Paper 161-2008

Price and Revenue Optimization for Banking

Keith Renison, SAS Institute Inc., Portland, OR
Lorry Hardt Jr., SAS Institute Inc., Fort Worth, TX

ABSTRACT

The May 2007 issue of Finance Tech online magazine quoted Bank of America's SVP of pricing strategy, Dan Malouff, as saying "Price optimization is absolutely necessary to deal with the margin compression that we have experienced." Though the quotation seems ubiquitous for most industries—including travel and retail—the concept of price optimization is nascent for financial services (FS). In analyst circles and internal FS analytic practices alike, the concept of price and revenue optimization is quietly—if not secretly—being discussed, as it proves itself to be the next opportunity to leverage analytics for a competitive advantage.

This paper is designed to bring the concepts of customer price sensitivity, demand elasticity, profitability, and price differentiation into a financial services context. This paper will articulate how bid win/loss history can be used to determine demand response to price, its effect on profitability, implications for risk, and how applied analytics can be used to optimize tactical pricing decisions.

INTRODUCTION

This paper is intended to demonstrate the fundamentals of price and revenue optimization (PRO) in a financial services context. We expect you to walk away with an understanding of why it's important, why it's different from other approaches, and where you should begin. Put differently, our intention is not to provide a comprehensive guide to implementation, nor will we defend our approach to a mahogany-clad room filled with men wearing smoking jackets (you know who you are). We will indeed cover some technical concepts as well as SAS® analytic techniques for the technical folks, but will wrap it in a business context. There is a lot of veiled buzz about this in the banking world and surely others will, and are, taking these concepts further than we can here. Indeed, we have set out to provide you with just enough information to break the surface; the rest is up to you.

PRO, as mentioned in the abstract, is not a particularly new idea. A business based on products that expire (airlines or hotels) or products that are valued differently, depending on who is buying them (retail), has a motivation to adjust pricing in order to eke out every bit of revenue humanly possible. A vacant hotel room means no revenue; a customer willing to pay more to rent the Mustang when all you have is the Pontiac G6 means a lost opportunity to charge more for a product with virtually the same overhead. Ever wonder why retail outlet malls are located on the fringe of town? It's because they know if you're willing to drive to the middle of nowhere, you care more about price than someone paying full price at the downtown location. One thing all these businesses have in common is that they understand a customer's ***willingness to pay*** is a critical part of your ability to set a price for your products.

With the financial industry's recent full-court press toward organic growth and retention, as well as the big, bad word "commodity" being used in reference to financial products, it's no wonder we find ourselves looking for new ways to improve profits through pricing. Enter PRO. Banks have a unique opportunity to take PRO further than any of the other aforementioned industries. Why? Because banks know far more about with whom they are doing business, and whether their offers are accepted or declined at a particular price. The following section will outline the basic pieces of this phenomenon, and send you on the road to optimized profitability.

PRICE SENSITIVITY

There are several contexts from which to describe price sensitivity, but for simplicity's sake, price sensitivity can be defined as the price a customer is willing to pay at a given time. So why is sensitivity so vital to price optimization? Imagine for a moment a customer who does not choose your offer because they perceive it as being too expensive. However given what you expect from your potential relationship, you could have offered a lower price and still turned a profit. In other words, you overbid. Conversely, imagine your customer accepted your offer, only for him or her to have been willing to accept a higher price (you underbid). Other pricing techniques often consider things like cost, expected revenue, risk of default, market conditions, minimum margin, etc. However, few banks incorporate the most critical piece: the customer. You can go through all the work of determining the best margin to clear, only to completely miss the mark with your customer's expectations.

BID WIN/LOSS HISTORY

As described above, few industries are as uniquely positioned as financial services to know to whom, how, and when an offer is proposed, as well as the outcome. This is often discussed as bid win/loss history in the academic literature (Phillips, 2006). As we are learning, this history is the linchpin to understanding price sensitivity. Using analytic techniques paired with historical offer data, it's possible to estimate the probability that a customer will accept your offer at a given price. Of course with any predictive modeling, you must make sure your population serves as a quality proxy of future consumer choice, which of course has many complexities we won't dive into here.

MULTIDIMENSIONAL SEGMENTATION

In the larger context, this problem is very multidimensional in nature. Common sense tells us each person brings a distinctive set of decision criteria, so we are caught in a balancing act: to get as close as possible to modeling each unique intersection to support pricing decisions, without unnecessary levels of detail or excessive processing. Plus, many offers are made to customers with whom we have not yet done business. This forces us to assume their future behavior based on people like them. So, at the onset, we stick to the most common dimensions that can affect price sensitivity—and typically within the purview of a financial institution—which are the customer segments to which they belong, the products they are considering, and the channels through which they came. Certainly, dimensions such as life-stage, credit score, service satisfaction, their uncle's advice, and the like, all can play into how much a customer is willing to pay. Indeed, the greater detail you have about these aspects will allow you to apply optimization techniques at more detailed intersections.

FITTING A MODEL TO THE DATA

The probability of a customer accepting an offer at a given price is often not linear. A population typically doesn't gradually lose interest in a product as you raise prices. Rather, it's common to see a high level of acceptance at the lowest prices, followed by a steep drop in demand as you cross some threshold (often near the competitive rates). You will also typically find a small proportion of customers willing to pay any price you offer them. Not you, right? Therefore, a good first step is to fit bid win-loss history data to the SAS logistics function, represented visually as a reverse S-shaped curve, with price ascending on the X-axis and the probability estimate on the Y-axis. Here's an example:

```
DATA BASICWINLOSSDATA;
INPUT OFFER PRICE BINARYSCORE;
DATALINES;
1       13.8    1
2       15.0    1
3       13.9    0
4       15.2    1
...
47      14.9    1
48      15.4    0
49      15.1    1
50      15.2    1
;
PROC LOGISTIC DESCENDING DATA=BASICWINLOSSDATA;
MODEL BINARYSCORE = PRICE / LACKFIT;
OUTPUT OUT=NEW P=PRED L=LOWER U=UPPER;
RUN;

PROC SORT DATA=NEW;
BY PRICE;
symbol1 i = join v=none l=1 w=3 c=black;
symbol2 i = join v=none l=4 w=1 c=red;
symbol3 i = join v=none l=4 w=1 c=green;

PROC GPLOT DATA=NEW;
PLOT PRED*PRICE LOWER*PRICE UPPER*PRICE / OVERLAY;
RUN;
QUIT;
```

The first section of the code sets up an oversimplified sample dataset depicting the acceptance, or rejection of an offer to customers in a specific segment, for a specific channel, for a specific product, for a specific time period. Yes, this is a single cluster and we essentially think all these people are practically the same, relatively speaking. In the real world, we would need far more than 50 observations; however, this sample should allow us to adequately articulate our point. A "1" represents acceptance of the offer, and a "0" indicates they took their business elsewhere. For the price, we picked a range of interest rates not far off from current credit card rates in the market today.

The second section runs a binomial logit function, sorting the data in descending order. The descending option in PROC LOGISTIC is a way to denote the model to utilize 1 as the event instead of 0. The application of MLE (maximum likelihood estimation) within logistic regression gives the ability to predict the chances of accepting an offer (our dependent variable) based on the values of price.

The last section of code is used to produce the plot shown in Figure 1. This shows us that as we raise prices, the probability of customer acceptance declines. There is often a steep drop signifying the customers' reservation price, or prevailing market price, followed by a slow flattening (these are the customers who will pay anything). The dashed lines represent the upper and lower confidence bounds of the prediction.

Figure 1

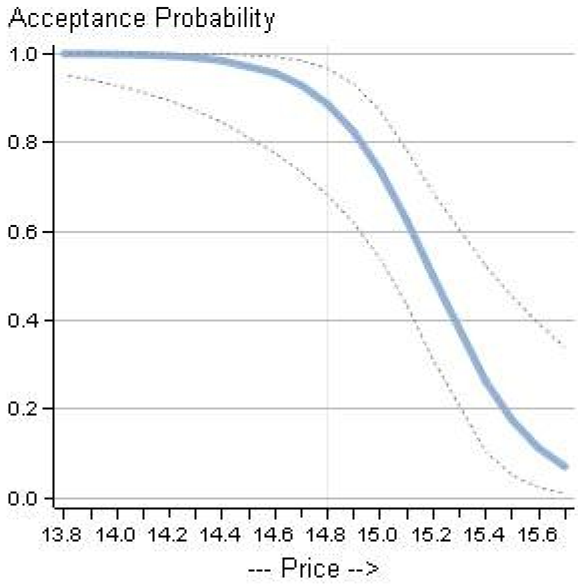

COST, REVENUE, AND PROFIT DYNAMICS

Much like differences in willingness to pay, differences in customer behavior and various risk factors also determine the expected profitability of the relationship if they accept your offer. The following sections explain these elements in greater detail.

ACTIVITY-BASED MANAGEMENT AND INCREMENTAL COST

Activity-based cost modeling (ABM) techniques are becoming a norm in banking (Datamonitor, 2006). We also believe that ABM is one of the most accurate ways to burden your customers with cost. An ABM model can provide a transparent way to take your organization's resources and drive them through the work you do to service your customers and subsequently allocate those costs to the ways in which your customers behave. There are, of course, various techniques and managerial accounting practices that define allocation rules, but for the sake of pricing, only the incremental cost should be included in your optimization model. This is because general overhead typically allocated to customers to make ABM costs match your general ledger dilutes the cost incurred by gaining that relationship, and saved by not gaining it. Put differently, general overhead will always be there, whether you add a new customer or not, and it shouldn't be part of the pricing decision.

REVENUE AND PROFITABILITY

New methods in modeling profitability have finally come to the fore. Traditionally, true customer profitability has been obscured by broad-sweeping, top-down allocations (Celent, 2006). Whatever the reason, for years it has gone like this: Put everyone into a group; determine the cost of that group; then, spread cost and revenue to that group. The problem we find is that this method drops the customer detail for good, lost in the abyss of our own segmentation. Using solutions like SAS® Profitability Management however, we can now very quickly use detailed customer transactional data to drive behavioral cost that comes from an ABM system like SAS® Activity-Based Management, as well as direct revenue and cost from other systems (such as FTP systems) into each customer intersection using easy-to-define business rules. This provides a wealth of detailed customer profitability information at the place where it's incurred—at the transaction level—with information as to what behaviors produced it. It can then be aggregated to any level, for any dimensional intersection, regardless of our previous segmentation. This of course becomes one of the other critical inputs to your optimization routine. Though we don't discuss this at length in this paper, this data also serves as critical input to predictions about future or lifetime value (also known as customer lifetime value, or CLV for short). CLV can also be used as an alternative or comparative input to a PRO system since it can serve as a useful description of future behavior. At the risk of sounding too obvious: If your profitability numbers are based on broad estimations rather than driven by actual customer transactions, then the value of your optimization outputs is diminished. As they say, "garbage in, garbage out."

CRITICAL IMPLICATIONS FOR RISK

Continuing to use our credit card example, the optimization process can become quite complex when incorporating credit worthiness. Before we can make the link into the optimization, some background related to credit risk management might be helpful. In the simplest terms, credit risk is the potential loss due to the counterparty's inability to meet its obligations. The overall result can be derived by some form of the following expected loss equation:

*EL=PD*LGD*EAD (Where; EL = Expected Loss, PD =Probability of Default, LGD= Loss Given Default, EAD = Exposure at Default).*

Within the retail banking industry, credit scoring is used to aggregate credit card holders into groups based on certain attributes. This can be anything from a FICO score (which denotes an expected PD) to a direct PD calculation. From there, most offers are made group-by-group (such as credit cards). By simulating the PDs over time, we can now also calculate the unexpected loss (UL) which is the amount greater than expected loss due to uncertainties in the estimates. This UL can impact the amount of minimum regulatory capital required for the portfolio of credit card customers, and ultimately must be rolled into the net profitability figure used in your pricing decisions.

OPTIMIZATION

"Consider everything. Keep the good. Avoid evil whenever you notice it." (1 Thess. 5:21-22)

Optimization is defined by Mikhail Atallah as a "…computational problem in which the object is to find the best of all possible solutions. More formally, find a solution in the feasible region which has the minimum (or maximum) value of the objective function." (1998)

In this section, we will elaborate more on the notion of building a framework of optimization routines that aid in pinpointing the best price that will maximize profit. For our example, we have taken some liberties at merging our bid history data with notional revenue and profitability numbers. As mentioned above, these data would come from profitability models that take into account a complete view of our expected customer relationship. Our dataset looks like this (Table 1):

Obs	PRICE	PRED	LOWER	UPPER	TOT_REVENUE	TOTAL_PROFIT
1	13.8	0.99924	0.95112	0.99999	46884.33	-3077.66
2	13.9	0.99873	0.94044	0.99997	47200.00	-2736.52
3	14.0	0.99788	0.92754	0.99994	47499.12	-2394.91
4	14.1	0.99646	0.91201	0.99987	47770.48	-2052.72
5	14.2	0.99411	0.89339	0.99971	47995.43	-1709.86
6	14.3	0.99019	0.87116	0.99934	48143.09	-1366.46
7	14.4	0.98372	0.84471	0.99851	48162.85	-1023.07
8	14.6	0.95584	0.77635	0.99264	47447.76	-344.10
9	14.7	0.92834	0.73256	0.98394	46398.24	-18.57
10	14.8	0.88576	0.68037	0.96580	44571.26	283.44
11	14.9	0.82271	0.61696	0.93041	41678.45	542.99
12	15.0	0.73527	0.53728	0.86917	37498.64	735.27
13	15.1	0.62439	0.43482	0.78222	32056.09	836.68
14	15.2	0.49873	0.31182	0.68599	25774.35	837.87
15	15.4	0.26276	0.10457	0.52102	13758.07	620.11
16	15.5	0.17581	0.05197	0.45359	9265.37	474.70
17	15.6	0.11322	0.02453	0.39328	6005.15	344.19
18	15.7	0.07099	0.01126	0.33905	3789.49	239.95

Table 1

Figure 2

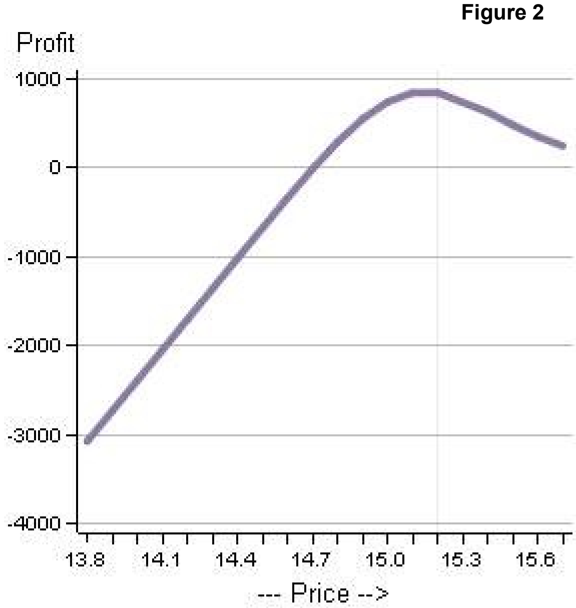

The price optimization dynamic is best understood by following the incremental revenue and profitability plots that correspond to our acceptance probability.

Starting at the lowest price on the X-axis, we can see that for this customer intersection, profitability is actually negative (Figure 2). Meaning, we have the highest acceptance probability, but our incremental cost exceeds the incremental revenue, and if we offer this price, we will lose money. As we begin to raise the price however, we see demand drop, which reduces our incremental cost, and the increase in revenue (by raising the price) brings our profitability up above zero. This would be your "bottom line" price with maximum demand, as indicated by the vertical line in Figure 1, at 14.8. Eventually, incremental profitability peaks at the price point where a balance is made between demand, revenue, and cost. In this case, the price that will optimize profitability without constraints is 15.2, as shown in Figure 2.

It's also important to note that revenue also has its own peak, often before profitability (Figure 3). Meaning, if someone has a marketing campaign that involves a price that is optimized for revenue, it's rarely going to generate the maximum profit. Likewise, any campaign optimized for demand must be tempered by the significant trade-off in profitability. In fact, understanding optimal demand, revenue, and profitability together would allow you to determine the "cost" of the campaign by subtracting the potential optimal profit from the potential profit at the price specified by the campaign.

Of course, we cannot always charge whatever price we want, and we therefore must be prepared to impose constraints on any of the dimensions or measures in play. For example, we may have other strategic marketing objectives in place, or certain fair lending rules may allow only certain kinds of differentiation, or may require caps on rates, etc. In our example, we have included arbitrary thresholds that were used to demonstrate potential key business drivers. Currently, there are no tradeoffs between the constraints nor are our data particularly sophisticated; therefore, the process for determining the optimal price does not require SAS/OR® coding logic. Rather, a set of basic filtering and ranking procedures can be utilized. Given the optimization problem within the context of this particular banking scenario, the objective function can be defined as the process for maximizing profits ($\pi_{a,r}$) subject to total revenue (r) greater than \$35,000 and the probability of acceptance (a) greater than 80%. In this case, the following code would be sufficient:

```
%LET LOGITOUT=DEMAND_AND_PROFIT;
%LET OBJFUNCTION=TOTAL_PROFIT;
%LET CONST1=(TOT_REVENUE > 35000);
%LET CONST2=(PRED > 0.8);

PROC SQL;
CREATE TABLE CONST AS
       SELECT *,
       (MAX(&OBJFUNCTION )) as MAX
       FROM &LOGITOUT
       WHERE &CONST1 AND &CONST2;
CREATE TABLE MAX AS
       SELECT *
       FROM CONST
       WHERE &OBJFUNCTION = MAX;
RUN;
QUIT;
```

Figure 3

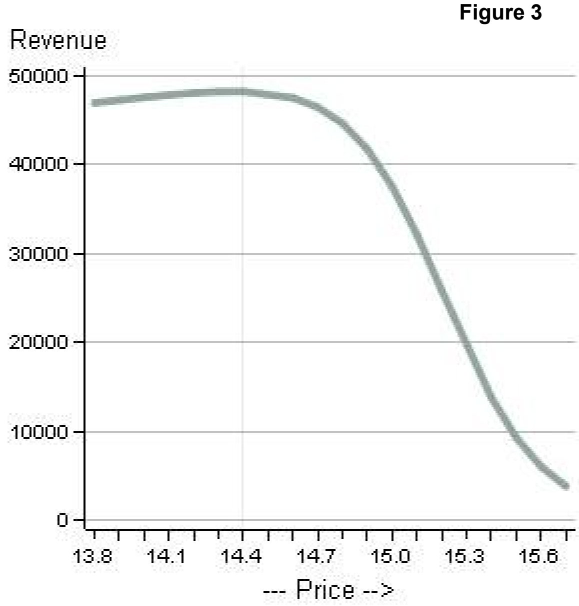

The result of which gives us the following result (Table 2):

Obs	PRICE	PRED	LOWER	UPPER	TOT_REVENUE	TOTAL_PROFIT	MAX
1	14.9	0.82271	0.61696	0.93041	41678.45	542.988	542.988

Table 2

The objective function in our optimization problem above designates maximum profit as our ultimate goal within some basic constraints. In reality, integrating risk analytics and simulation introduces trade-offs between constraints and greatly increases the complexity of the problem. The application of market and credit risk management will be a necessary component as retail banks move into pricing optimization because of the sheer number of financial offerings that have associated levels of risk, as well as the number of behavioral factors that affect each model.

For example, one may think that higher levels of interest rates on credit cards would yield higher levels of profit, but each increase in a basis point could also increase the probability of default. Now, a bank must weigh how much risk it is willing to take on to achieve a given level of profit. Likewise, those attracted to lower rates may prepay balances prematurely, reducing the overall expected profitability. This scenario—the more practical integration of risk factors, profitability, and the probability of acceptance—becomes now a prime candidate for more sophisticated analytic techniques used by the likes of SAS/OR, which we will leave for future discussion.

DIFFERENTIATION

To be clear, differentiation is not discrimination. Differentiation is the use of legitimate methods for charging different prices for the same product to different customers. For example, the airline industry effectively charges higher prices to business travelers than to leisure passengers for the same flights. They do this by having lower prices for early booking, and higher prices for late booking, with the understanding that people booking late are typically business travelers and are willing to pay more for their ticket. Most differentiation strategies are based on customer behavior, that is, letting customers demonstrate their own sensitivity to price based on how they buy or use your products. By calculating the optimal price at each intersection (presumably by running optimization routines at each intersection of customer segment, by channel, by product), small increments of profit are obtained that otherwise wouldn't have been. Here is a short walkthrough of a few practical strategies.

CHANNEL USE

Customers have come to expect that different channels are more expensive than others. The most obvious example is the Internet shopper. Customers shopping for a money market online are going to be more sensitive to price than a customer who walks into the retail branch of their primary bank. Customers browsing credit card rates from a mobile phone may be younger, and may be more fascinated by the novelty of applying for credit from a phone than they are concerned about interest rates.

DYNAMIC RELATIONSHIP PRICING

Dynamic relation-pricing is essentially a way to allow the customers' behavior to dictate how much they pay for a particular product or combination of products. Essentially, price differentiation is achieved by specifying rules by which a customer must interact with, or use the product in order to earn a particular price. Free ATM use in exchange for exclusively using Internet banking, or free trades with high deposit balances are all ways to charge a different price to different customers for the same product, based on their willingness to exhibit these behaviors.

To put relationship pricing in a price optimization context, imagine a bank considering a campaign to offer a 5.2% yield on a checking account; only, the customer must meet stringent behavior criteria, such as converting to online-only statements, use their debit card at least 10 times a month, log into Internet banking once a month, and have a direct deposit from their employer (yes, this product does currently exist). What we've done is to begin the process of creating a product theoretically available to everyone, but only certain customers—customers whose behavior and profitability we can predict—actually sign up for it. By analyzing historical bid win/loss history for this customer intersection, we could determine, in advance of setting the price, that all customers who behave in this way are relatively price-sensitive, and have historically moved their deposit balances to online-only institutions. By using price optimization in conjunction with relationship pricing (differentiation), we may find that the price that will produce the optimal profitability for this specific intersection is 4.4%. If we go ahead with the original plan of 5.2%, we will be able to quantify the potential loss in profit from that decision.

TARGETED MARKETING

Similar to retailers using coupons, we may make a particular product available to all customers at a given price, but we may target marketing efforts only through specific channels or to specific customers in order to attract the intersection to which we are selling. This adds an additional optimization layer, such as determining which specific range of prices (the product of your PRO analytics) should be offered to which customers, through which channels, while considering different types of marketing constraints (such as the offer medium, campaign-cost constraints, contact constraints, etc.). This type of problem is frequently solved by solutions like SAS® Marketing Optimization. Put simply, PRO techniques can tell you ***what*** price to offer to whom in order to optimize some objective function, whereas marketing optimization solutions tell you ***which*** price to offer, as well as ***when*** and ***how*** it should be offered.

CONCLUSION

Ultimately, developing a price and revenue optimization approach in banking is not trivial, and there are, of course, a multitude of aspects we didn't go into. However, we think we have provided a nice collage of concepts and considerations on this topic to get you started.

That being said, we hope that you take away the following concepts:

- that financial institutions are indeed applying these techniques and developing systems as the data to support them comes into their own
- you must have sound profitability information based on accurate, behavior-driven cost and revenue that incorporates risk analytics
- bid win/loss history is a critical component that can no longer be ignored
- and that there are several mature SAS products out there that together, can round out even the most complex pricing problems.

REFERENCES

Atallah, Mikhail, J. 1998. *Algorithms and Theory of Computation Handbook*. Boca Raton, FL. CRC Press, LLC.

Datamonitor. 2006. "Activity-Based Costing and Performance Management." *ComputerWire Business Intelligence Marketwatch Series*. http://www.sas.com/reprints/datamonitor_0406.pdf.

Feig, Nancy. 2007. *Compressed Margins, Advanced Technology, Mature Data Stores Lead Banks to Price Optimization*. Bank Systems & Technology. http://www.banktech.com/showArticle.jhtml?articleID=199204030

Grealish, Alenka. 2006. Celent Report. "Customer Analytics Triathlon: From Data to Knowledge to Wisdom." http://www.celent.com/PressReleases/20060613/CustomerAnalytics.htm.

Phillips, Robert. 2005. *Pricing and Revenue Optimization*. Standford, CA: Stanford University Press.

RECOMMENDED READING

SAS Institute. 2007. SAS Institute white paper. *"Optimization with SAS/OR®."*
http://www.sas.com/apps/whitepaper/index.jsp?cid=3584

SAS Institute. 2003. SAS Institute white paper. *"Are All Your Customers Profitable to You?"*
http://www.sas.com/apps/whitepaper/index.jsp?cid=3450

CONTACT INFORMATION

Your comments and questions are valued and encouraged. Contact the authors at:

Keith Renison
SAS Institute Inc.
Portland Oregon, USA
+1 (503) 292-3673
keith.renison@sas.com

Lorry Hardt Jr.
SAS Institute Inc.
Fort Worth, Texas, USA
+1 (817) 349-6037
lorry.hardt@sas.com