Automated Business Intelligence for Instant Game Supply

Dr. Stephen Wade

We show that a set of business rules combining pack-level data, store-level parameters and game-level performance metrics can be used to completely automate the specification of instant tickets to be supplied to retail stores. Stores served by the automated method in a 26-week test increased sales and were relatively protected against the clogging effect of a slow-moving game, compared to stores served by conventional means.

By Dr. Stephen Wade, Washington's Lottery

Background

The Lottery sales representative who calls on a particular store knows things about that store that no one else knows. An analyst at Lottery headquarters, who can compare that store to hundreds of similar stores, may know things about the store that the representative does not know. Both kinds of knowledge may be of use in handling the routine, week-to-week business of the Lottery, such as supplying each retailer with instant tickets. Thousands of small decisions (like "Should we send the store another pack of Lucky Sevens this week?") are made somewhere in the Lottery every day. Is there any way for the representative's knowledge and the analyst's knowledge to influence these small decisions? And would that have a positive effect on the business?

Washington's Lottery now has experience that answers both of these questions, "Yes." This experience was gained in a project called the Business Rule Test, or BuRT. BuRT combined three distinct levels of information in attempting to determine the "best order" of instant tickets to send to a store. The most granular or "lowest" level of information dealt with individual games and particular packs of tickets, associated with a store. At a distinctly higher level, sales representatives provided information about how they wanted the store to look, in terms of categories of games, and headquarters described the sales history of each store in terms of these categories. At the highest level, we created performance metrics both for categories of games and for individual games, at a statewide level.

We used these three levels of information in precisely-defined and consistent ways to determine how many packs of each game to send each of 60 retailers, each Monday for 13 weeks. A computer program was at the heart of BuRT. Each week this program calculated the number of tickets of each game likely to be needed at each retailer. It estimated the number of tickets of each game on hand in each store. It specified an order that would cover the retailer's needs for a set period of time.

Sales representatives have direct control over a set of parameters that factor into the calculations for each store. The need for each game is calculated according to the statewide popularity of the game in the previous week, and the usual trade in each retailer. Acceptance of the system by both sales representatives and retailers has been high – although the original 'test' set of retailers numbered only 60, more than 300 retailers are currently being served by the test system.

Although BuRT was conceived more as an enhancement to process efficiency than to sales, in fact the 60 test stores sustained an increase in instant ticket consumption, compared to retailers who have not participated, during the first six months of the program.

We thus consider that the business rules expressed in the test software are sound, and provide a sufficient basis for full automation of instant game orders. The next section of this article describes these business rules, and the final section discusses outcomes.

Business Rules

Games belong to categories. Currently in Washington we distinguish 10 categories of games, differentiated on price, play style, and prize type.

We calculate each retailer's average rate of sales in each category. We represent this in tickets per week. Each week we re-calculate the average rate of sales over the past several weeks.

Representatives choose, for each retailer, the maximum number of different games to stock in each category. We call this the width of the category.

Representatives choose, for each retailer, how many weeks' worth of tickets should be on hand in each category. We call this the depth of supply in the category.

We calculate the number of tickets needed by the retailer in the category as the average rate of sales, multiplied by the desired depth of inventory.

We determine which games in each category to supply, and how many tickets of each game, according to the relative popularities of the games at a statewide level.

It is in the selection of games and the apportionment of tickets that the headquarter analyst's knowledge is applied. Each week, each game on the market can be described as having a certain level of activated inventory (tickets available for purchase) and a certain rate of consumption (based on validation of winning tickets). The quotient: [rate of consumption]/[tickets available for purchase] is the turn rate of the game. The whole category of tickets can also be described by a turn rate. A game that has "average" popularity has a turn rate like that of the category. A highly popular game may have a turn rate twice or three times the category average. For convenience we express popularity as in index: [turn rate of a game]/ [turn rate of its category], where the "average" game has a popularity index (PI) =1.

Our aim is to fill each retailer's need by selecting the most popular tickets among those available to ship, and providing quantities of each game according to its current popularity. For example, suppose that the representative has indicated that we are to supply a retailer two games in a particular category. From among the games we can ship, we select the two with the highest PI: suppose one has PI =2, while the next-best has PI=1. Suppose this retailer needs 300 tickets in the category. We would aim to assure that the retailer has 200 tickets of the more popular game, and 100 of the other game.

When we launch a new game, we temporarily assume that it is the most popular game in the category. We assign it a PI of 2 in the first week, and use its market PI thereafter. This results in all retailers (who sell that category of tickets) getting a supply of the game in the first week; resupply follows depending on how the game performs. We anticipate the launch of a new game by trimming the inventory level of the game it will likely displace. At each store, in each category, there is a "least popular game in the category". This will not be the same game in all stores, since the width of the category varies from one store to another. In the week before launch of a new game in the category, we supplied only as many tickets of that least popular game as seemed likely to be consumed in the week.

We may or may not make a shipment to this retailer, depending upon how many tickets in these games appear to be on hand.

In evaluating the stock on hand, we are concerned only with the games that are currently shippable. A retailer may or may not have other games, and the automated system does not try to manage that. We evaluate each pack of tickets exclusively by looking at the validations it has produced. If no ticket from the pack has ever been validated, we consider that all the tickets in the pack are available for sale. If some tickets have been validated within the past week, we go on to calculate the total number of tickets likely sold (by multiplying the total number of validations, the overall odds of the game, and a factor that corrects for unclaimed winners in each category). If the difference between the size of a full pack and the number of tickets sold is greater than the retailer's need, we consider that the partial pack is sufficient to support the trade and we do not send another pack. On the other hand if a pack has produced no validations within the past week, we consider that it may not be able to support the trade in the next week. We do not count any tickets in this pack as available for sale.

In specifying the number of packs to ship, we round up. If we calculate a need of 210 tickets and the tickets come in packs of 100, we send three packs. We try to err on the side of too much rather than too little inventory.

The Test and its Outcomes

Stores not served by BuRT were served by our current business system. In this system, one of several inventory specialists, with access to gaming system pack information, reviews each account periodically and places an order as needed. These specialists have access to the same data used by BuRT, but typically make their decisions on the basis of whether and when packs have changed status (for example, from "activated" to "settled").

For the test, we implemented these rules as Visual Basic for Applications code within Microsoft Excel. In the first weeks, BuRT distributed explanatory calculations and "suggested orders" by e-mail to twelve sales representatives, each of whom had five stores in the test. The representatives then entered the orders for fulfillment by our warehouse. The representatives soon asked whether the orders could be uploaded directly to our shipping system. This has become our standard procedure: full automation, with informational reports going to the representatives.

Sales representatives in Washington receive incentive pay based largely on instant sales at their stores. Consequently they tend to scrutinize any innovation for its impact on sales. Their willingness to manage their retailers through the category-level controls provided by BuRT was thus a significant vote of confidence. Nor, evidently, was this confidence misplaced: objective measures confirm an increase in consumption of instant tickets from the test stores, compared to all other stores during the same period of time.

In particular, we found that during the first 21 weeks of 2011, 77% of the week-to-week variation in instant ticket consumption in the 60 test stores could be accounted for by a predictor based on consumption in all other stores that were active throughout the year. This predictor and its one-standard-deviation confidence limits are shown in the chart below, for all of 2011 up through early December. Predicted consumption each week is shown by open circles, actual consumption by solid triangles. During the first part of the year, the actual sales were sometimes above and sometimes below prediction. After the start of BuRT in June, actual consumption is consistently higher than predicted. The size of the increase is 5.7%.

Many of the sales representatives were initially skeptical that state-level information on the relative popularity of different games would be helpful in their particular stores. Experience shows that state-level information is useful; this does not contradict the idea that slightly more granular information may be more useful still.

Although representatives were free to adjust depth of inventory as they wished, most kept their settings close to the 1.5 or 2 weeks provided as default values. The effect of this is that stores that sell a high volume of tickets find just what they need when they get their weekly shipments. Stores that sell a low volume of tickets (relative to the size of a pack) may

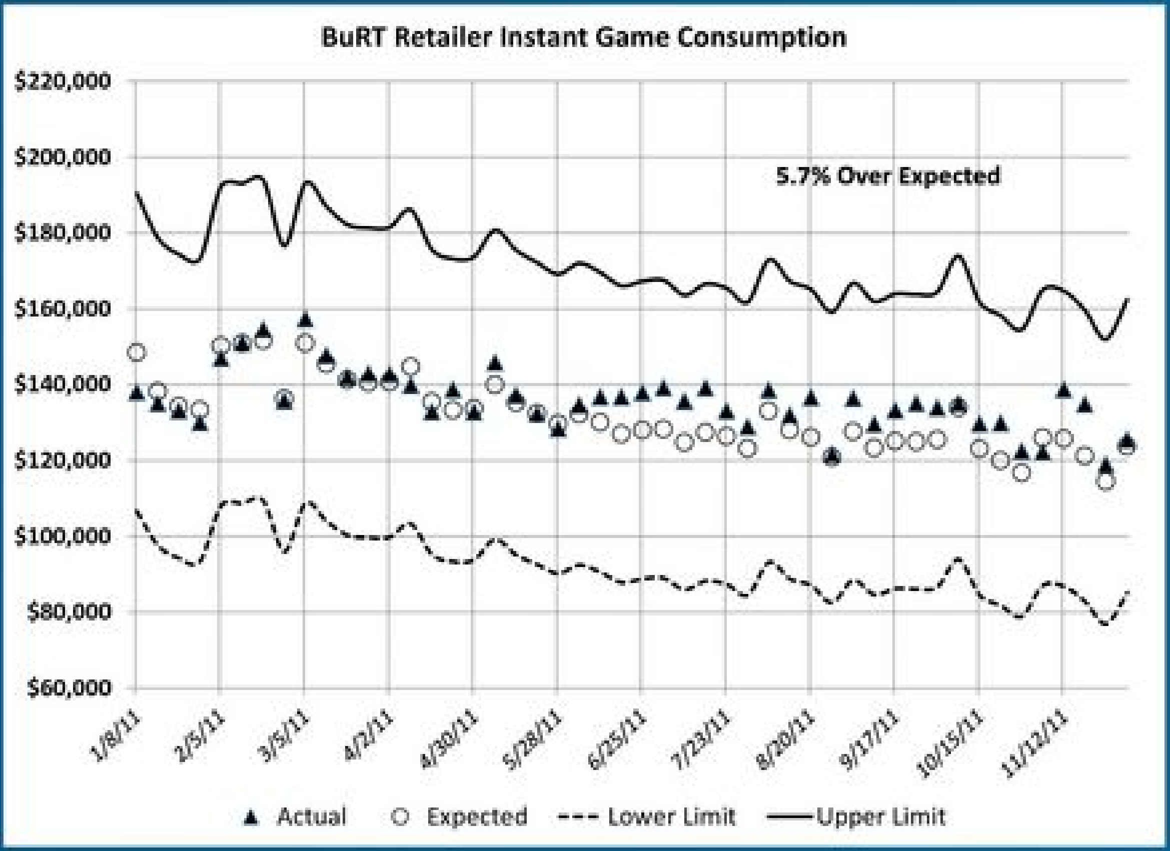

Chart 1

sometimes in fact have several weeks' supply.

No process other than physical inspection of stores is likely to be completely accurate in determining when the last ticket of a pack is truly gone. If a player buys a whole pack and holds the winning tickets without validating them, BuRT would treat that pack as "full". If players buy half of a pack and then no more, BuRT would treat that pack the same as one that was fully validated. Fortunately, some degree of uncertainty about individual packs seems to be tolerable when there are appropriate controls in the system, such as the ability of the individual sales representative to adjust depth of inventory in each category, at each store. In a store where players frequently buy and hold whole packs of $20 tickets, the representative can specify deeper inventory for this category. If in the same store $1 tickets move slowly, the representative can specify thinner inventory there.

An intended effect of these business rules is that the most popular games receive the greatest exposure across the many retailers. Conversely, when a particular game proves to be significantly less popular than others in its category, it may be offered for sale at fewer locations (since not all retailers show all available games), and it is not supplied to any location in amounts greatly in excess of need. These actions should tend to prevent "clogging" stores with slow-moving inventory.

BuRT performed according to this intent. One game launched during our test (Number 1049, "Roaring 20s") did not develop a typical popularity profile as indicated by the Popularity Index (PI) metric: it never achieved "average" popularity, and over 16 weeks on the market its PI averaged 0.6. Taking returns from retailers into account, only about 40% of the tickets printed in this game will ultimately be activated. We can contrast how inventory levels of this game developed in the state generally, and in the stores that participated in BuRT during the life of the game.

We use Activated Inventory (AI, as described above) to represent the abundance of a ticket in retail distribution. We tracked the abundance of all games in the same category as game 1049 ($2 non-extended play games) launched after March 1, 2011. Just before BuRT started in June 2011, these games had similar abundance in all stores statewide.

The slow-moving game 1049 launched in the week ending 08/10/2011. Statewide, it accounted for 28% of the AI in the category four weeks after launch. In the BuRT stores it never accounted for more than 21% of AI. Overall, the "share-of-pipeline" of game 1049 in the BuRT stores was about 70% of that in other stores. The BuRT stores were thus relatively protected against the "clogging" effect of this game.

In summary: we have shown that a set of business rules combining pack-level data, store-level parameters and game-level performance metrics can be used to completely automate the specification of instant tickets to be supplied to retail stores. Stores served by the automated method increased sales and were relatively protected against the clogging effect of a slow-moving game, compared to stores served by conventional means.

Washington's Lottery expects to benefit from increased operational efficiency as all our sales representatives learn how to manage their instant ticket business in each retailer by game categories. Although the Microsoft Excel implementation built for the test is not meant to support routine business, we are confident that the simple business rules identified above can be implemented on a suitable platform to enable routine and complete automation of the order-specification process. ■