



A Modern Approach to RFM Segmentation by Roy Wollen

Executive Summary

Despite the current love affair with predictive models, direct marketing's three-variable formula, Recency-Frequency-Monetary Value (RFM), still has a place in modern database marketing. RFM is not a replacement for inferential statistics. But in the real world, RFM can still be useful when models are not practical. RFM also provides a management summary of customer behavior based on purchases and plays a role in policing the black-box results of predictive models to ensure quality *before* a campaign is implemented.

Segmentation of customers

What comes to mind when you read these descriptions of customer segments?

- Advocates
- Repeat buyers
- Gift givers
- Too good to be true
 - Habitual returners
 - Fraud
- Trial buyers
- Once was enough
- Dormant
- Defected
 - About to Defect
 - Revolving Door
- Cry in your pillow (“please come back”)

Each of these customer segments are rooted in purchase behavior. Purchase behavior is the best predictor of repeat purchasing and loyalty. While it is measured in different ways, depending on industry and customer lifecycle, all database marketers covet its empirical facts on how often a buyer renews their subscription/membership, visits your site, shops at your store. Further, purchase behavior is about how much they spend, the products/services they buy, and in what combination or sequence. Purchase behavior codifies both the tenure- as well as recency your relationship with your customers.

The good news for today's marketers is that the most important purchase behavior is already on your customer database. What can you do with these data? You can:

- Identify groups of customers
- Target them for campaigns
- Promote repeat purchase and loyalty
- Defend against attrition/defection
- Acquire customers who resemble the best ones



Segmentation gives direct marketers a quantifiable way to distinguish between the best and worst customers on file. Purchase behavior is the most powerful way to segment your customers by historic value.

You can also gain insight with purchase behavior from outside sources of information such as co-op databases, which are a collection of hundreds of direct marketing lists pooled together; individual response lists and any merge/purge processing itself (with its resulting intra- and inter- matches). These data provide contextual dimensions, and will help you realize there is more going on in your customers' lives than just their relationships with you.

Other kinds of segmentation bring you closer to why your customer buys from you, particularly from primary research which adds their attitudes and experiences. This information, along with demographics, explains their motivations and brings customer segments to life. This area of research is the most interesting and strategic, and has impact on benefits customers seek, media consumption habits and advertising strategy. It is critical input to decisions around what to say and how to say it.

Recency – Frequency – Monetary Value (RFM)

Recency Frequency Monetary Value (RFM) is a quick, descriptive way to segment a marketing database on purchasing behavior that direct marketers have used with success since the 1930s.

What is RFM? How do you use it? Is it superior to inferential statistics such as predictive models and decision trees? Generally, there is no contest against statistical modeling; RFM will *lose* just about every time. While this whitepaper may have just come to a screeching halt, the balance will argue that RFM still has an important place in today's modern direct marketing as a complement to predictive models, particularly for non-statisticians (Figure 1).

Figure 1: RFM advantages over Predictive Models

RFM	Predictive Models
Easy for managers to understand	Black Box
Can build it yourself	Need to hire a or trained data miner
Can build it yourself, now	Requires a build process, with analysis and validation datasets
Portable across industries	May be applicable to only one company
Somewhat effective at mitigating the confounding effect of seasonality	Would need a model for each season; ideally one model for each campaign
RFM definition is stable and does not need to be rebuilt or redefined	Typically would need to be rebuilt every 2 years when predictive power decays, or in reaction to a competitor or marketplace shift.
Applies to all customers. Supports sortation of all customers in the database (by RFM quartiles, quintiles, deciles, duo-deciles, centiles)	Doesn't always apply to all customers (why score customers you know you won't promote?)
Can use RFM across the organization for reactivation, cross-sell	Additional model may be required for reactivation and cross-sell programs

To be fair, predictive models have critical advantages over RFM. RFM by definition only utilizes 3 predictor variables, whereas predictive models can interrogate hundreds or even thousands. Models employ as many independent variables as necessary to maximize prediction. While many are collinear, the more independent variables, the more power in predicting future purchase behavior.

The Definitive article on superiority of predictive models vs. RFM can be found within a library of articles on direct marketing published by Wheaton Group.
http://www.wheatongroupllc.com/library/01_01_01.asp

If you're forced to choose, choose 'both'

I choose to use both predictive models and RFM side by side. My rationale is that predictive models do a superior job of predicting sales and my first priority is to make money. Predictive Models have a better return on investment than segmentations based on RFM.

However, RFM is applicable when modeling isn't practical. For example, perhaps your database is too small to warrant the investment in building the model. The fees for building predictive models can run in the tens of thousands, while RFM is essentially a do-it-yourself proposition; and, as a result, much cheaper. Finally, building a model takes time, especially since you need to go back in time, build a regression equation on a group of customers, validate the result on an equivalent group, implement the model by scoring today's database, and then pull the trigger. RFM values, by contrast, can be applied to your database by the time you get to the end of this paper.

My favorite role for RFM is as management tool. By predictor variables, models gain predictive power, but lose the ability to explain the reasons why it works. This is a black box for management, as well as a missed opportunity to understand the key business drivers of the dependent variable (response, sales per campaign, loyalty). It's true there have been admirable steps forward in pairing up Regression models for predictive power with Principal Component Analysis and tree-analyses to explain its building blocks. However, these are additional investments in themselves which require time and money. My suggestion is to use RFM for understanding, in concert with predictive models for campaign execution.

Validate a predictive model with RFM before you contact customers

How can you validate that a model is correct before you pull the trigger? What are some basic diagnostics you can run before you put the offers into the marketplace?

When brings you very attractive F-Statistics and R^2 compelling you to roll the dice on your next campaign, you should verify their work as your first step. First, you'll be reviewing the decisions they made with regards to sample sizes, statistical and sampling techniques, treatment of outliers, treatment of incomplete information (blanks, nulls, garbage), transformation of variables and so on. But you should also ask for a crosstab comparing their "score" (outcome of the model) versus RFM.

This crosstab should contain not only population counts, but ratios such as LTD Dollars/Buyer, Average Order Size, Average days since last purchase, LTD Orders/Buyer. You'll be looking for a correlation between the best scores and the best ratios. If there are attractive ratios in the basement of your predictive model, you may be missing out on opportunities with good customers.

A Modern Take on Recency



It is part of Green Bay Packer folklore that coach Vince Lombardi starting a training camp, attended by professional athletes, snorting “This, is a football!” The implication was start from scratch and re-learn the game (his way).

I’ve always related to that sentiment in a more positive light and interpreted it to mean learn everything you can about a particular topic, whether that’s a pulling guard sweep to the weak side (football jargon) or customer segmentation.

As business people who leverage databases to be better direct marketers, what can Recency, the R of the RFM model, teach us about our customers and marketing programs?

First, a definition. Recency (R) is defined as the time since last purchase, or meaningful transaction, that your customer makes. The more recent the last action, the higher the likelihood your customer will respond the next e-mail, phone solicitation, direct mail campaign, etc. It is operationalized as the number of months since the last purchase. But the unit of measure can easily be changed to weeks or days for online businesses. You should always use the update date or the high date¹ on the database, not “today’s” date.

What is the Date of Last Purchase that makes sense for your business?

- Most Recent purchase at the store
- Most Recent website purchase
- Most Recent catalog purchase (your catalog)
- Most Recent purchase in division X of your corporation
- Most Recent purchase across all divisions (aka “Corporate Recency”)
- Most Recent catalog purchase (across all catalogs in a co-op databases)
 - Z24, Prefer Network, Abacus, I-Behavior, Next Action, b2bBase, MeritBase
- b2b Recency has its own special application of recency
 - When the individual made their last purchase
 - When the company made its last purchase (Most recent date across all individuals)
- Last point earning or award activity on the web
- Last installment date on a continuity club

Recency has important business applications beyond segmentation. It is a key business dimension and can be triangulated (same customer - two recency dates) to see customers with a deeper understanding (See Figure 2 on the next page).

¹ In database marketing, a database is sometimes called “Names through x” where x is the most recent purchase date observed or the update date.

Figure 2: Recency can be applied beyond RFM Segmentation when it's "triangulated"


Recency Date	Triangulated with	Application
Last Purchase Date	Frequency – Monetary Value	RFM Segmentation
Last Web Date	Last Store Visit	Channel Preference Each channel's recency, in what combination, will describe a customer segment
Last Web Date	Corporate Recency	Reactivation If identical, web was most recent. If not, customer bought in another channel and (depending on size of difference) candidate for reactivation
Division 1 Recency	Division 2 Recency	Cross-Sell experiments
Individual Recency (b2c)	Household Recency	Cost Reduction Mail one per HH
Individual Recency (b2b)	Site Recency or Parent Company Recency	<i>Topic for another whitepaper</i>
Your company's recency	Co-op Recency	Reactivation
Non Buyer (no purchase recency)	Co-op Recency	Acquisition Better Recency on Co-op indicates better target
Non Buyer (no purchase recency)	List Recency ("selects")	Acquisition Better selects indicates better target (e.g., "hotline" names bought within last 90 days)
Recency	Products Purchased	Product Recommendations "People who bought x also bought y" recommender systems
Last Ship Date for Series	# of installments	Tenure promotion
Last Order Date	# of remaining items in series	Expected date when supply runs out
Last Subscription Date	# of issues in subscription	Renewal Campaigns
Last Response Date	NCOA move date	Winback campaign
Last Response Date	Job Change date	Bring us with you campaign
Last Purchase Date	Item Shipment Date(s)	Fulfillment issue?
Last Purchase Date	Pattern of Purchase Dates	Seasonality, Velocity, Segmentation (gift giver?)
Last Purchase Date	Time going by without repeat purchase	Retention campaign
Last Purchase Date	First Purchase Date	Tenure and Loyalty analysis

Show me the Data

A report card on a season's worth of campaigns can illustrate the power of Recency on its ability to segment customers on response rate and sales. This simulation is a powerful way to understand the range of quality of your customers (See Figure 3). I'm presenting this as a management summary, not as a suggestion on how to execute your campaigns.

Figure 3: Response Summary by Recency


Response Summary by Recency						
Last Purchase Recency	Customers	6-Month Season			% Resp.	Sales/Piece
		Campaigns	Visits	Revenue		
0 - 3 Months	170,862	1,215,314	49,401	\$21,797,541	4.1%	\$17.94
4 - 6 Months	128,238	1,034,895	21,715	\$8,552,928	2.1%	\$8.26
7 - 12 Months	202,443	1,436,914	21,004	\$8,102,213	1.5%	\$5.64
13 - 18 Months	178,411	912,021	7,217	\$2,636,827	0.8%	\$2.89
19 - 24 Months	154,214	592,846	3,812	\$1,259,712	0.6%	\$2.12
25 - 36 Months	294,001	641,028	3,105	\$1,071,216	0.5%	\$1.67
37 - 48 Months	141,888	226,515	857	\$239,701	0.4%	\$1.06
49 - 60 Months	46,071	72,184	215	\$74,715	0.3%	\$1.04
Total	1,316,128	6,131,717	107,327	\$43,734,853	1.8%	\$7.13



It is clear to see that “sales per piece” and its first cousin “response rate” march to the beat set by Last Purchase Date, the row in this report. As managers, we’re looking for ways to distinguish between good and bad customers. Clearly the recent customers (0-3 months since their last purchase) are superior to the dormant customers (no purchase in 2 years).


This can be charted on dimensions including revenue per buyer as well as contribution margin per buyer (especially as it relates to a threshold for breaking even). See Figure 4.

Figure 4: The power of Recency illustrated



Customer segments with more recent purchases garner more attractive results (sales per piece) from a 6-month campaign.

Charting **sales per piece** illustrates the range from a high of \$18 (for 0-3M customers) to a low of \$1 (for 19M+).



Margin per buyer can also be charted to illustrate which customer segments are above the breakeven threshold, where the campaigns cover costs.

Those segments that are below breakeven should be suppressed; the investment re-allocated to better segments or pocketed as earnings.

A Modern Take on Frequency

Frequency (F) defined as how often a buyer has a dollar earning transaction. It is typically a “Life to Date” field, and thus would be the accumulation of all transactions from an original date to the update date or. This value includes the original transaction (e.g., enrollment event, first purchase, etc).


The modern take on Frequency is that it is more interesting to understand the difference between 1x and 2x, versus 10x and 11x. That may seem odd, especially since the 10 time buyers are your loyal advocates, and they are disproportionately important to the health of your business. That’s true, however, when planning campaigns the trial buyers need all the help they can get from your customer database and beyond. (Using RFM, I’m defining 0-3M 1x as a trial customer, whereas 19M+ 1x is a goner for most businesses).

What else can we use for 1x buyers? I’ve found several things important clues as to which one time trial buyers will blossom into advocates and which won’t. Here are some idea starters:

- Depth: # items
- Breadth: # of departments (distinct “aisles” in the store)
- Price points
- Demographics, Geography (Zipcode, clusters) is more important for 1x buyers (and critical for 0x buyers)
- Contextual information
 - From co-op databases
 - 1x for you; 10x for your competitors
 - 1x for you; new to catalog shopping
 - From the merge/purge
 - Match to rentals (1x with lots of matches is far more promising than 1x with no matches)
 - “Multi-Buyers”: 1x multi buyer is a customer who did not match another list in a merge/purge
 - 2x multi matched 1 other list
 - 3x multi matched 2 other lists
 - From the a b2b site
 - 1x buyer in a loyal, high-revenue site is more promising (and may already be a specifier-decision maker) than a 1x buyer in a new 1 buyer site.
- Source of the customer
- Self-selection (did they come to you, or did you bribe their first purchase with a \$-off coupon?)


Figure 6: Response Summary by Recency and Frequency

Response Summary by Recency - Frequency							
Recency	Frequency	Customers	6-Month Season			% Resp.	Sales/Piece
			Campaigns	Visits	Revenue		
0 - 3 Months	3x+	90,802	713,249	37,508	\$17,147,029	5.3%	\$24.04
0 - 3 Months	2x	29,595	188,862	5,360	\$2,115,413	2.8%	\$11.20
0 - 3 Months	1x	50,465	313,203	6,533	\$2,535,099	2.1%	\$8.09
Subtotal 0 - 3 Months		170,862	1,215,314	49,401	\$21,797,541	4.1%	\$17.94
4 - 6 Months	3x+	58,251	513,520	14,231	\$5,835,189	2.8%	\$11.36
4 - 6 Months	2x	23,920	188,644	3,139	\$1,137,168	1.7%	\$6.03
4 - 6 Months	1x	46,067	332,731	4,345	\$1,580,571	1.3%	\$4.75
Subtotal 4 - 6 Months		128,238	1,034,895	21,715	\$8,552,928	2.1%	\$8.26
7 - 12 Months	3x+	81,972	666,277	12,293	\$4,944,773	1.8%	\$7.42
7 - 12 Months	2x	40,365	280,555	3,764	\$1,371,733	1.3%	\$4.89
7 - 12 Months	1x	80,106	490,082	4,947	\$1,785,701	1.0%	\$3.64
Subtotal 7 - 12 Months		202,443	1,436,914	21,004	\$8,102,208	1.5%	\$5.64
13 - 18 Months	3x+	60,541	375,304	3,695	\$1,406,123	1.0%	\$3.75
13 - 18 Months	2x	37,628	192,206	1,465	\$531,979	0.8%	\$2.77
13 - 18 Months	1x	80,242	344,511	2,057	\$698,731	0.6%	\$2.03
Subtotal 13 - 18 Months		178,411	912,021	7,217	\$2,636,832	0.8%	\$2.89
19+ Months	3x+	129,024	398,968	2,687	\$949,888	0.7%	\$2.38
19+ Months	2x	127,003	313,473	1,816	\$598,560	0.6%	\$1.91
19+ Months	1x	380,147	820,132	3,487	\$1,096,891	0.4%	\$1.34
Subtotal 19+ Months		636,174	1,532,573	7,989	\$2,645,339	0.5%	\$1.73
Total		1,316,128	6,131,717	107,327	\$43,734,864	1.8%	\$1.78




Breaking down Recency ranges by Frequency ranges will add to your understanding of your customers. Not all 1x buyers are bad (particularly the more recent ones). Not all 3x+ buyers are loyal (particularly the less recent ones).

Figure 7: The power of Recency and Frequency illustrated



Again, segments with better RF scores garner more attractive results (sales per piece) from a 6-month campaign.

Sales per piece now range from a high of \$24 (for 0-3M 3x+) to a low of \$1.34 (for 19M+ 1x).



Charting **margin** per buyer gives marketing managers more information to make campaign decisions.

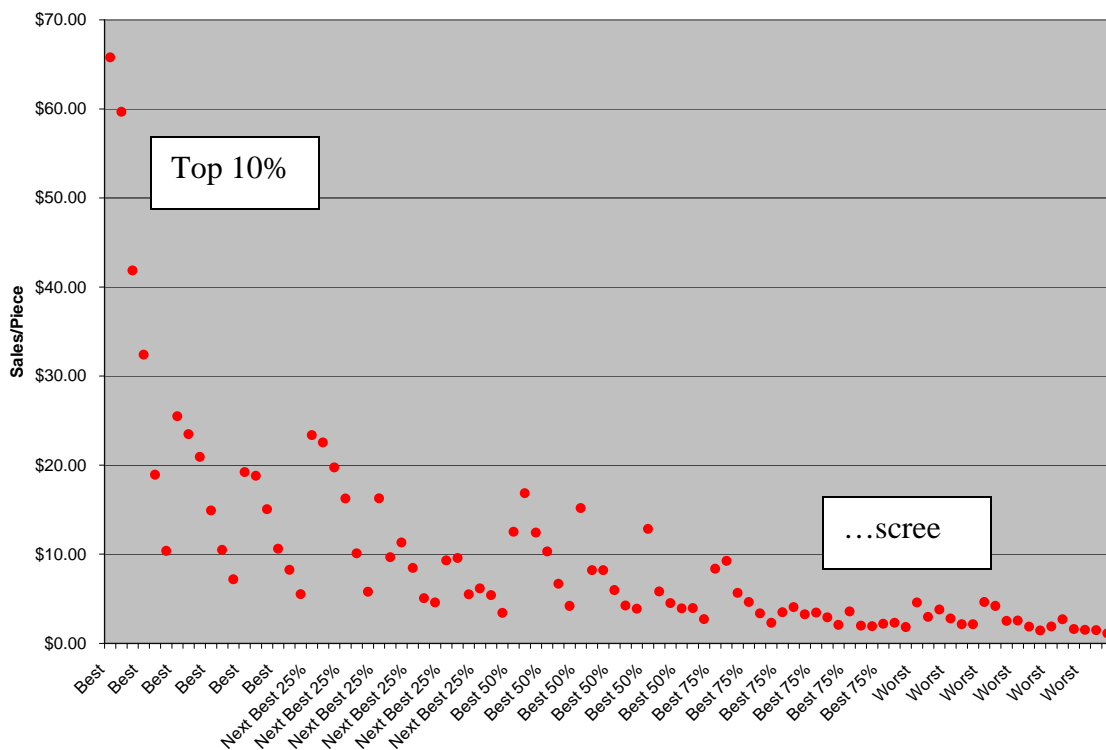
Viewing the chart with RF scores on its x axis now illustrates that the 1x portion of the 13-18M segment is below breakeven, but repeat buyers in the same recency ranges are above.

A Modern Take on Monetary Value

Monetary Value (M) is the sum of all revenue earned. A judgment is used to decide between “Life to Date” dollars, “Average” dollars, or some dollar amount over time (“0-12 month” dollars).

Adding the third leg of the RFM triumvirate makes our charts resemble what’s called a scree-plot. Scree plots borrow their name from the junk that falls off very high mountains. The RFM chart now looks like giant peaks of high customer value compared with run off of lower value segments. See if you agree: *compared to your best customers, everybody else is just scree.*

Figure 8: The power of RFM illustrated



Some might call this the proverbial “80/20” rule. My colleague Chris Pickering calls it “Pickering’s Law of Disproportionate Value” which is only fair given that the definitive book on RFM is called: Libey and Pickering on RFM and Beyond, published by MeritDirect Press and can be found at <http://www.meritdirect.com/mdpress.htm>

RFM Model

The biggest complaint about RFM models is the use of cells, and the unwieldy number some RFM models produce. A large number of cells defeats the purpose of ease of use for management understanding.

Recommendation

- Create a continuous RFM score for each customer on your database when predictive models are not practical (How? See below).
- Use a cell-based approach to provide an easy-to-understand summary of purchase behavior, track customers over time plus demystify and police the black-box results of predictive models.

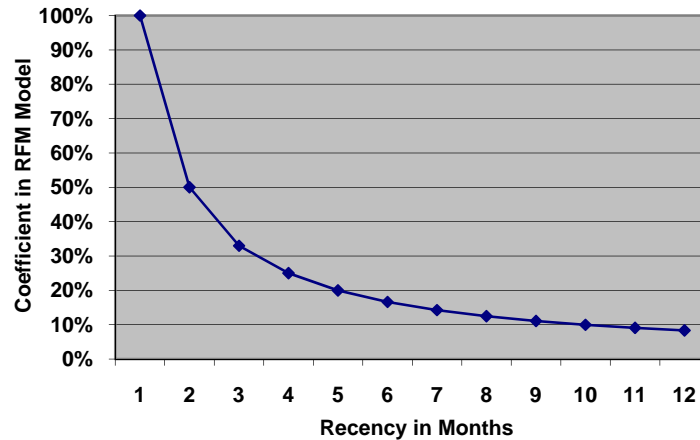
Rather than create a cell-based approach to RFM (where R would have month ranges such as 0-3, 4-6, 7-12..., F would have ranges like 1x, 2x, 3x+, and so on), the breakthrough thinking is to create a continuous RFM score for each customer on your database. This score is one value for each customer, and lends itself to sortation of the customers by value, and the decile summaries and gains charts that follow.

The formula is taken from the seminal article on the subject called "*A Direct Mail Customer Purchase Model*" by Connie L. Bauer, Journal of Direct Marketing, Summer, 1988.

The simple version of the model is **RFM Score** = $\frac{1}{R} * F * \sqrt{M}$

..where R is Recency of Last Transaction in Months, F is Number of Purchases, Transactions or meaningful buying events, and M is a sum of monetary value from the customer.

In this formula, Recency is a negatively accelerating curve that can be illustrated as follows:



Mathematically, Recency has the most powerful effect on the RFM score (e.g., Best recency is 1/1 and will not affect the RFM model, Next best recency is 1/2 and will reduce the RFM score by 50%, Next best is 1/3 and will reduce it by 66%, Inactive customers might have 1/13 and will reduce the score by more than 90%)

It's important to keep zeros out of this equation, so either remove non buyers or add the value 1 to each component as needed. A square root is used to scale down the effect of M, into the neighborhood of Recency and Frequency values, and ensure that it doesn't **bully** the RFM score result.

Steps in building an RFM model

1. Remove outliers, fraud, employees, suppressions, opt-outs
2. Decide how you will treat non-buyers
3. Build the model, QC a handful of records
4. Create an RFM score for each customer on the database
5. Rank the entire database on this RFM score
6. Divide into equal groups (deciles, duo-deciles, centiles)
7. Repeat after each update

Best Practices for using RFM

- Use common sense. With durable goods and large ticket items (cars, furniture) RFM may actually work in the reverse (where less recent is better than more recent). This may also be true with seasonal businesses. Also, be flexible in how you operationalize the values within the RFM model. For example, web marketing and e-commerce have its own RFM. Recency may be measured in days not months. Frequency may turn into how frequently a visitor returns to your site, what they purchase or view and where they click. M might be extended to include what's in the shopping cart (before abandons).
- Seasonality makes defining the R ranges a challenge. Does your 0-12M range include "Back to School", "Valentine's Day", "Holiday Shopping"?
- Augment RFM where you can, if it helps you understand your business dynamics. Bob Kestnbaum pioneered the concept of RFMP adding a product dimension to RFM. Kestnbaum called it "FRAC" where "Frequency" was first variable since he thought it was most predictive, "A" was "Average Dollars" and "C" was "Category of purchase." His rationale is sound: The best predictor of future Product A purchases is past Product A purchases. This addresses the challenge of seasonality. Your key business levers are also viable candidates to extend the simple RFM summary. For b2c direct marketing, this might be RFM*Income*. For b2b, this might be RFM*Industry*. Both might see benefits in an RFM*Channel* segmentation scheme.
- Extend RFM for acquisition marketing. RFM can also be put to work to acquire new customers. There is some current thinking on calculating RFM for each zipcode on file to target prospects based on observed RFM for customers that live in the neighborhood.

RFM is not a replacement for inferential statistics. It should add value in its unique ability to add management understanding, ensure data quality and substitute for predictive models when models are not practical. RFM has a role in modern direct marketing.

About Database Insight, Inc.

Database Insight, Inc. offers practical advice on how to improve your direct- and interactive marketing by deepening your understanding of your customers.

We offer consulting for direct marketing clients through exceptional people, forward thinking solutions and analytics. We bring to bear 20 years of experience understanding customers via database analysis.

For information on hard-copy or digital reprints, e-mail Roy@DatabaseInsight.com

Database Insight has an office within the historic Civic Opera Building at 20 N. Wacker Drive in downtown Chicago. <http://www.lyricopera.org/about/house.asp>

