The Business Analytics Dispatch

CFO Perspectives on Financial Data-Driven Decisions

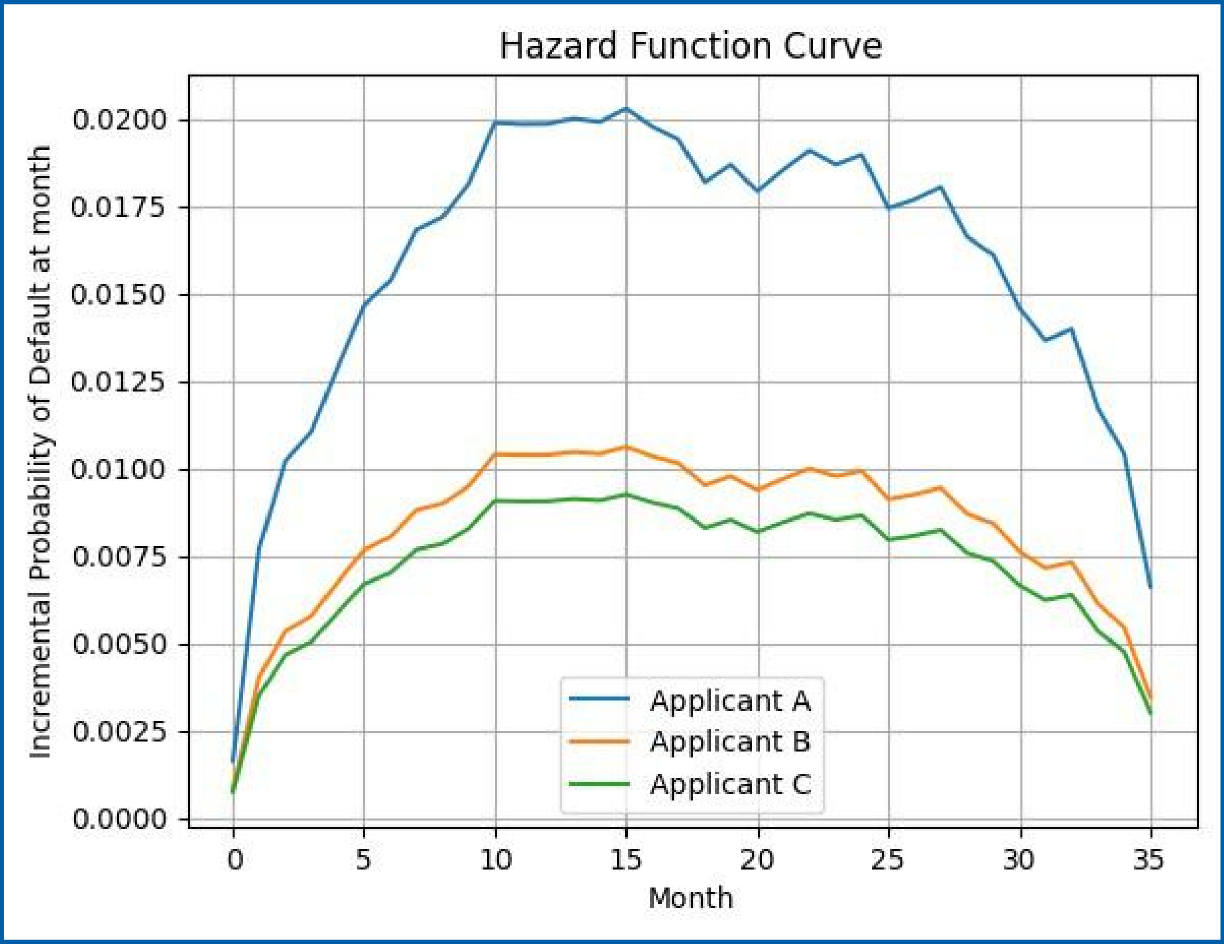

Contents

Unlocking Consumer Loan Pricing: A Deep Dive into Survival Regressi...

Sunday, August 27, 2023

In the evolving landscape of consumer lending, fintech companies have revolutionized borrower experiences, introducing real-time approvals and swift fund transfers. While tree-based classification models like XGBoost currently dominate credit scoring, survival regression algorithms are an intriguing alternative.

(Quick note: These survival algorithms extend beyond consumer credit to products with recurring payments, such as subscriptions or memberships.)

Traditionally applied in healthcare analyses, survival regression models predict the time until an event, with the “event” here signifying a potential default on a loan payment. Unlike traditional risk categorization, this model evaluates the probability of default for each payment, giving rise to a series of default probabilities known as the “hazard function curve.”

Essentially, the model provides the likelihood that a payment (be it a loan installment, gym membership fee, subscription, etc.) won’t occur as planned.

This approach deviates from static pricing for risk buckets, offering dynamic loan pricing that considers variations in risk among borrowers. Instead of assigning a risk rating and fixed price, the model calculates probability-adjusted cash flow for each payment, which aligns with fixed income cash flows.

The Cox Proportional Hazards algorithm anchors the model. Trained on historical loan data and inspired by academic research on credit risk pricing, it evaluates real-time default curves for prospective borrowers based on financial attributes.

The figures below illustrate three applicants (A, B, and C) using cumulative hazard and survival functions. A higher forecasted cumulative hazard curve implies a lower ending survival probability for the applicant.

The loan origination process involves financial analysis and interest rate application, constructing risk-weighted cash flow series. The targeted Net Present Value considers all costs, providing a comprehensive view of profitability.

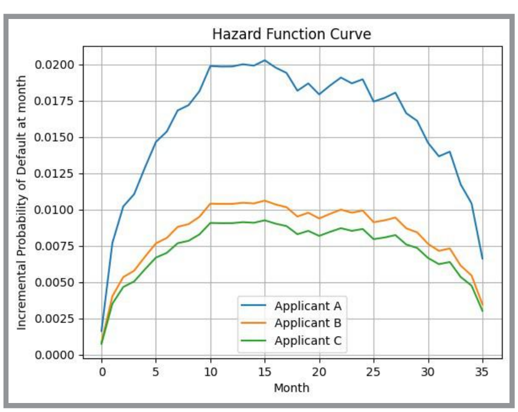

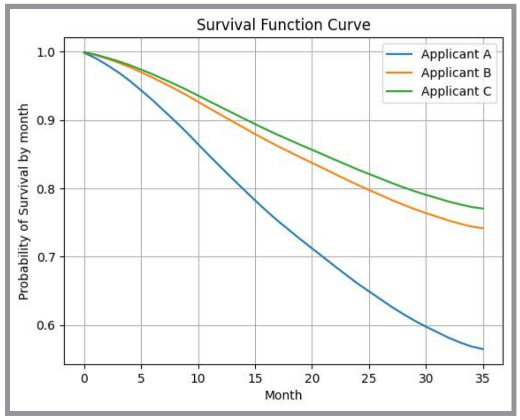

The two figures display the same three applicants: A, B and C, in two ways, using the cumulative hazard function and the survival function. The higher the forecasted cumulative hazard curve throughout the months, the lower the ending survival probability of the applicant. Hazard and survival curves visually represent the model's approach for a 36-month installment loan with three applicants. Survival regression algorithms redefine consumer loan pricing, introducing precision and adaptability to risk assessment. Feel free to challenge conventional methods with this groundbreaking model!

See my posting on Algorithmic Crypto Trading.

See my posting on Operations Research and Data Science.

FAQs about this Blog Post

1. **How do survival regression algorithms differ from traditional credit scoring models like XGBoost?** Survival regression algorithms, traditionally used in healthcare analyses, predict the time until a specific event occurs, such as defaulting on a loan payment. Unlike traditional risk categorization, which assigns a risk rating and fixed price, survival regression models evaluate the probability of default for each payment, providing dynamic loan pricing that considers variations in risk among borrowers. This approach offers probability-adjusted cash flow for each payment, aligning with fixed income cash flows, and is anchored by the Cox Proportional Hazards algorithm.
2. **What is the significance of the hazard function curve in survival regression models?** The hazard function curve in survival regression models represents the likelihood that a payment, such as a loan installment or subscription fee, won't occur as planned. It provides insights into the probability of default for each payment and allows for the evaluation of real-time default curves for prospective borrowers based on their financial attributes.
3. **How do survival regression algorithms redefine consumer loan pricing?** Survival regression algorithms introduce precision and adaptability to risk assessment by providing dynamic loan pricing based on the probability of default for each payment. Instead of static pricing for risk buckets, these models calculate probability-adjusted cash flow for each payment, allowing for more accurate assessment of borrower risk and profitability. This groundbreaking approach challenges conventional methods and offers a more tailored and informed approach to lending.

How We Replaced an Implementation of Workday Adaptive Planning Ente...

Wednesday, September 27, 2023

Excel's powerful capabilities, integrations and flexibility make it a favored tool for all financial and accounting professionals. Like many middle market companies, we considered moving from an Excel dominated financial planning and reporting process to an "enterprise grade" solution. A very difficult decision, we set aside Excel for a unified financial planning tool, also known as Enterprise Planning Management (EPM) systems.

After a review of solutions and recommendations, we decided to move our financial planning to Workday's Adaptive Planning (WAP). Our financial forecast in Excel is a complete system: It handles recurring revenue waterfalls, consolidations by products and business units, eliminations between business units, balance sheet forecasting, among other complexities. Nevertheless, the transition, despite a good plan on paper, became never ending.

We faced two core problems, which we thought we could overcome. First, the precision and complexity of our Excel forecasting model was hard to replicate in WAP. Second, our lack of deep knowledge in WAP modeling, forced heavy reliance on consultants and a time-consuming iterative process to make any headway. To minimize the obstacles and make some use of WAP, we paused our forecasting transition efforts and focused on WAP as a reporting tool. We had modest success, but we ended up having a hodge-podge system of exceptions and frequent error checking that was worse than the status quo.

During this failed transition period, the analytics team, which is part of our finance team, dramatically increased its expertise and capabilities in PowerBI. (While I am going to focus on PowerBI, I encourage the finance pros reading this to think about this solution using whatever business intelligence platform that is available. This should work with any BI platform.) PowerBI's integrations with Excel and our accounting system (Microsoft NAV) provided the light-bulb moment for moving forward with an in-house automated financial reporting system connecting our Excel forecasts to accounting results and producing polished reporting in real-time.

In order to get there, we assigned a skilled data analyst to work directly with accounting and FP&A to create an ETL (extract, transform and load) template in PowerBI that could take our GL-coded accounting records and match them to financial reports that were business friendly and consistent with our forecasting templates. Here are the key success criteria that made this possible.

1. Our data analyst had PowerBI, SQL skills needed for the entire buildout.
2. We were lucky that our data analyst also had solid accounting/finance knowledge to work directly with FP&A and accounting teammates. However, this could have been another team member working in tandem.
3. The financial reporting templates were already matched to our excel forecasting outputs. This line-for-line matching eliminated the need for another ETL template, but that could have been created if necessary.
4. Our data analyst spent time mapping GL codes to our financial reporting templates.

Without this, the ETL development would have been impossible.
5. In addition, the data analyst methodically mapped our eliminations entries between subsidiaries and hierarchical entities.
6. Then, it was time for record matching so that financial reporting template, forecast and GL Codes could be connected in sample data with a clear line of sight to each other.
7. Finally, the ETL template was ready to be programmed and tested.
8. PowerBI reporting dashboards were then developed and tested with initial data flows. Here the finance team compared PowerBI financial reports to our previous reports. Checking for accuracy at the line-item, subtotal, and total levels. Any errors were traced all the way back to GL-codes to ensure the fixes could be implemented in the ETL template.
9. We then iterated step 8 until multiple periods showed no errors and everything tied out to the most important GL line items such as net income, fixed assets, total revenue, cash balance in every grouping variation we needed (e.g., consolidated, product, business unit, geography, etc.).

The above process took about 120 days to get through Step 8 and then another 60 days (2 reporting cycles) to get through Step 9. All of this was achieved with one resource dedicated to the project and all other FP&A and accounting teammates being on call as needed.

With our financial reporting now published in an automated way, we have dramatically reduced the processing time and eliminated exceptions handling for information flows from accounting to financial reporting. While the EPM also promised financial modeling automations, we never went back to that. Instead, we have improved our Excel-based forecasting models in ways that would be hard to replicate in a new system given the resources we have and the connections of these models to our PowerBI reporting system.

If you are considering an EPM, especially for reporting, it might be worth looking at your existing business intelligence platform for an easier and more manageable solution.

See my post on Lifetime Customer Value here.

FAQs about this Blog Post

1. **Why did the company decide to transition from Excel to an Enterprise Planning Management (EPM) system, specifically Workday's Adaptive Planning (WAP), for financial planning?** The company considered transitioning from Excel to an EPM system to enhance its financial planning and reporting processes. Despite Excel's capabilities, the decision was made to move to an "enterprise-grade" solution for improved precision and efficiency. After reviewing various options, Workday's Adaptive Planning (WAP) was chosen for its promising features and capabilities.
2. **What were the core problems faced during the transition to Workday's Adaptive Planning (WAP), and how did the company address them?** The transition encountered two main challenges: difficulty replicating the precision and complexity of the Excel forecasting model in WAP, and a lack of deep knowledge in WAP modeling leading to heavy reliance on consultants and a time-consuming iterative process. To address these challenges, the company paused the forecasting transition efforts and focused on using WAP as a reporting tool. However, this approach resulted in a hodge-podge system of exceptions and frequent error checking.
3. **How did the company leverage PowerBI to create an in-house automated financial**

reporting system, and what were the key success criteria for its implementation? The company assigned a skilled data analyst to work directly with accounting and FP&A to create an ETL (extract, transform, and load) template in PowerBI. Key success criteria included the analyst's PowerBI and SQL skills, solid accounting/finance knowledge, mapping of GL codes to financial reporting templates, mapping of elimination entries, and record matching to connect financial reporting templates, forecasts, and GL codes. The process involved programming and testing the ETL template, developing PowerBI reporting dashboards, and iterating until accuracy was achieved. This approach significantly reduced processing time and eliminated exceptions handling, resulting in a more efficient financial reporting process.

Decoding Consumer Balance Sheets: A Deeper Dive Beyond Savings Rates

Friday, October 27, 2023

Navigating the landscape of consumer finance, especially in the realm of excessive debt, prompts questions about the financial robustness of consumers and its potential impact on economic trends. In the post-COVID era, media discussions often revolve around the consumer savings rate, a metric influenced by stimulus measures and changing consumption patterns. However, a recent revelation, supported by alternative data points, challenges conventional perspectives on consumer finances. This analysis delves into the nuances of consumer balance sheets, exploring the interplay between savings rates and the substantial cash build-up in checking accounts.

Alternative Data Insights:

While the savings rate serves as a valuable indicator, it falls short in revealing the depth of cash accumulation. Contrary to widely reported savings rates, a closer look at the Federal Reserve's Currency and Checkable Deposits data uncovers a more robust and sustainable financial position for the average US consumer. Comparing the cash availability evolution for the Bottom 50% and Top 50% of households reveals a significant uptick, with the former experiencing a 2.5x increase since January 2020 and the latter boasting a more substantial 3.5x surge.

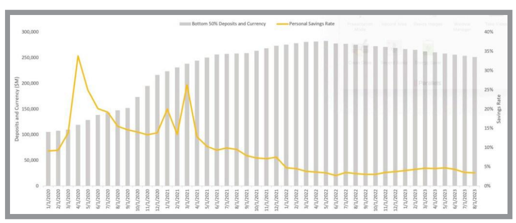

Savings Rate vs. Currency and Checkable Deposits Bottom 50% of US Households

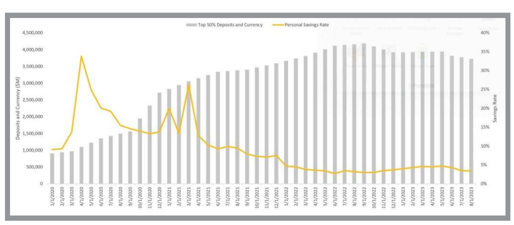

Charting the Course:

The provided charts depict the evolution of Currency and Checkable Deposits for both household groups. Notably, both segments began utilizing their accumulated cash, with the Bottom 50% initiating consumption in June 2022 and the Top 50% following suit in October 2022.

Average Consumer Balance Sheets:

Analyzing these data points underscores the resilience of the average US consumer balance sheets, with ample cash reserves and a prolonged trajectory before returning to pre-COVID levels. However, the sustainability of these balances varies by wealth decile, with wealthier households demonstrating a more protracted cash preservation period.

Erosion of Cash and Wealth Disparities:

It is crucial to acknowledge that these observations represent averages across all households, and the erosion of cash will likely manifest from the bottom up. The bottom 50% has already experienced negative growth, contrasting with the top 50%, signaling potential disparities in the impact of economic shifts. Less affluent households may face recessionary pressures while the broader economy remains relatively stable.

Forecasting Economic Trends:

While consumer balance sheets are not projected to be a significant driver of economic slowdown in the short term, factors like hiring trends, wages relative to inflation, and industrial output are expected to play more substantial roles in shaping the economic landscape. The intricate dynamics of wealth distribution and consumer behavior necessitate a comprehensive understanding for accurate forecasting of a recession or a “soft landing” scenario in the coming years.

See my blog post on Consumer Cash Balances and Credit Card Delinquencies here.

FAQs about this Blog Post

1. **What alternative data points challenge conventional perspectives on consumer finances, particularly regarding savings rates and cash accumulation?** Contrary to widely reported savings rates, a closer examination of the Federal Reserve’s Currency and Checkable Deposits data reveals a more robust and sustainable financial position for the average US consumer. This analysis uncovers a significant uptick in cash accumulation, with both the Bottom 50% and Top 50% of households experiencing substantial increases since January 2020, indicating a deeper reservoir of cash than previously acknowledged.
2. **How do the provided charts depicting the evolution of Currency and Checkable Deposits for different household groups illustrate the consumer financial landscape?** The provided charts illustrate the utilization of accumulated cash reserves by both household segments. While the Bottom 50% began consuming their accumulated cash in June 2022, the Top 50% followed suit in October 2022. These charts underscore the resilience of consumer balance sheets, with ample cash reserves and a prolonged trajectory before returning to pre-COVID levels, albeit varying by wealth decile.
3. **What insights does the analysis offer regarding the erosion of cash reserves and wealth disparities among different household segments?** The analysis highlights the potential erosion of cash reserves, particularly among less affluent households, as negative growth has already been observed in the bottom 50%. This suggests potential disparities in the impact of economic shifts, with less affluent households facing recessionary pressures while the broader economy remains relatively stable. Understanding these dynamics is crucial for accurate forecasting of economic trends and addressing wealth disparities in the long term.

Unlocking Synergies: Elevating Data Science with Operations Researc...

Wednesday, November 15, 2023

Introduction:

Who’s on a quest to develop advanced data science capabilities? One of my analytics team’s strategic expansion brought together diverse talents in statistics, applied math, and engineering. This case study explores the integration of operations research, fostering collaboration and knowledge diversity within analytics.

Objective:

Our primary goal was to blend diverse skill sets, creating an environment conducive to innovative problem-solving. While the envisioned integration remained a future prospect, immediate focus shifted to operations research for its promising prescriptive capabilities.

Operations Research Focus:

Econometrics was another area if interest for time-series analytics, but operations research, tailored for data science programming and extensive datasets, emerged as a focal point. Excelling in solving objectives within specified constraints, it offered optimal solutions that set it apart from traditional machine learning models.

Prescriptive Analytics vs. Predictive Analytics:

Distinguish prescriptive analytics (operations research) from predictive analytics (machine learning). The former provides optimal solutions based on defined constraints, while the latter predicts outcomes based on historical data.

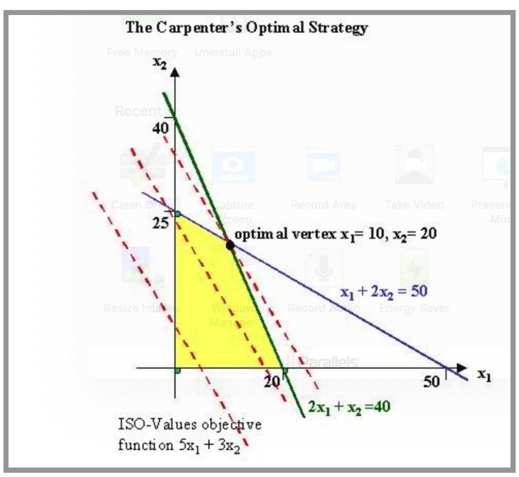

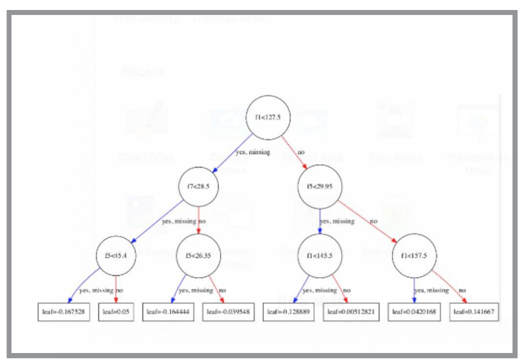

Transportation Example:

In a transportation scenario, predictive models analyze historical data for efficient routes. Operations research, however, prescriptively determines the least-cost path based on constraints, suggesting routes not traveled before. One other aspect of operations research and linear programming models is that they also handle revenue and expense variables quite well.

Methodology and Insight:

While both approaches may lead to similar conclusions, their methodologies diverge significantly. Predictive models embrace uncertainty, offering likely outcomes, while operations research precisely calculates optimal solutions, evaluating all possible choices.

Data Science Synergy:

Understanding this nuanced difference empowers an analyst or data scientist to approach problem-solving flexibly. Predictive models shine in uncertainty, providing choices based on learned experiences. Operations research excels with known inputs and complex combinations, delivering reliable solutions.

Conclusion and Future Prospects:

This case study illuminates the ongoing journey in cultivating a collaborative data science environment. As the capabilities of a team evolve, the prospect of adding talents like the previously mentioned econometrics, which excels in time-series forecasting, holds the promise of elevating capabilities to tackle even more complex challenges. Unleash the potential of data science synergy with operations research expertise! #DataScience #OperationsResearch #AnalyticsSynergy #PrescriptiveAnalytics #PredictiveAnalytics #CaseStudy

See my blog post on using machine learning to trade crypto futures here.

See my blog post on using a temporal future transformer neural network algorithm to a forecast time series here.

FAQs about this Blog Post

1. **What is the primary objective of integrating operations research within the analytics team, and why was it prioritized over other areas such as econometrics?** The primary objective was to blend diverse skill sets within the analytics team, fostering collaboration and innovative problem-solving. While econometrics was also considered for time-series analytics, operations research took precedence due to its promising prescriptive capabilities tailored for data science programming and extensive datasets. Operations research excels in solving objectives within specified constraints, offering optimal solutions that set it apart from traditional machine learning models.
2. **How does operations research, as a form of prescriptive analytics, differ from predictive analytics, particularly in a transportation scenario?** Operations research, or prescriptive analytics, provides optimal solutions based on defined constraints, such as determining the least-cost path in transportation scenarios. Unlike predictive analytics, which analyzes historical data to predict outcomes, operations research suggests routes not traveled before and handles revenue and expense variables effectively using linear programming models.
3. **What insights does the case study offer regarding the methodology and approach of operations research compared to predictive analytics?** While both approaches may lead to similar conclusions, their methodologies diverge significantly. Predictive

models embrace uncertainty, offering likely outcomes based on learned experiences, whereas operations research precisely calculates optimal solutions, evaluating all possible choices. Understanding this nuanced difference empowers analysts and data scientists to approach problem-solving flexibly, leveraging predictive models for uncertainty and operations research for known inputs and complex combinations.

Unveiling the Enigma: Contrasting Consumer Cash Reserves with Escal...

Tuesday, December 19, 2023

A recent analysis sheds light on the intriguing interplay between burgeoning consumer cash reserves and the surprising surge in credit card delinquencies. Despite the Federal Reserve's reports revealing a remarkable 2.5x increase in cash holdings for the bottom 50% of households, a deeper dive into Transunion's credit data exposes an unexpected trend in delinquency rates among recently issued credit cards.

Visualizing the Cash Peak:

Illustrating the ascent and descent of consumer cash, a Federal Reserve chart showcases the savings rate versus currency and checkable deposits for the top 50% of U.S. households. Notably, cash holdings peaked in September 2022, and while a decline is underway, the rate of decrease suggests a prolonged period before reaching pre-COVID levels.

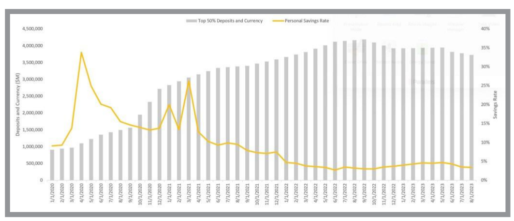

Delving into Credit Card Delinquencies:

Contrary to expectations, Transunion's October 2023 report unravels a concerning pattern in credit card delinquencies, particularly among recent issuances. The analysis, organized by vintage (issuance date), indicates a noteworthy acceleration in delinquency rates over shorter periods. Notably, the most recent vintage (Q4 2022) surpasses the pace of all its predecessors.

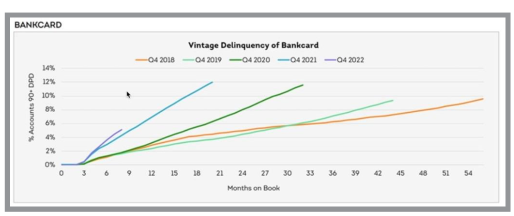

Savings Rate vs. Currency and Checkable Deposits Top 50% of US Households. Source: Federal Reserve, 2023

Analyzing the Discrepancy:

The conundrum arises – why are credit cards experiencing escalating delinquencies despite consumers holding substantial cash reserves? Several factors contribute to this apparent paradox. The Federal Reserve's data, reflecting the bottom 50% of households, conceals the nuanced distribution of cash by wealth. A deeper dive into Transunion's data suggests that stratifying credit risk tiers may unveil higher and faster delinquency rates among lower-grade credit consumers.

Unwinding Positive Effects of COVID Economics:

The favorable effects of COVID-related economic stimuli, such as increased wages, stimulus packages, newfound credit availability, and savings from stay-at-home orders, are now unwinding for lower credit tiers. This segment had the opportunity to spend more and accumulate debt during the pandemic, with banks readily extending new credit accounts. However, as the economy reverts to normalcy, these consumers face regular demands on their cash, leading to a resurgence of credit card bills with historically high-interest rates. Consequently, the combination of heightened financial commitments and mounting credit card debt is fueling a surge in delinquencies, despite the apparent abundance of cash.

In unraveling this financial paradox, it becomes evident that the intricate dynamics of consumer behavior and economic shifts necessitate a comprehensive understanding for stakeholders in the financial landscape.

See my blog post on consumer savings rates and cash balances here.

FAQs about this Blog Post

1. **What does the recent analysis reveal about the relationship between consumer cash reserves and credit card delinquencies?** The analysis highlights a surprising trend where despite a substantial increase in cash reserves among consumers, there has been a surge in credit card delinquencies. While Federal Reserve reports show a significant rise in cash holdings, Transunion's credit data uncovers an unexpected acceleration in delinquency rates among recently issued credit cards, particularly in lower credit tiers.
2. **What factors contribute to the discrepancy between rising cash reserves and increasing credit card delinquencies?** Several factors contribute to this apparent paradox. Firstly, while the Federal Reserve's data indicates an overall increase in cash reserves, it may mask the nuanced distribution of cash by wealth. Additionally, the unwinding effects of COVID-related economic stimuli, such as increased wages and stimulus packages, have led to heightened financial commitments for lower credit tiers. Moreover, the readily available credit during the pandemic allowed consumers to accumulate debt, leading to a resurgence of credit card bills amid regular demands on cash.
3. **How do economic shifts and consumer behavior interact to fuel the surge in credit card delinquencies?** Economic shifts, coupled with consumer behavior, play a significant role in driving the surge in credit card delinquencies. As the economy transitions towards normalcy, consumers, especially those in lower credit tiers, are facing increased financial demands alongside mounting credit card debt accumulated during the pandemic. Despite the apparent abundance of cash reserves, these consumers are struggling to manage their financial obligations, leading to a surge in delinquencies despite holding significant cash reserves.

Unlocking Value Creation: The Power of Lifetime Customer Value in O...

Friday, December 22, 2023

You might see it in various places as CLV (Customer Lifetime Value) or LTV (Lifetime Value). Lifetime Customer Value, or LCV, is what I call this metric. Fairly interchangeable in my experience, people who use these metrics regularly will know what you mean when you refer to any one of them. LCV's compact measurement of the value of an individual customer unit is powerful. It's something that I have been using for years and would like to share some insights into it.

Conceptually, LCV can be calculated as the net value of a customer over their lifetime. In theory, if you calculate it to a net value including allocated overhead costs for every single customer and then summed all of those individual values up, you should be very close or equal to the enterprise value of a business calculated in a discounted cash flow. These two things (Sum of LCV and DCF EV) equate because the LCV is basically taking the net present value of each individual customer's cash flow, and when summed up, it should equal the net present value of the company's cash flow. As mentioned, this assumes you calculate LCV on a net basis with accurate allocation of overheads.

If you look at it from the other direction, you could say that LCV equals EV divided by all active customers. Cable television companies often use (at least I think they still do) the EV divided by active customers to derive a value per subscriber metric for valuation purposes. This is a very easy way to examine the relative strength of the individual subscribers across companies by comparing the relative value per subscriber between companies.

The importance of these equalities is that LCV, as an operational metric usable in all areas of the organization, is tied to value creation for the entire business. If marketing pushes Customer Acquisition Cost (CAC) down, then LCV goes up, and the business should gain value. If the cost of goods sold goes up, then LCV goes down, and so does the value of the business. If an organization embraces this metric, they can push shareholder value creation alignment into many corners of a company.

When it comes to LCV, one of the main areas of focus in most cases, I find, is CAC. CAC can be volatile, especially in a world of digital marketing, where competitive forces can turn against you and make marketing very expensive in short and even sustained periods of time. As a result, if you have sticky pricing and overall operating expenses over the course of a year, CAC tends to be the part of LCV that causes the most fluctuation.

When it comes to marketing, CAC does this in two main ways. The marketing expense can fluctuate in or out of your favor, which drives LCV up or down. But given that marketing fluctuations can also translate into higher or lower customer acquisition counts, the LCV can compound its impact on the overall valuation – the sum of the LCVs as described above.

The two-by-two below illustrates the concept at a very high level.

	Spend More Marketing	Spend Less Marketing
Marketing Costs Down	LCV increases Customer counts increase EV increases	LCV increases Customer counts decrease EV remains steady
Marketing Costs Up	LCV decreases Customer counts flat EV remains steady	LCV decreases Customer counts decrease EV decreases

Another area of interest is LCV as a contribution calculation *before overhead* versus LCV as a net calculation *after overhead*. The table below shows the differences between the two calculations at a high level.

LCV-Net	LCV-Contribution
-CAC	-CAC
+Revenue	+Revenue
-COGS	-COGS
-OPEX (allocated)	N/A
-Financing Costs (allocated)	N/A
= LCV Net	**=LCV-Contribution**

I already touched on the net calculation and how it is connected to EV. The contribution calculation gives you LCV down to the contribution margin level, which is to say LCV-Contribution is the lifetime customer value that can be used to pay all overhead and financing costs of the company. This is particularly useful if the organization has steady overhead costs that don't increase quickly with the customer base. It gives the business operators a sense of how many customers they can add to the business on a marginal basis profitably. It allows the operators to take aggressive approaches to CAC and Cost of Goods/Services because every customer with some value at the contribution level will drive growth in EV. With that said, such an approach would lower net LCV over time as lower contribution clients would dilute the average LCV.

To sum it up, Lifetime Customer Value (LCV) is a powerful tool that goes beyond just numbers. It tells us the long-term value of each customer and how it connects to the overall value of a business. By paying attention to LCV, companies can make smart decisions that impact everything from marketing costs to the value of the entire business. In the fast-paced world of digital marketing, where things can change quickly, understanding and using LCV gives businesses a reliable way to plan for the future. The simple matrix and the difference between net and contribution calculations show how flexible and useful LCV can be. So, as businesses delve into LCV insights, they can uncover new ways to improve their strategies, build better relationships with customers, and set the stage for lasting success.

See my blog post on lifetime customer value here.

FAQs about this Blog Post

1. **What is Lifetime Customer Value (LCV), and how does it differ from other similar metrics like CLV and LTV?**
   Lifetime Customer Value, or LCV, represents the net value of a customer over their lifetime with a business. While it may be referred to interchangeably as Customer Lifetime Value (CLV) or Lifetime Value (LTV), LCV offers a compact measurement of individual customer value. Conceptually, LCV aligns with the enterprise value (EV) of a business calculated through discounted cash flow (DCF) analysis, assuming accurate allocation of overhead costs. This alignment signifies that LCV essentially captures the net present value of each customer's cash flow, reflecting the overall value of the company.
2. **How does LCV relate to shareholder value creation and operational alignment within a company?**

LCV serves as a critical operational metric that influences value creation across various aspects of a business. For instance, if the marketing team successfully reduces Customer Acquisition Cost (CAC), LCV increases, leading to enhanced shareholder value. Conversely, an increase in the cost of goods sold would decrease LCV and diminish the business's overall value. By embracing LCV, organizations can foster alignment in shareholder value creation throughout different departments.

3. **What are the key considerations when analyzing LCV, particularly regarding fluctuations in Customer Acquisition Cost (CAC) and overhead expenses?**
   CAC plays a significant role in determining LCV, especially in dynamic environments like digital marketing, where competitive pressures can affect marketing expenses unpredictably. Fluctuations in CAC not only impact LCV directly but also influence customer acquisition counts, consequently affecting the overall valuation of the business. Moreover, there's a distinction between LCV calculated as a net value after overhead and LCV as a contribution value before overhead. The latter provides insight into the lifetime customer value available to cover all overhead and financing costs, offering operators guidance on sustainable customer growth strategies. However, aggressive approaches to CAC and cost management may lower net LCV over time, highlighting the trade-offs involved in maximizing customer value.

An AI Crystal Ball? How We Predict Future Outcomes Using a Temporal...

Monday, January 22, 2024

Our data science and analytics teams handle and apply lots of data for insightful decision-making. Last year, I presented the data science team with a challenge: use historical data to predict a key business driver for each of the next 8 periods. We wanted to have a data-driven preview of what we might see in the in each of the next eight periods so that we could anticipate the actual outcome and make better decisions with a an eight-period.

The data science team went to work researching ways we could do this and tested a few different methodologies. We have lots of input data from our own and public sources to feed any model we wanted to test, which worked well for us. With that said, we had low expectations about finding a predictive model that produced anything reliable.

Testing different algorithms is always our approach. For the semi-technical readers, before settling on Temporal Fusion Transformer (TFT), the algorithms we tested included ARIMA, VAR, GARCH, ARCH models (univariate), Prophet, NHits, and Nbeats. TFT is an attention-based deep learning neural network algorithm. Using a mix of inputs, it produces a forecast over multiple periods in a future time horizon that you can determine. You can predict days, weeks, months quarters (really any interval is possible) into the future. Your choice.

The picture below shows the concept of how TFT works.

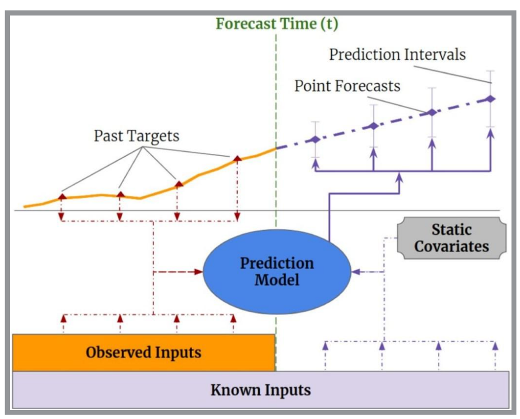

Source: Bryan Lim, Sercan Ö. Arık, Nicolas Loeff, Tomas Pfister. "Temporal Fusion Transformers for interpretable multi-horizon time series forecasting." International Journal of Forecasting. Volume 37, Issue 4, October–December 2021, Pages 1748-1764.

A continuous improvement process best describes how we developed and continue to refine the model. It's a never ending process of improvement, as a true crystal ball is never achieved.

These are the four stages of development we went through after choosing TFT as our algorithm:

- Stage 1a: Selecting all logical observed inputs and test how they drive the model. We started with over 100 and the final model only used 20. Go to Stage 1b as needed.
- Stage 1b: Refining the time intervals of the observed inputs. Since the inputs might come in varying time intervals (daily, weekly, monthly and quarterly, etc…), we needed to find methods to standardize them. You should choose an interval that matches the decision-making forecast you are producing, if you can. Go back to Stage 1a as needed.
- Stage 2: Model iteration and improvement. Complete back testing. Examine early predictions. Go back to Stages 1a and 1b as needed. At this point, you have probably settled on one or two of the most promising algorithms.
- Stage 3: Begin using in production and comparing predictions to the future periods as they unfold. Learn and refine by going back to any previous stage as needed.
- Stage 4. Continuous improvement loop. Write long-term road map. Test new inputs as they are presented. Continuous scrutiny of the predictions against what actually happens – learn and make changes by going back to any previous stage as needed.

*Note that at any of the stages of development, you can use a **TFT encoder decoder** to measure the importance of different inputs in the algorithm to learn which ones have the most impact on your prediction.*

Below are the results of our model. The orange line is the actual result of the key driver and the blue line is the prediction of the key driver that was made 8 periods ago. The area to the right without the orange line is the next 8-period forecast. So, at Period 19, we can use the blue line forecast to take action based on what Periods 20-27 tell us. When we reach Period 20, we evaluate the updated forecast and we make a decisions accordingly for the future periods. This way, we have a rolling 8-period prediction/decision cycle.

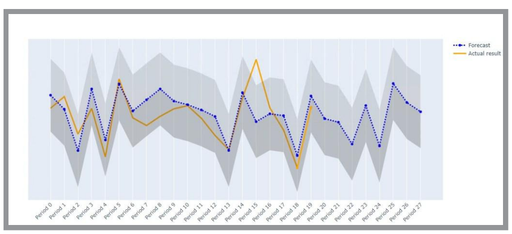

As you can see the model has been refined to a level that it makes useful predictions and handles volatility of the prediction with some reliability. Right now, we don't use this to predict the future down to the exact number, which would be ideal, but we use it for a directional understanding of where things are headed so we can make better decisions at the current decision point.

See my blog post on survival regression for predicting consumer credit risk here.

FAQs about this Blog Post

1. **What methodology did the data science team employ to predict key business drivers for future periods, and how did they refine the process?**
   The data science team tackled the challenge of predicting future business drivers by testing various algorithms, ultimately settling on the Temporal Fusion Transformer (TFT), an attention-based deep learning neural network algorithm. Before arriving at TFT, they experimented with other methods such as ARIMA, VAR, GARCH, ARCH models (univariate), Prophet, NHits, and Nbeats. Following the selection of TFT, the

team underwent a continuous improvement process consisting of four stages: selecting observed inputs, refining time intervals, iterating and improving the model, and finally, incorporating it into production and continuously refining based on real-world outcomes.

2. **How does the Temporal Fusion Transformer (TFT) algorithm function, and what role does it play in the predictive process?**
   TFT is a deep learning neural network algorithm designed for multi-horizon time series forecasting. It utilizes a mix of inputs to produce forecasts over multiple periods into the future, allowing users to predict days, weeks, months, quarters, or any desired interval. The algorithm's effectiveness lies in its ability to handle volatility and provide reliable predictions, albeit not down to the exact number. Instead, it offers a directional understanding of future trends, empowering decision-makers to make informed choices at each decision point.
3. **What does the continuous improvement loop entail, and how does the model adapt to changing circumstances?**
   The continuous improvement loop involves ongoing refinement and enhancement of the predictive model. It begins with selecting logical observed inputs and testing their impact on the model, followed by refining time intervals and iterating to improve the model's accuracy through backtesting and early prediction examination. As the model is deployed in production, it is continuously compared to actual outcomes, allowing for further learning and refinement. At any stage, the use of a TFT encoder decoder helps measure the importance of different inputs, enabling the team to prioritize those with the most significant impact on predictions. This iterative process ensures that the model remains adaptive to changing circumstances and evolving data patterns, facilitating better decision-making over time.

Key considerations for SaaS (or any recurring revenue) financial mo...

Thursday, February 15, 2024

Build your components so the are easily expandable in time and detail

In SaaS, decoding revenue dynamics is pivotal for pushing the business forward. Let’s talk about the elements of financial modeling tailored for SaaS companies:

1. Revenue Insights:

MRR (Monthly Recurring Revenue): This quantifies the predictable monthly revenue, offering immediate insights into short-term revenue trends. In my experience, I build monthly forecasts and report on the business against the forecast monthly. Having a predictable MRR with less than 1% variance to the rolling 90-day forecast is achievable and ideal. (Of course, early in the business these variances could be higher.)

ARR (Annual Recurring Revenue): An annualized view of MRR, guiding long-term planning and providing a comprehensive overview of revenue trajectory. Often ARR is used to give investors a sense of how much revenue the business has on the books that will repeat for the following year. This gives comfort to investors who see this as a baseline of revenue helping fund the company. Personally, I think contracted backlog is a more interesting way to look at this same element of SaaS, but I will cover that another day.

Churn Rate (aka cancellation rate): Measuring customer subscription cancellation, influencing MRR and ARR. Managing and reducing churn is crucial for cost-effective customer retention. Churn has two important modeling conventions that you should consider: first, does your cancellation rate change with the age of the client or contract. This is heavily influenced by the contract duration, but if you have no duration, this is an important factor to consider. Second, when modeling, it is often easier to model client counts as retention, which is (1-churn%). Always be sure that you are applying this correctly as there is a difference between churn-to-date and churn since the last period.

2. Cost Projections:

COGS (Cost of Goods Sold): Direct costs related to delivering the software service, impacting gross margin and signaling operational efficiency. Accurate forecasting is vital for profitability projections. Cloud services and direct IT support of software delivery and up-time fit into this bucket.

Operating Expenses: Day-to-day operational costs affecting operating margin and overall profitability. Monitoring ensures business efficiency and agility. This includes more typical overheads like rent, sales and marketing costs, R&D and management.

3. Customer-Centric Metrics:

CAC (Customer Acquisition Cost): Evaluating the average cost to acquire a customer. Discrepancies between CAC and customer LTV (Lifetime Value) indicate marketing or sales process inefficiencies. CAC should include all sales and marketing costs, including sales overhead for things like a CRM software, pre-sales scheduling and sales management. If you leave these items out, you are really looking at marketing acquisition cost. It's useful in some cases to do this, but CAC, especially when you are running dynamic LTV analysis.

Retention Rate: Depending on how you want to use this, it could be a very granular financial model component. Otherwise, it can simply be the percentage of retained ARR over a specified period. The latter example again is an important metric to help convince investors you have a stable source of revenues.

LTV (Lifetime Value) aka LCV (Lifetime Customer Value): One basic approach is to calculate this as total gross profit from a customer throughout the lifecycle of the client. Personally, I like to be very granular with this and I use specific components of the above for the analysis: -CAC, +churn adjusted revenues, -churn adjusted COGS = LTV Contribution and -allocated overhead = Net LTV. In addition to the final LTV values, I look at the following ratios: LTV Contribution to CAC and LTV Net to CAC. Note: Churn adjusted revenues and expenses are very useful when you have client with changing cancellation rates over time. Pro tip: You can also look at this by subscription cohorts if you sign up a lot of contracts each month.

4. Financial Health Analysis:

Cash Flow: Tracking cash movement for informed management of working capital and expense management.

Break-Even Analysis: Predicting profitability by determining the sales volume needed to cover costs. Essential for strategic pricing and sales strategies. In this area, it's useful to look at the count, value and consistency of new contract additions in the forecast to determine when the business becomes profitable.

Understanding these components offers a complete view of the current and future operation, empowering leaders to make informed decisions aligned with growth objectives. The interplay of revenue insights, cost projections, customer-centric metrics, and financial health analysis forms the bedrock for a robust SaaS financial model.

See my blog post on PowerBI for FP&A here.

FAQs about this blog post.

1. **Question:** Why is Monthly Recurring Revenue (MRR) considered a vital metric for SaaS companies?
   Answer: MRR quantifies predictable monthly revenue, providing immediate insights into short-term revenue trends. It is crucial because it helps in building monthly forecasts and tracking business performance against these forecasts, ensuring stability and growth.
2. **Question:** How does Customer Acquisition Cost (CAC) contribute to financial modeling for SaaS businesses?
   Answer: CAC evaluates the average cost to acquire a customer and is essential for assessing marketing or sales process efficiencies. Including all sales and marketing costs, such as CRM software and pre-sales scheduling, ensures accurate financial modeling, especially when conducting dynamic Lifetime Value (LTV) analysis.
3. **Question:** What role does Churn Rate play in SaaS financial modeling, and how

should it be interpreted?
Answer: Churn Rate, also known as cancellation rate, measures customer subscription cancellations, influencing both MRR and ARR. Understanding churn is crucial for cost-effective customer retention. When modeling churn, it’s important to consider if the cancellation rate changes with the age of the client or contract and ensure correct application to avoid discrepancies between churn-to-date and churn since the last period.

Crypto futures trading can produce serious returns with predictive ...

Tuesday, February 27, 2024

Predictive models underpin many trading systems. In this post, I discuss the application to the emerging world of crypto futures.

Tradery Labs
I recently had the pleasure of doing some advisory and coaching work with a startup called Tradery Labs.

Tradery Labs is bringing futuristic predictive-modeling techniques into a highly honed system that will democratize the use of predictive algorithmic trading. The company's goal is to give an investor the tools needed to build and test their own algorithms without the need for data scientists and programmers. More on that in a future post.

I recently sat down with Tradery's head of modeling, Angel Aponte, to talk shop about his latest models in crypto futures.

Some Background on Bitcoin Futures
Unlike stocks where you can "sell short" and bet against the value of a stock, there is no concept of "selling short" *actual*bitcoin, you can only buy it or not hold it. The futures market for bitcoin changed all of that in 2017 and enabled traders to financially take a positive or negative position on bitcoin. These bitcoin futures enable traders to align their investment with their view on where bitcoin is headed. That is, traders sell futures when they expect bitcoin to decline and traders buy futures when they expect bitcoin to increase.

What are the goals of latest Tradery Labs algorithm?
Tradery has some lofty goals: the current algorithm targets a 50% annual return on investment and tries to achieve the following objectives:

1. Beat a buy and hold on the base asset
2. Have more winning months than losing months
3. The largest winning month needs to be bigger than the largest losing month
4. Worst case drawdown (singular decline in value) of 20%
5. Make money when the market goes up and when it goes down
6. Product profit overall

Like all predictive model builders, Tradery is in continuous improvement mode. The model never gets to perfection. In financial markets, this is especially important because the volatility of markets creates opportunities for new trends with new causes to develop that new information can be used to retrain a predictive model to improve its accuracy.

Testing Models
Predictive modelers always test different approaches to achieve their objectives. Tradery tested both statistical and deep learning methods for this latest project. Statistical techniques use mathematical models to predict outcomes, while deep learning methods use an algorithm to learn from the available data to make predictions. Both techniques have a rich ecosystem of free software libraries that enable flexible model building. The

key is to have reliable, reproducible tests, that you can iterate over quickly, and then validate those results in the real world. All strategies that test successfully need to be followed by months of real-world results, before trading them live with real money at stake.

Testing does not necessarily produce clear cut winning models. Modern techniques are so advanced that the top models are usually comparable. Tradery finds that models vary in deciding when to put capital at risk and then de-risk (ie, buying and selling). Some techniques do better in uptrends, some do better in downtrends, and some perform best when the market trading in a range (i.e., generally moving sideways). Over a period of time, these are the only three options that a market can be in, which makes choosing on a small differences between models challenging. In this exercise, Tradery's winning model is based on statistical techniques, not deep learning, which might be a surprise to people.

Model Performance

Tradery backtests its model using historical data that the model has never seen before to see how well it does in its predictions. This means that the model gets tested using data it has never seen before and its ability to make predictions is repeatedly tested to rate how well it will perform in real life.

The picture below shows how the model has done in backtesting. You can see that the model performed very well predicting the outcomes using data it had never seen before. Green bubbles and dotted lines means good, while read means bad.

Notice the outsized winning trades and the lack of outsized losers at the top of the picture. Those are the upward pointing green arrow heads. There are two in particular that are well above the mixed green and red arrows that are in a tight range. Those two arrow heads show two trades that drove the big uptrends in performance.

In the lower portion of the picture, note the positive returns in green for market movements going up and down. This means the model is picking the right position based on expectations of increasing and decreasing prices.

Also, note that the model positioned the trades incorrectly in a sideways market as shown in red on the left. You can see the red dotted line (zoom in) which shows losses in sideways market movement.

In this data set, the model was not wrong about any big swings, which is why there are no outsized losing trades, as stated earlier. However, this particular model did not make winning decisions when the market moved sideways.

Findings and Takeaways

Angel Aponte provided some insights into important observations from the process.

The market adapts and evolves over time.

Therefore, a model's performance will degrade over time. In addition, more traders are coming into these emerging futures markets and those new entrants create dynamics that change rapidly. As a result, the team must continually test new models and retrain existing ones.

Another important finding indicates that faster and more frequent trading is not necessarily better.

One might think that high velocity trading is a natural outcome of these kinds of models, but the trading signal that the model seeks can get noisy in short intervals making quick decisions unreliable.

In addition, the cost of commissions is an important factor when trading algorithmically. You can have a highly accurate model that loses money on the trading commissions, so including that cost in back tests is important. This is another costs that works against high-velocity trading. A model needs to cover its transaction costs to be successful.

Visit Tradery Labs here.

See my posting on Survival Regression.

See my posting on Operations Research and Data Science.

FAQs about this Blog Post

1. **What is the primary objective of Tradery Labs' latest algorithm in crypto futures trading?** Tradery Labs aims for a 50% annual return on investment with its latest algorithm. The algorithm strives to surpass a simple buy and hold strategy, maintain more winning months than losing ones, ensure the largest winning month exceeds the largest losing month, limit the worst-case drawdown to 20%, and generate profits regardless of market direction—whether it's bullish, bearish, or sideways.
2. **How does Tradery Labs test its predictive models, and what insights does it gain from testing?** Tradery Labs rigorously tests its models using historical data unseen by the algorithm. This allows for an evaluation of its predictive capabilities in real-life scenarios. Through testing, Tradery Labs discovered that model performance can degrade over time due to the evolving nature of the market. Additionally, it found that faster and more frequent trading isn't necessarily advantageous, as high-velocity trading can introduce noise into the trading signals, leading to unreliable decisions. Furthermore, the cost of commissions is a crucial consideration, as even a highly accurate model may struggle to cover transaction costs, particularly in high-velocity trading scenarios.
3. **What are the key differences between statistical and deep learning techniques in the context of Tradery Labs' algorithm development?** Tradery Labs explores both statistical and deep learning methods for its algorithm development. Statistical techniques rely on mathematical models to predict outcomes, while deep learning methods utilize algorithms to learn from available data for predictions. Despite the rich ecosystem of free software libraries supporting both approaches, Tradery Labs found that the winning model in this instance was based on statistical techniques rather than deep learning. Additionally, it observed that different models excel in various market conditions, with some performing better in uptrends, others in downtrends, and some in sideways markets, adding complexity to the selection

process.

Recurring revenue modeling can be tricky, using cancellation curves...

Tuesday, March 12, 2024

Recurring revenue modeling can be tricky, using cancellation curves can improve precision and results

How to Think About Cohort-Based Subscriber Forecasting Using...

SALVATORETIRABASSI.SUBSTACK.COM

In a recent post on SaaS financial modeling, I covered some of the main drivers that play a role in the construction of financial forecasts for SaaS and related business models. One of the most important aspects of such financial forecasts is the build out of contracted revenues. In general contracted revenues can be quite predictable, which makes the recurring revenue model so attractive to investors.

The Basics

In a **basic format**, the recurring revenue forecasting for a good financial model will have the following components to calculate the monthly revenue:

1. Average revenue per subscriber
2. Number of subscribers, beginning of the month (past bookings)
3. Number of subscribers added in the month (new bookings)
4. Composite cancellation rate (the expected % of existing subscribers who will cancel in the month)
5. Number of subscribers lost in the month (2*4, cancellations or churn)
6. Net number of subscribers (2+3-5)
7. Revenue for the month (1*6)

The image below shows a recurring revenue forecast based on the above calculations. It is necessary to understand that in this kind of model, the limited variations in average revenue and cancellation rates lend themselves to a **composite** view of the revenue build. If these underlying simplifications are reliable, the above methodology works just fine.

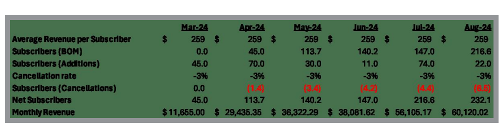

	Mar-24	Apr-24	May-24	Jun-24	Jul-24	Aug-24
Average Revenue per Subscriber	$ 259	$ 259	$ 259	$ 259	$ 259	$ 259
Subscribers (BOM)	0.0	45.0	113.7	140.2	147.0	216.6
Subscribers (Additions)	45.0	70.0	30.0	11.0	74.0	22.0
Cancellation rate	-3%	-3%	-3%	-3%	-3%	-3%
Subscribers (Cancellations)	0.0	(1.4)	(3.4)	(4.2)	(4.4)	(6.5)
Net Subscribers	45.0	113.7	140.2	147.0	216.6	232.1
Monthly Revenue	$ 11,655.00	$ 29,435.35	$ 36,322.29	$ 38,081.62	$ 56,105.17	$ 60,120.02

Pro-tip: Unless you have some strong need, I allow subscribers to be calculated in *fractions and avoid any rounding functions for subscriber counts. I find partial clients (even though there is no such thing) makes models easier to manage because rounding functions sometimes have unintended consequences and also require maintenance and awareness of their use when other people are using your model.*

More complex subscriber calculations

But, what if average revenue per subscriber changes for each new cohort of subscribers and the cancellations vary based on the age of the client. In this case, the value of the existing contracted backlog and the forecast of future contracted backlog becomes much more complex. You can stick to the above methodology, but with cancellations being age dependent, you could be in for hidden surprises and also leave your operations teams with a less refined set of objectives when they are trying to reduce cancellations.

One way to resolve this complexity is to look at a cohort-based backlog, which accounts for the average revenue variation by specifically assigning a revenue amount to a cohort and also assigning a cancellation percentage to each cohort based on its age. In this kind of model, each cohort is assigned a date of birth (sometimes called a vintage) so that it can be tracked uniquely throughout time.

A Look At Cancellation Curves

The image below shows what the cancellations would look like in a **cohort-based format**. (I am intentionally ignoring revenue variations, but this would use a similar methodology to accommodate that variation.) Notice how each month of the model needs to have a cancellation percentage for each cohort.

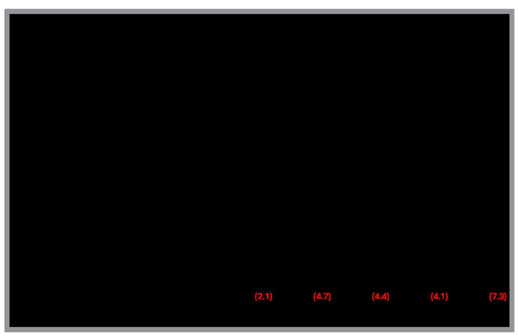

Compared to the **basic format** at the beginning of this post, the **cohort-based format** has turned into a matrix instead of being a single vector (line) of the spreadsheet. In fact, to do this precisely, each line of the **basic format** should become a matrix. Then instead of multiplying lines in Excel, you multiply across matrices to get to revenue.

Using rough math, the composite cancellation rate in the matrix is about 3% over the March to August time frame. However, you can see that the Aug-24 ending revenues in the **cohort-based format** ($59,420) are slightly lower than the **basic format** ($60,120). Now you might think that the $700 (1.1%) is not a big deal, but over time and with increased volume this variance will grow and lead to weaker forecasting. While I would love to use a simpler model for expediency, it does not stand to scrutiny when you want to have reliable forecasting of revenues.

Summary

Tracking recurring revenues is tricky and precision comes with model complexity. I find that the complexity is worth it because it instills confidence in your audiences over time and also provides the operations teams with very specific data about handling the execution on their end. For example, in the **cohort-based format** above, but not shown here, I would easily provide a forecasted cancellation count by age of the subscriber, which enables the operations team to manage their targets very specifically during the subscriber lifecycle journey.

One final note: this post only deals with the build up of subscribers in the future. If you have existing subscribers, you can use the same methodology but you should not mix the existing cohorts with the projected ones. The matrices go in different directions and they are hard to combine. Manage them in separate files if needed. I hope to do a post on that in the future.

See my previous post on SaaS revenues here.

FAQs for Recurring Revenue Modeling using Cohorts:

1. Why is cohort-based forecasting important in recurring revenue modeling?Cohort-based forecasting is crucial in recurring revenue modeling because it allows for a more accurate representation of revenue streams by considering variations in average revenue per subscriber and cancellation rates based on the age of the client cohorts. This approach provides a more granular and precise understanding of revenue projections, enabling better decision-making and operational strategies.

2. How does cohort-based forecasting differ from basic recurring revenue modeling?In basic recurring revenue modeling, calculations are simplified by using composite averages for revenue per subscriber and cancellation rates. In contrast, cohort-based forecasting assigns specific revenue amounts and cancellation percentages to each cohort based on their unique characteristics, such as date of birth or vintage. This results in a more detailed and nuanced analysis of revenue trends over time.

3. What are the benefits of using cohort-based forecasting in revenue modeling?Utilizing cohort-based forecasting in revenue modeling offers several advantages, including enhanced accuracy in predicting revenue fluctuations, better insights into subscriber behavior over time, and the ability to provide operations teams with specific data to optimize customer retention strategies. While this approach may introduce complexity, the precision it brings to forecasting can lead to more reliable financial projections and improved operational efficiency.

Managing a Loan Portfolio with Great Analytical Tools

Tuesday, March 19, 2024

Using Vector ML Analytics to Drive Success in a Consumer Loan Portfolio

In the dynamic realm of consumer lending services, precision and efficiency in financial reporting serve as linchpins for gaining a strategic advantage. Managing a loan portfolio with great analytical tools is a key component to achieve precision and efficiency. Recognizing this imperative, my team embarked on a journey to enhance our financial reporting capabilities for our consumer loan portfolio. Aiming to stay ahead of the curve, we partnered with Vector ML Analytics, a software company specializing in consumer lending financial analytics. The capabilities they provided proved instrumental in reshaping our approach to financial reporting, with a particular focus on portfolio analytics, warehouse reporting, securitization reporting, and servicer reporting. In addition, through our partnership with Vector ML, we were able to take our cash flow forecasting down to the weekly level from a previous monthly forecast. This drove significant working capital efficiencies for our asset-backed warehouse utilization.

One great feature of Vector ML Analytics that we liked was the spreadsheet interface. This is a practice we have seen in generic FP&A tools, but it makes for a familiar work environment for all the analysts touching the system. The image below shows the main screen for Vector ML Analytics platform. You can see the spreadsheet interface in the background but in the foreground you see all the functional areas: FP&A, Credit Risk, Treasure and Capital Markets.

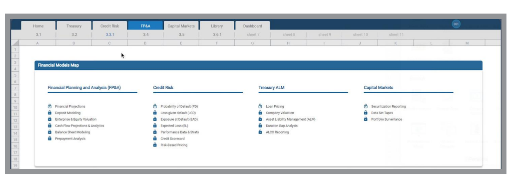

Phase 1: Data Mapping and Integration

Our collaboration with Vector ML Analytics commenced with a pivotal phase centered on data mapping and integration. Through close collaboration, our teams meticulously dissected and comprehended our existing reporting processes. This collaborative effort involved gathering and integrating crucial data and inputs into the Vector platform, including loan data tapes and assumptions essential for portfolio assets and liabilities. This foundational phase set the stage for a transformative overhaul, marked by frequent collaborative sessions and a comprehensive analysis of our prevailing methodologies.

Phase 2: Implementation and Deployment

With a thorough understanding of our unique requirements, Vector ML Analytics transitioned to the core phase of the project—implementation and deployment. This phase focused on customizing and configuring the Vector platform to not only replicate but also enhance our reporting processes. The emphasis was on leveraging automation to refine portfolio analytics, warehouse reporting, securitization, and servicer reporting. This customization aimed to integrate advanced analytics and automation, elevating our reporting capabilities beyond mere replication to empower strategic decision-making.

Enhanced Reporting Capabilities

Asset Reporting: A sophisticated cash flow engine was developed to forecast cash flows accurately under diverse scenarios, incorporating variables such as prepayments, delinquencies, and defaults. This capability provided us with deeper insights into our portfolio's dynamics, facilitating more informed strategic decisions. Concentration risks are a key area for monitoring a consumer loan portfolio. Below you can see a geographic concentration dashboard by state.

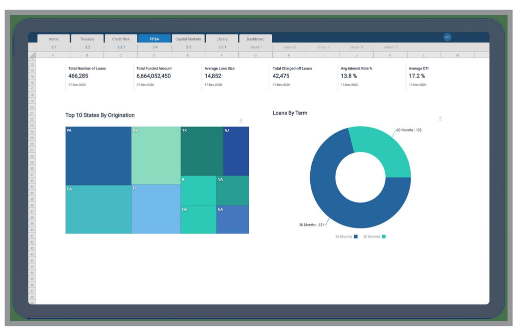

Liability Reporting: Significant advancements were made in liability reporting through the automation and enhancement of warehouse borrowing base reports and monthly servicer reports. By optimizing eligibility and concentration limits and establishing customizable performance-related triggers, Vector ML Analytics provided us with a dynamic and responsive reporting framework, enhancing risk management and compliance.

ROI Enhancement through Improved Reporting:

The collaboration with Vector ML Analytics yielded tangible improvements in speed, accuracy, and overall efficiency, embodying the ethos of "better, faster, cheaper, smarter" financial analytics.

1. 1. Speed and Efficiency Gains: Automation and bespoke features significantly accelerated our reporting processes, enhancing operational efficiency and reducing

opportunity costs.

2. 2. Accuracy and Precision: Advanced analytics and a robust cash flow engine heightened the precision of our financial reports, minimizing the risk of errors and providing a solid foundation for strategic planning and risk management.
3. 3. Cost Reduction: Streamlined and automated reporting processes led to considerable cost savings, enabling more efficient resource allocation and enhancing overall cost-efficiency.
4. 4. Smarter Decision-Making: Detailed insights empowered us to make more informed, data-driven decisions, facilitating strategic asset management and portfolio optimization.

Shocking the portfolio with different assumptions was a key requirement for us. This allows the portfolio manager to test different scenarios for risk and enable planning for adverse economic conditions outside of the company's control. Below you can see the impact of three different scenarios on portfolio balance, actual interest income and amortization.

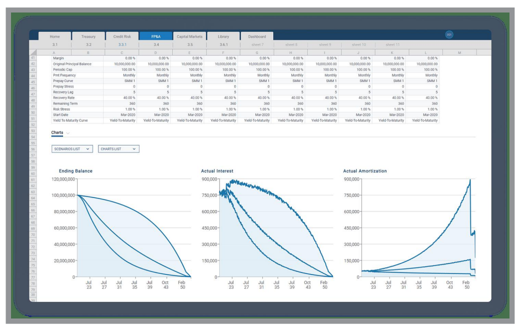

Summary

The synergy between our team and Vector ML Analytics underscores the transformative potential of technological innovation and specialized expertise in financial reporting. This partnership streamlined our reporting processes, provided critical insights for strategic decision-making, and significantly enhanced the ROI of our financial reporting endeavors. This case epitomizes the value of tailored financial analytics solutions in bolstering operational efficiency and strategic planning in the financial services sector.

Visit Vector ML Analytics here.

Visit my posting on survival regression for consumer credit risk here.

FAQs for Enhancing Financial Reporting in Consumer Lending Services:

1. How did the partnership with Vector ML Analytics enhance financial reporting capabilities for the consumer loan portfolio?The collaboration with Vector ML Analytics brought about significant enhancements in financial reporting capabilities by focusing on portfolio analytics, warehouse reporting, securitization reporting, and servicer reporting. Through advanced analytics and automation, the partnership enabled a shift towards weekly cash flow forecasting, driving working capital efficiencies and strategic advantages in asset-backed warehouse utilization.

2. What were the key phases involved in the transformation of financial reporting processes with Vector ML Analytics?The transformation journey with Vector ML Analytics comprised two pivotal phases:

- Phase 1: Data Mapping and Integration – This phase involved dissecting existing reporting processes, gathering crucial data inputs, and integrating them into the Vector platform to lay a solid foundation for the overhaul.
- Phase 2: Implementation and Deployment – Focused on customizing and configuring the Vector platform to enhance portfolio analytics, warehouse reporting, securitization, and servicer reporting through automation and advanced analytics.

3. What tangible benefits were realized through the collaboration with Vector ML Analytics in financial reporting?The partnership with Vector ML Analytics resulted in tangible improvements across various aspects:

- Speed and Efficiency Gains: Automation accelerated reporting processes, enhancing operational efficiency.
- Accuracy and Precision: Advanced analytics improved the precision of financial reports, minimizing errors for strategic planning and risk management.
- Cost Reduction: Streamlined processes led to cost savings, enabling efficient resource allocation.
- Smarter Decision-Making: Detailed insights empowered data-driven decisions for strategic asset management and portfolio optimization, embodying the ethos of “better, faster, cheaper, smarter” financial analytics.

Can I pay you to stop using TikTok? Or, will you pay to stop your w...

Tuesday, March 26, 2024

Is social media a product status trap? What is the price to stop using TikTok?

This week is a little turn away from my typical analytics discussion. We are going to talk about product economics and how consumers get value out of social media. While reading, you can ask yourself questions about the products in your business and how the buyers think about value.

TikTok is a status signaling product. Basically, being one of the cool kids – as seen by your social network – has value.

Generally speaking, consumers like status signals. They pay for them.

Take a platinum card or better yet a black card from Amex. The consumer who opts into the black card, also known as the Centurion Card, wants to signal success, wealth, access, among other things. This simple black titanium card holds sway over the owner of it because it signals a status and the owner is willing to pay for that status. Amex does give major perks in return, so it's not just an expensive stamp of metal. Still, the user gets tremendous emotional value out of the status it signals.

(Funny anecdote – in a flex of status signal deflation, a friend of mine once bent and folded his brother-in-law's brand new black card at a bar after he paid for drinks. Takes some strong fingers to bend a titanium card. I wonder if Amex delivered a new card that night as part of their lux service?)

Back to TikTok.

Consumers use social media to signal status. Again, the cool kids. Now, you might say it is free, but it costs a consumer time and ownership of their data, generally speaking, so it does have a cost.

The social media signal itself has value because everyone in the group values being seen by communicating the right signals about how great their lives are. Not dissimilar from the black card, except that almost anyone can generate some social media status.

At the same time, social media has its haters. I suppose the Centurion Card does too, but social media is ubiquitous, so we hear more about those haters.

As it turns out, based on some recent research, most social media users probably would opt out of their social media shackles – for a price. But, who is paying whom might surprise you.

Imagine getting paid to free yourself from the habit of social media. Admit it: It's an attractive proposition. Get paid to not have to worry about creating status. What a relief.

Recently on Freakonomics radio, I listened to a podcast ("Are You Caught in a Social Media Trap?") where host Stephen Dubner interviewed two economists, Leonardo

Bursztyn, professor of economics at the University of Chicago and Benjamin Handel, professor of economics at the University of California, Berkeley. The pair of academics studied the effect of payments on social media use.

Now, economists tend to be mathematical types using experimentation that has definitive outcomes. For example, I pay you a $1 in an experiment and either you do something or you don't. To understand the social media trap and its value they had to use surveys which are not transactional and definitive but hypothetical. (The costs were too high to actually pay people in a large research group, I suppose.) Regardless of the this fact, the respondents attitudes on payments to shut off social media had consistency and logic to them.

Before we delve into the results, ask yourself this: How much would I have to pay you to stop using TikTok or your favorite social media?

Now ask yourself: How much would I have to pay you to stop if everyone that matters to you on social media also stopped?

Chances are if you want to breathe a permanent sigh of relief from the social media pressure you feel, you probably would accept less to have everyone in your social circle dump the platform you use. To me, having my whole network exit at once has more value and less individual pain so it's easier to accept at a lower price. Get rid of the FOMO. Get time back in your life. Delay myopia.

The economists surveyed college students across different campuses to find out the hypothetical value of cutting off from social media for a price. They chose this group because of the ubiquity and usage of social media and they focused their questions on TikTok.

So what were the magic numbers? According to Handel and Bursztyn, the interviewees an average would accept $50 to deactivate their own TikTok account for a month. I can't say if $50 sounds low or high, but there is a price to get these students to stop for a month.

So how much less would the college students accept to take themselves and their social media group off TikTok for a month? It should be less than being the only one to accept exile for pay. After all, if everyone is off, the social status signal doesn't really exist anymore. No FOMO.

Now, it turns out that these same respondents did not want any payment to turn off TikTok, **if their social group also went off the TikTok**.

Instead, the average respondent would **pay** for this to happen. I didn't expect that. The surveyed students said they would **pay$30** to have their friends' accounts and their own account deactivated for a month.

So, there is a social media product trap. You need to be there to signal status and have to suffer through the FOMO if you are not there. You want to get paid for excluding yourself. But, when everyone one else disappears with you, there is no harsh reality of exclusion and on average the surveyed students would be willing to pay for that freedom.

Status signal products have a cost and consumers can feel trapped inside them. If there were no trap, these alternative scenarios would have no value to the consumer and they wouldn't consider taking money (or they would set a very high price) for behavioral change.

How do other products you use fit into this paradigm? Luxury products in general have some form of a status trap. What about in your business? Are there products or features of products that buyers would rather do without? Does a product have an embedded aspect that actually decreases its value and therefore price?

Some things to ponder….

Listen to Are You Caught in a Social Media Trap?

Visit the researchers Leonardo Bursztyn and Benjamin Handel.

Visit Freakonomics radio.

FAQs:

1. Why do people pay for status-signaling products like the black card from Amex or engage in social media to signal status?

Answer: People pay for status-signaling products or engage in social media to signal status because it fulfills their desire to be perceived positively by others. Status symbols like the black card convey success, wealth, and exclusivity, while social media allows individuals to showcase curated aspects of their lives to gain social approval and validation from their peers.

2. Is there a downside to using social media for status signaling?

Answer: Yes, there can be downsides to using social media for status signaling. While it offers a platform for self-expression and connection, it also comes with the pressure to maintain a certain image, leading to feelings of inadequacy, comparison, and addiction. Additionally, concerns about privacy and data ownership are prevalent in the digital age, raising ethical and psychological considerations.

3. How do economists measure the value of cutting off from social media?

Answer: Economists often use surveys to measure the hypothetical value of cutting off from social media. While direct experimentation may be challenging due to practical and ethical reasons, surveys provide insights into individuals' attitudes and willingness to accept compensation for behavioral changes. By examining responses across different demographics, economists can gauge the perceived value of social media use and its implications on consumer behavior.

0104189894

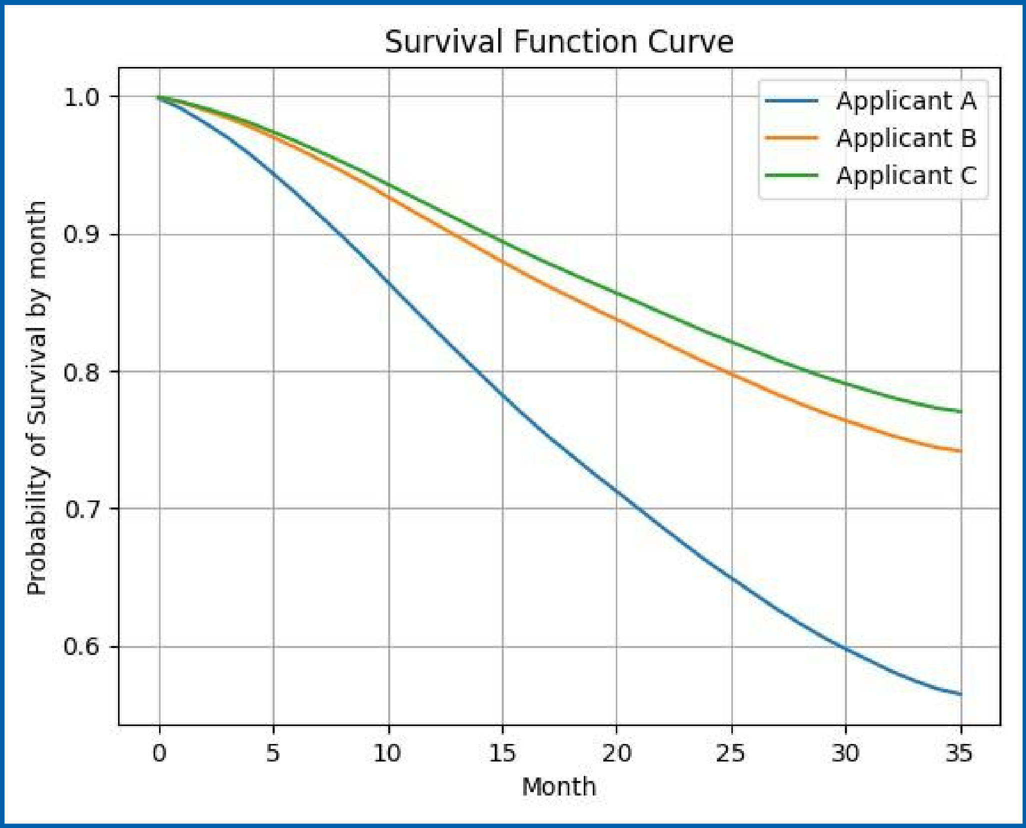
Survival Function Curve
Applicant A
Applicant B
Applicant C
Probability of Survival by month
1.0
0.9
0.8
0.7
0.6
0
5
10
15
20
25
30
35
Month