Managing Churn By The Numbers

Dennis Dance (E-mail: dldance@meta3.net), D. L. Dance & Assoc.

Ray Phelps (E-mail: phelpsr@millisaps.edu), Millisaps College

Gayla Dance (E-mail: dancegf@millisaps.edu), Millisaps College

Abstract

Churn (a.k.a. customer turnover) is a significant problem in the telecommunications industry. Given the high cost of acquiring new customers, mining existing databases may provide information to promote customer retention. In everyday language this article presents a method whereby a manager using Microsoft Office™ is able identify those customers most likely to churn.

The article steps through an example using real world (slightly disguised) cellular phone data. Without using technical language, the reader learns how to generate individual unit profiles and baseline to determine if and when a particular customer is likely to churn.

The process presented is benchmarked using a blind test data sample. The process scored an impressive 80% accuracy in distinguishing active versus disconnected cases.

1. Introduction

In the telecommunications industry, nothing seems more certain than death, taxes and customer churn. After all, customers come and go. Churn—a.k.a. customer turnover—is a serious problem. For most upper level management, acquiring new customers seems more important than keeping existing onesⁱ. And this strategy worked well as the industry faced a staggering increase in the demand for telecommunication goods and services. Churn appears to have been accepted as a fact of life. It is just “one of the costs of doing business” and we just have to budget for it.

Without a doubt, costs are estimated and associated with the focus on “customer acquisition.” The fact remains that these new customers don’t come cheap—some estimate that it will cost about \$300 for every new mobile (e.g., PCS) client. And these high acquisition costs – combined with continuing support and service costs – typically do not allow carriers to make a profit from customers during the first year.ⁱⁱ

Customer churn is a serious issue. Indeed, “the U.S. churn rate will exceed the new subscriber add rate by 150%.”ⁱⁱⁱ One result of this trend will be for firms to expand their current emphasis on “customer acquisition” to include “customer retention.” The purpose of this article is to use a step-by-step approach so that you can develop a churn forecast. Our goal will be to use everyday language to present a method whereby any manager can use the tools in Microsoft Office™ to manage churn and increase customer retention.

Most likely Bill Gates feels that every PC either has, or should have, the ubiquitous Microsoft Office. So, let’s manage churn with it. While more sophisticated tools exist (including those used in our own firm), you can certainly begin managing churn with something as simple as MS Office. The results may not be as robust as those produced by more powerful processes. But you can start managing churn today. Later you may want to consider other processes using neural nets, advanced statistical packages, exotic algorithms, etc.

Certain caveats are in order. Apologies are made to statisticians everywhere. This paper uses a barebones pragmatic approach to managing churn. Liberties are taken that would make a grown statistician cry. The bottom line, however, is that this simple approach works. And that is all that matters.

2. What You're Going To Do

Suppose you have 50,000 active units and 2% monthly churn. What you're going to do first is identify the data you need for your database. It will contain data you think increases or decreases churn. Next you want to take a random sample from your customer care system. You want to extract both active and cancelled units. You then calculate the number of units that cancel service over time. You use these percentages to develop profiles for active units and disconnected units. Once an active unit baseline is developed, you then forecast what is likely to happen in the next 90 days. Finally, you begin the salvage process.

2.1. Step 1: Identify Your Data

There are at least two issues involved before creating any data set. The first is to determine what data to use. The second is to make certain that each field (or column) contains the expected data. For example, the zip code field contains nothing but zip codes.

Ideally, a quick meeting with representatives from your Information Technology (IT) department and your customer care system will confirm that you can be swamped with data.

Before creating your dataset you must attempt to identify what you believe are your churn drivers.

Let us define "churn drivers" as factors that "increase" or "decrease" customer retention. Initially, you are going to have to make a "best guess" about these variables. Your goal is to identify data indicating your behavior and that of your customer.

Your company's behavior might include such things as the following:

- Marketing regions you create (e.g., north, south, general office, etc.)
- Rate plans
- Promotions offered (e.g., free weekends, free minutes, etc.)
- Customer segments you identify (e.g., consumer, government, industrial, etc.)

Your customer's behavior would include such factors as:

- Payment history (e.g., 0 days late, 30 days late, etc.)
- Features selected (e.g., caller ID, call forwarding, etc.)
- Number of roaming minutes.

Don't agonize over these. You are not going to have to assign these to either your company or your customers. All you need to find are those data fields that are the result of someone's behavior.

Of course you can also use attitudinal, demographic and psychographic data. Some of these may be of value. We have found that effective segments may be created using items based on behavior.

Identify 15 – 30 data items. Many fewer than this will paint an incomplete picture. Many more than this will become a size and computational burden for you and MS Office.

This step is the hardest and arguably the most important. You are searching for data factors that are relatively independent of each other and represent someone's behavior. Discussion of churn driver selection is beyond the scope of this paper. For this paper, we cheated and used our proprietary system to select them. It identified 20: 14 are associated with increasing churn and six are associated with decreasing churn. Regardless, strive for data fields that have a relationship similar to that in Figure 1. That is, the lines diverge. They are not parallel to each other. They move relatively independently of each other.

Figure 1. Two Independent Data Fields

2.2. Step 2: Select A Random Sample

You can use MS Excel to select units at random. Suppose you want to select a random sample of 1,000 units from the 50,000 in your customer care system. Just use the MS Excel function “=RANDBETWEEN(1,50000)”. Because you’ll have to eliminate any duplicates, you may want to generate 1,100 random numbers to get 1,000 unique ones. Then extract the units so identified.

At the risk of your eyes glazing over, you can ask your IT Department to create the extract for you. Or you can put fear into their hearts and make their eyes bulge by telling them you want to use ODBC and Microsoft Query. Discussion of this is beyond the scope of this paper. But by using these feature you can extract your data directly from your customer care system without ever leaving MS Excel. MS Access has the same capability via ODBC. ODBC is Open DataBase Connectivity, which is a standard database access method developed by Microsoft Corporation. Are you grinning yet?

Whether you have 10,000 or 200,000,000 units to track, using a sample size of 1,000 units yields a sampling error around 3%. Using a 10,000-unit sample results in about a 1% sampling error. So stick with the 1,000 since a 10-fold increase in sample size improves your accuracy by only 2%^{iv}.

The sample data used for this paper is real world data that has been slightly modified to maintain data confidentiality.

Two fields that must be included in the extract are “Time” and “Status”. “Time” is the number of uninterrupted days a unit has been active. “Status” is an indicator which when “0” means the unit is active and “1” cancelled.

2.3. Step 3: Create Category Breakpoints

The next activity is to take your extracted data and place one data field by itself on an Excel tab. Assume you have 20 data fields; then you would use 20 different spreadsheet tabs. Sort each data field into ascending sequence by time and then into descending sequence by status.

You'll have two types of data: continuous data and categorical data. Continuous data includes things such as bill amounts, minutes used, etc. Categorical data represents specific choices such as a specific feature, specific rate plan, etc.

First examine your continuous data. Selecting one data field at a time, divide the sorted data into say 10 groups of roughly 100 units each (i.e., $1,000 = 10 \times 100$). For this example, use "Average Account Balance". Each group has a "first value" and a "last value" as shown in Table 1.

Table 1. "Average Account Balance" As Categorical Data

Group	First Value	Last Value	Count
1	-\$89.10	\$18.53	71
2	\$19.95	\$25.46	126
3	\$25.78	\$51.56	101
4	\$52.00	\$120.34	102
5	\$124.05	\$209.40	98
6	\$209.52	\$237.47	69
7	\$239.40	\$254.40	133
8	\$254.49	\$358.80	100
9	\$360.55	\$596.76	100
10	\$598.63	\$13,305.78	100

If you total the "Count" column you'll see that Table 1 has exactly 1,000 units. This will always be the case, even if you have duplicate data values since you want to know how many units are in each group.

Now consider your categorical data. For this example, use "Billing Method". You offer customers three payment options:

- Billing Method 1 – Paper bill
- Billing Method 2 – Credit card
- Billing Method 3 – Bank draft.

In the extract's "Billing Method" column, you would have 1,000 units with either a "1", "2", or "3" indicating the customer's payment option. In situations such as this (i.e., categorical data), you will have to create three new columns, one for each billing method. Alternately, decide which one you believe is a churn driver and focus on it.

Table 2. "Billing Method 1—Paper Bill" As Categorical Data

Group	First Value	Last Value	Count
1	0	0	64
2	1	1	936

Your experience tells you that the paper bill is a problem, so you select it and ignore the other two. You now create your category groups as shown in Table 2. It shows the first value and the last values are the same within each group. All this means is that only "0" or "1" appears in the data field. Zero means "No" and one means "Yes – the customer receives a paper bill". In both cases please note that all 1,000 units from the sample extract are included. Thus, if there are 20 churn variables, then you'll have 20 spreadsheet tabs, each with the same 1,000 units.

2.4. Step 4: Calculate Group Probabilities

Now select one of the churn drivers (e.g., "Average Account Balance"). Recalling Table 1, ten category groups were created. Select each of the groups in turn and compute the percent of active units remaining over time as shown in Table 3.

The steps to do this are listed for the churn driver "Average Account Balance" and its category Group 1 as follows:

- If not done previously in Step 3, sort the churn driver tab into ascending sequence by “Time” (i.e., column 3 in Table 3) and then descending sequence by “Status” (i.e., column 2);
- Using Table 1, mark each unit with its group number (i.e., column 5);
- For all groups except Group 1, set column 6 to “0”;
- For all Group 1 units, indicate their status as “0 – Active” and “1 – Cancelled” as shown in column 6;
- Using Table 1, obtain the number of units in Group 1 (i.e., 71);
- Every time a Group 1 cancelled unit (column 2) is encountered in column 6, subtract one from the number of active units in the group (column 7);
- Finally, calculate the percent of active units remaining in Group 1 (e.g., $98.59\% = [71 - 1] / 71$) as shown in column 8.

The percentages calculated in column 8 are in fact empirical probabilities. That is, column 8 shows the estimated probability that a particular Group 1 unit is active at a particular point in time (i.e., column 3) given that the unit has a value (column 4) in the range specified by its group category (i.e., Table 1).

All these steps must be repeated for each category associated with each churn driver.

Table 3. First 20 “Average Account Balance” Group 1 Percentages

Account	Status	Time	Average Account Balance	Group	Group 1 Status	Number Active Remaining	Prob(Active)
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
		0				71	
10029103	1	4	\$3	1	1	70	98.59%
10027895	1	23	\$34	3	0	70	98.59%
10026983	1	25	\$81	4	0	70	98.59%
10029686	1	29	\$45	3	0	70	98.59%
10029738	1	44	\$23	2	0	70	98.59%
10026636	1	47	\$47	3	0	70	98.59%
10029070	1	55	-\$15	1	1	69	97.18%
45407	0	58	\$31	3	0	69	97.18%
45318	0	61	\$62	4	0	69	97.18%
45289	0	62	\$52	4	0	69	97.18%
45249	0	63	\$18	1	0	69	97.18%
44981	0	69	\$42	3	0	69	97.18%
44951	0	70	\$43	3	0	69	97.18%
44968	0	70	\$28	3	0	69	97.18%
10024764	1	71	-\$14	1	1	68	95.77%
10024631	1	72	\$0	1	1	67	94.37%
44858	0	72	\$59	4	0	67	94.37%
44775	0	75	\$71	4	0	67	94.37%
10015986	1	76	\$21	2	0	67	94.37%
300	0	76	\$256	8	0	67	94.37%

2.5. Step 5: Organize Churn Driver Probabilities

At this point in the process you have computed ten columns of probabilities for the churn driver “Average Account Balance”. The reason you have ten columns of probabilities is that the churn driver has ten category groups as shown in Table 1. Because you also followed this process for the churn driver “Billing Method 1 – Paper bill” you have two columns of probabilities for it as implied by Table 2.

Referring to Table 4, the demonstration example has 20 churn drivers with each one having between two and ten category groups for a total of 69 category groups. Table 4 also groups the churn drivers into classes such as “Rate Plan”.

Table 4. Number Of Churn Drivers And Group Categories

# Churn Drivers	# Category Groups	Classification	Type Data
1	2	Billing Method	Categorical
4	8	Credit Classification	Categorical
1	2	Customer Area	Categorical
3	24	Dollars	Continuous
4	8	Marketing Area	Categorical
1	10	Minutes	Continuous
3	6	Rate Plan	Categorical
1	5	Rate Plan Promotion	Categorical
2	4	Unit Feature	Categorical
20	69	Total	

Organizing the churn driver probabilities is nothing more than collecting the group probabilities calculated in Step 4 onto one spreadsheet tab. An example is Table 5. This organizing of the probabilities makes Step 6 easier.

In Table 5 notice that the first 20 Group 1 probabilities (column 6) are all 100%. Now look at the turquoise circle (column 5). This is the only Group 1 unit encountered. All the others are in Group 2. The reason these Group 1 units are all 100% is the blue circle in column 2. The turquoise circle is the first Group 1 unit encountered. But its status (column 2) is Active (i.e., “0”). Thus, none of the Group 1 units has cancelled service as of day 76 (column 3): they are all still active.

2.6. Step 6: Specify Individual Unit Profile Probabilities

Every unit record extracted has a field for every churn driver you identify. However, the contents of that field may be zero or empty. For example, Table 4 shows three rate plans are churn drivers. For the sake of discussion let’s say these are churn drivers:

- Rate Plan 2
- Rate Plan 3
- Rate Plan 4.

A particular customer’s unit can be in only one rate plan at any point in time. Consider the following two cases:

- The unit is in Rate Plan 1 (not a churn driver), then data field “Rate Plan 1” = 1 and
 - Data field “Rate Plan 2” is “0”
 - Data field “Rate Plan 3” is “0”
 - Data field “Rate Plan 4” is “0”.
- The unit is in Rate Plan 3 (a churn driver), then data field “Rate Plan 3” = 1 and
 - Data field “Rate Plan 1” is “0”
 - Data field “Rate Plan 2” is “0”
 - Data field “Rate Plan 4” is “0”.

Table 5. Churn Driver Probabilities For “Billing Method 1 – Paper Bill”

Account	Status	Time	Billing Method 1 - Paper Bill	Group	Prob(Active Group 1 Unit)	Prob(Active Group 2 Unit)
(1)	(2)	(3)	(4)	(5)	(6)	(7)
10029103	1	4	1	2	100.00%	99.89%
10027895	1	23	1	2	100.00%	99.79%
10026983	1	25	1	2	100.00%	99.68%
10029686	1	29	1	2	100.00%	99.57%
10029738	1	44	1	2	100.00%	99.47%
10026636	1	47	1	2	100.00%	99.36%
10029070	1	55	1	2	100.00%	99.25%
45407	0	58	1	2	100.00%	99.25%
45318	0	61	1	2	100.00%	99.25%
45289	0	62	1	2	100.00%	99.25%
45249	0	63	1	2	100.00%	99.25%
44981	0	69	0	1	100.00%	99.25%
44951	0	70	1	2	100.00%	99.25%
44968	0	70	1	2	100.00%	99.25%
10024764	1	71	1	2	100.00%	99.15%
10024631	1	72	1	2	100.00%	99.04%
44858	0	72	1	2	100.00%	99.04%
44775	0	75	1	2	100.00%	99.04%
10015986	1	76	1	2	100.00%	98.93%
300	0	76	1	2	100.00%	98.93%

These two cases plus the group number are the key to use when specifying individual unit profile probabilities.

You determine each unit's profile by multiplying all the churn driver probabilities together. The resulting answer is that unit's estimated probability of remaining active. Suppose you have only two churn drivers. Then you have two cases to consider when performing this multiplication:

- The categorical data “Billing Method 1 – Paper Bill”
- The continuous data “Average Account Balance”.

Consider a unit that has the following profile:

- “Average Account Balance” of \$18
- The customer receives a paper bill
- The customer has been active for 63 days.

Use Table 3 to obtain the probability for the “Average Account Balance”. Account 45249 (column 1) matches this profile. Specifically, Table 3 shows Time is 63 days (column 3) and “Average Account Balance” is \$18 (column 4). Also, notice the group number is “1” (column 5). Thus for this churn driver in category group 1 the estimated probability to use is 97.18% (column 8).

Use Table 5 to find the “Billing Method 1 – Paper Bill” probability. Account 45249 (column 1) is in Group 2 (column 5) with a probability of being active set at 99.25% (column 7).

Using these two churn drivers only, the estimated probability of Account 45249 being active on day 63 is 96.45% (i.e., 97.18% X 99.25%).

But what if Account 45249 uses a bank draw draft? In this event the data field “Billing Method 1 – Paper Bill” would be zero indicating this payment method is not used. When this occurs use the probability 100% because “Bank Draw Draft” is not a churn driver. That is, “Billing Method 1 – Paper Bill” has no effect. Since you are multiplying all the probabilities together, 1 X any number is just the number (e.g., 1 X 5 = 5). Your result is 97.18%.

2.7. Step 7: Establish Today’s Units’ Profiles

After performing all the multiplications to develop each unit’s profile at “Time”, you can graph the estimated probabilities versus time and get a chart similar to the one in Figure 2. The green dots show active units while the red dots show cancelled units.

Figure 2. Active Versus Cancelled Unit Probabilities

Long ago when you first started reading this article, the comment was made that you had 2% monthly churn or 24% annual churn. Another way to look at this is 76% of your units do not churn.

So what’s the blue line? The blue line is drawn so that it should show 76% of all dots above the line and 24% of them below the line.

2.8. Step 8: Compute Baseline Probabilities

A reasonable question is, “How do you know where to draw the blue line in Figure 2?” The answer to this is “your churn rate”. This blue line becomes a baseline against which you can measure unit profile probabilities. Those profiles (i.e., red dots) below the baseline probabilities are likely to disconnect. Those profiles (i.e., green dots) above the baseline are likely to remain active.

The first step in generating the blue line is to separate the active units and the cancelled units onto two separate spreadsheet tabs. One tab has the active units and the other the cancelled units. Chart the profile probabilities for both groups on the same graph. They should look like the green and red dots in Figure 2.

Next you’re going “to smooth” the active units’ probabilities by using the “=AVERAGE(range)” spreadsheet function. The range you want to use is the first 30 active probabilities’. Then copy this “=AVERAGE(range)” to include all active unit probabilities in groups of 30. You are creating a running average with each average having 30 units.

Using the same ranges, calculate the standard deviation for each average calculated. Use the spreadsheet function “=STDEV(range)” for this. You are creating a running standard deviation with each having 30 units.

Now use the spreadsheet function “=NORMSINV(0.24)” where “0.24” is 24% which is your annual churn rate. In this instance the result is “-0.706302217” (not shown). (There are better and more sophisticated ways to develop the value to use as the argument for “=NORMSINV(argument)”. But they are beyond the scope of this paper.)

The X-values in Figure 2 are obviously “Time”. To calculate the Y-values for the blue line, use the formula “=AVERAGE(range) + STDEV(range) * NORMSINV(0.24)”. This is your blue baseline Y-values.

Finally, compare individual profile probabilities with the baseline. All probabilities greater than the baseline have a “retention” profile. All less than or equal to the baseline have a “disconnect” probability.

So how accurate is this approach in identifying whether a unit is already cancelled or active? Comparing against the baseline in Figure 2 was well as our own proprietary service, this probability approach correctly identified 83% of the units. Our own proprietary approach correctly identified 98%.

Some people like to see lift curves to compare how well a particular system or methodology profiles units. A detailed discussion of these type curves is beyond the scope of this paper. Regardless, Figure 3 compares three process curves:

- Proprietary (turquoise)
- Probabilistic (dark blue)
- Random guess (black).

The horizontal or X-axis shows the percent of customers that you would have to contact in a salvage campaign. The vertical or Y-axis is the percent of customers that you actually want to contact (i.e., only those that are going to churn). In the example here, of the 1,000 customers 700 are active and 300 are cancelled. The perfect curve would reach 100% on the Y-axis and 30% on the X-axis at the same time. That is, you would contact only 30% of your customers: those going to churn. This would mean that all profiles were correctly assigned and no mistakes were made. Unfortunately, few things in life are perfect. So the closer to an upside down letter “L” (i.e., “T”) a curve is, the more powerful and accurate the process.

If you were to guess (e.g., flip a coin) whether each of the units is going to remain active or cancel, you should get about half of the unit profiles correct. Any line above this diagonal black line is an improvement over guessing. Any curve below the black line is worse than simply guessing. Since two process lines are above the diagonal line, they are clearly an improvement over guessing.

Bottom line this simple probabilistic approach works. It’s better than guessing and, in fact, works quite well with 83% accuracy for this example.

2.9. Step 9: Forecast Unit Profiles For Today + 90 Days

Repeat the process described in Step 7 using Time + 90 days. By comparing these new probabilities at Time + 90 against the baseline, you can forecast which units have a “disconnect” profile over the next 90 days. These are the units to target for salvage.

2.10. Step 10: Repeat Each Month

Repeat all these steps each month. In general your churn drivers will slowly change during the year. So you may want to reexamine your churn drivers each quarter.

Because customers add/remove features, talk more/less, etc. you will need to reevaluate the probabilities each month. Naturally, you’ll have to recompute your 90-day forecast each month.

But the hard part is already been done. You already have the formulas, spreadsheet, database queries, etc. All you need to do is download the new month’s data values from your customer care system.

3. A Note On Sampling Error

The term “sampling error” in the previous discussion on the random sample does not mean you or anyone else made a mistake. Rather, “sampling error”, has reference to selecting individual units that may not be similar to the average unit. For example, if the average billed amount for your units is around \$45 per month and one of the units in the random sample has a bill for \$2,000, then this unit is very different from the normal or average bill. The greater the number of these types of units the larger the sample error because they are further from the average.

Also, you may be asking yourself how using only a relatively few of your customers in the calculations can possibly help you to predict what all of your customers will do. While the estimated probability values you calculated using a portion of your customer database may not and probably will not be exactly the same as if you had used all the customers, the values you compute will be close enough to make important decisions. How close you want these values to be depends on how large a difference between them you can tolerate. There is a definite trade-off; the closer the values must be, the larger the sample (portion) we will have to use.

The size of the random sample you want to use can easily be calculated by using this Excel formula: $\text{=ROUND}(9604/(\text{acceptable difference})^2,0)$ where “acceptable difference” is the sampling error (e.g., 3%) you want to achieve. In the calculations for this paper, we chose a sample size of 1,000. This means that the average difference in the unit’s true cancellation probability and the estimated probability is around 3%. We believe that we can tolerate this difference for our purposes. Let’s say you feel that a 1.5% average difference in values is necessary. Substituting a 1.5 in for the acceptable difference in the formula above shows that you will need a sample size of 4,268 (i.e., $\text{=ROUND}(9604/(1.5)^2,0)$). Notice this is over four times the random sample size we used for 3%. A 1% average difference in values will require sampling 9,604 units. Notice that a small increase in accuracy requires a tremendous increase in effort and cost for a much larger sample size.

4. Limitations You Need To Know

Several limitations need to be discussed with this approach. Please note the following list is not comprehensive.

4.1. Limitation 1: Churn Driver Selection

Arguably the single most important step is selecting the churn drivers to use for developing profiles. There are some tools available in MS Office to aid in churn driver identification. However, they are not particularly good at identifying “good” churn drivers.

4.2. Limitation 2: Churn Driver Market Impact

Identifying churn driver market impact is practically impossible with Office like suites. For example our proprietary service shows that the churn driver “Billing Method 1 – Paper Bill” is associated with a 50% increase in churn. Similarly, one of the rate plans is associated with an 80% reduction in churn. This represents important market information not obtainable using this probabilistic process and MS Office. You can, however, generate comparison charts to make relative comparisons.

4.3. Limitation 3: What-If Scenario Analysis

Our proprietary service offers the ability to quantify What-If analysis at the unit level by:

- Adding or removing churn drivers
- Increasing or decreasing churn driver values
- Comparing these results against a unique value to determine whether a unit can be salvaged.

What-If analysis cannot be performed in this manner with the probabilistic approach. However, manipulating probabilities can perform some limited What-If analysis. This approach will affect all units with the factor being manipulated.

4.4. Limitation 4: Cancellation Dates

This probabilistic approach can only tell you the likelihood a unit will churn. There is no way to calculate the binary Yes / No state if it will churn.

Likewise, there is no way to calculate a date or date window during which the unit will cancel service. However, you can calculate the date or date window when the unit will develop a disconnect profile. You can use this time window to begin some salvage process.

4.5. Limitation 5: Manageable Size

Spreadsheets can hold only so much information before memory problems are encountered. Likewise, while 10,000 units might be processed in a spreadsheet many more would be inconvenient. So at some point you must consider using a database such as MS Access. All the sample size calculations can be performed in the spreadsheet. The churn driver group probabilities would have to be transferred to MS Access tables. The individual profile probabilities would have to be calculated with queries. These queries are not as simple as spreadsheet formulas and are comparatively slow in executing. But they do work. MS Access removes any size restrictions on the number of units.

5. Summary

Using a tool such as MS Office Professional allows you to develop individual unit probability profiles for forecasting churn. While some significant limitations are noted, the process is simple and easy to use. But most importantly, it gives you a chance to manage objectively and proactively churn.

References

1. Hyers, Ken. "Stemming the Rising Tide of Wireless Churn". Cahners In-Stat Group. April 27, 2000. Online. Internet. <http://www.instat.com/catalog/cat-wp.htm>. April 25, 2001.
 2. Hyers, Ken. "Stemming the Rising Tide of Wireless Churn". Cahners In-Stat Group. April 27, 2000. Online. Internet. <http://www.instat.com/catalog/cat-wp.htm>. April 25, 2001.
 3. Hyers, Ken. "Stemming the Rising Tide of Wireless Churn". Cahners In-Stat Group. April 27, 2000. Online. Internet. <http://www.instat.com/catalog/cat-wp.htm>. April 25, 2001.
 4. Downing, Douglass and Jeffrey Clark. *Business Statistics*. Hauppauge, NY. Barron's Educational Series. 1985. 229-231.
 5. Brase, Charles Henry and Corrinne Pellillo Brase. *Understandable Statistics* 5th Edition. Lexington, Massachusetts: D. C. Heath and Company. 1995. 409-412.
-