out of balance – where observed losses are not close to predicted losses for significant segments of the book of business.

Solutions

If you stick with GLMs, remember the difficulties…

3. There is a loss of predictive power when frequency and severity models are combined into pure premium relativities.
   - Explore ways to improve the fit. Do your own research – will modeling pure premium directly result in a better model?
4. GLMs require a large investment of time and resources.
   - Plan around this. Make sure you have buy-in from all decision-makers in your organizations. Keep them informed. Look for ways to produce actionable results throughout the project, not just at the end.

Solutions

If you abandon GLMs, what else is there?

- Data mining techniques
- Minimum bias
- General Iteration Algorithms (Fu, Wu, 2007)
- Something else???

Solutions

A third approach is to find a supplement to GLMs. Again, consider the difficulties…

1. GLMs are parametric. Model assumptions impact the results.
2. GLMs have no good way to explore the universe of possible interactions.
3. There is a loss of predictive power when frequency and severity models are combined into pure premium relativities.
4. GLMs require a large investment of time and resources.

All you need to find is a nonparametric, nonlinear approach which quickly finds relevant local interactions.

Solutions

What possible candidates exist for accomplishing this? There are many nonparametric approaches and other tools to be found in the fields of data mining and machine learning…

- Neural networks
- MARS
- Decision trees
- CART
- Random forests
- Polynomial networks
- Principle components
- Kernels
- Bagging
- Boosting
- Bootstrapping & resampling
- Activity mining

Solutions

Some issues in developing a solution include…

- Getting the technical expertise in nonparametric solutions.
- One-size-fits-all data mining methods have shown moderate performance on insurance-specific data.
- Better results are found by ensembling multiple methods.
- Nonparametric methods tend to be greedy – significant risk of overfitting.

Section 6

Questions?

Contact Info

Christopher Cooksey, FCAS, MAAA

EagleEye Analytics

ccooksey@eeanalytics.com

www.eeanalytics.com