How to Calculate Customer Lifetime Value (CLV)

Revolutionize Business Strategies by Knowing Your Best-Fit Customers

In this white paper you will learn:

- What CLV is
- How it is calculated
- CLV modeling principles
- The use cases for CLV

Overview

Customer Lifetime Value (CLV) is the single most important metric for you to know because it demonstrates who your best customers are and what they have in common. This will give your business the intelligence and insights that can help optimize business practices to lower customer acquisition costs, reduce churn, and increase retention. Using this metric effectively revolutionizes the approach to data-driven decision making. It is also a great metric for tracking the health of the business and the quality of your customers.

In this white paper, you will learn what CLV is, how it is calculated, CLV modeling and its use cases, and principles to applying CLV to business practices. It will discuss practical applications for multiple areas of the business to apply predictive CLV to enhance business performance.

What is CLV?

Customer Lifetime Value (CLV) is the single most important metric for you to know because it demonstrates who your best customers are and what they have in common. Using this metric effectively revolutionizes the approach to acquisition and retention marketing. It is also a great metric for tracking the health of the business and the quality of your customers.

Why is CLV Important?

CLV is a metric that estimates how much value, usually revenue or profit margin, any given customer will bring to your business over the course of the total time they interact with your brand—past, present, and future. It consists of three different factors:

- Average order value
- Purchase frequency
- Lifespan as a customer

Rather than taking the average of these three variables for all your customers in aggregate, it's important to calculate CLV at the individual customer level. That's because your customers are individuals and they have different attributes from one another, so you need an individual CLV prediction for each one if you're going to use this metric to its fullest potential.
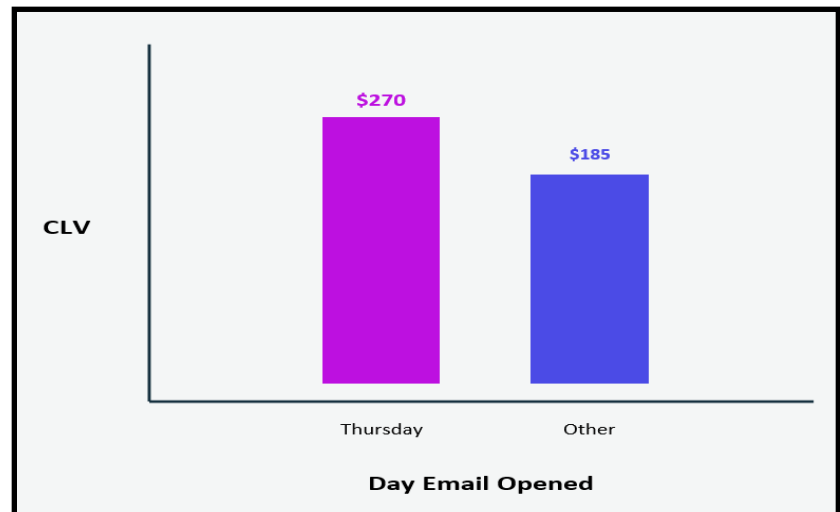
Why does this matter? Aggregate CLV simply does not allow you to adjust who you're targeting with re-engagement or acquisition efforts to maximize the benefits of each campaign. A good CLV model assesses the commonalities of all the customers unique to your business, then combines that information with per-customer behavior to predict future purchases and future dollars spent with your brand.

Developing a CLV model provides more than just one useful metric. In the process, your data analysis will also discover the most important characteristics that indicate higher or lower customer value—you will start to learn who the best-fit customers are for profitability. Characteristics are personal attributes that can be based on behavioral or demographic information, like whether someone opened an email on a Thursday or a Saturday, or what IP address they typically purchase from when shopping online.

What Can You Do With That Information?

Both the features and final output of a CLV model provide key data points for modern brands looking to focus resources on the most impactful touchpoints possible—but it's what you do with these data points that's most important. The CLV metric has many use cases, but the top five we have found to be the most impactful for DTC brands include:

- Data-driven customer segmentation
- Customer journey optimization
- Customer lifetime value modeling
- Digital marketing optimization
- Reducing churn & increasing retention

CLV model features can be the foundation for data-driven customer profiles. Your data scientist can apply a clustering algorithm to these features to identify customer segments. The resulting goal-driven segments can give you an idea of what your best customers engage with most, what they dislike, and how you might go about attracting new customers like them.

You can also combine CLV predictions and customer acquisition costs (CAC) to understand and optimize customer journeys. The goal is to get your customers to a specific conversion point with the minimum amount of friction or cost. The most effective way to accomplish this goal is to use a machine learning algorithm to examine the various paths your customers take. Then, the algorithm scores each path based on the costs to acquire each customer relative to their CLV.

Once you have these scores, your data scientist can use reinforcement learning to consider a sequence of touchpoints and determine which next interaction will maximize the chances of a certain outcome, like a repeat purchase. This means you can use CLV to predict how much revenue each of your customers is likely to bring in and encourage them to make more purchases over time. By optimizing communications and offers to customers, you can increase CLV strategically. Similarly, you can use CLV to identify your most profitable channels.

Conversely, digging into the least profitable customer journeys and profiles can help you identify and address points of friction in your existing marketing or product strategies. Don't forget to explore the model features associated with the strongest negative impact on CLV. Does purchasing a specific product often result in an increased interval before making another purchase, or even an untimely end to a customer lifespan? Perhaps some issues with a product are

causing dissatisfaction that can be addressed. Sending a survey to your lowest CLV customers is another good way to help you uncover potential areas for improvement in your business.
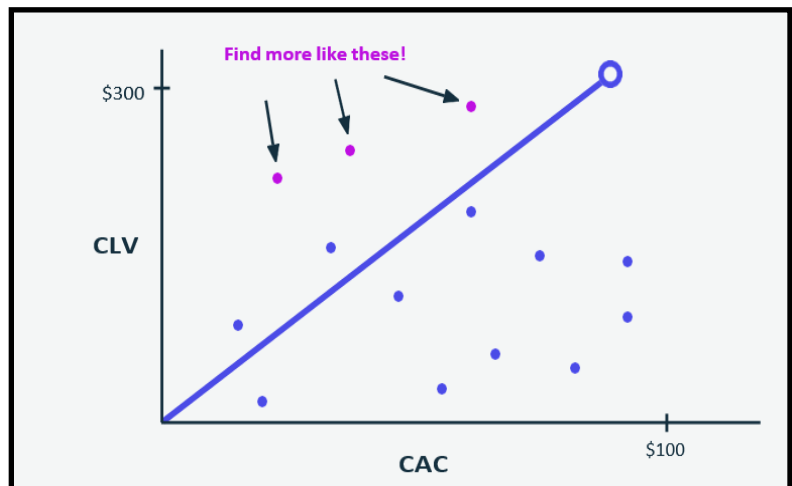
What Can CLV Solve for Your Business?

You may be thinking you've already got some of the above use cases covered, but it's very common to rely on aggregate survey information or intuition when creating resources like customer personas or journey maps. While these sources can lead to valuable internal conversations around better reaching your customers, you risk making critical strategic decisions based on an echo chamber unless you use data-driven CLV metrics to maximize the benefits of your marketing campaigns.

If you've ever felt like it might be impossible to accurately measure the ROI of your marketing campaigns, you're far from alone. Measuring ROI has been among the biggest challenges in an increasingly data-driven world. Predicting CLV and conducting regular benchmarking tests of your customer database and customer journeys are the best solutions to this challenge. It allows you to show exactly how much revenue campaigns are generating.

On the flip side, a newer challenge is to understand and respond to changes in paid social and search advertising platforms. For example, today's heightened privacy changes and uniform targeting strategies have made Facebook ads much more expensive, yet overall reach has actually declined.

CLV can help companies understand how marketing and customer experience efforts are performing without bias, develop strategies to rehabilitate under performing channels, and prove the value of all these efforts to executive stakeholders.

What's required to calculate CLV in your business?

Computing CLV at the individual level is not as easy as one might think. There's more to the process than firing up Excel, doing a little searching online, and coming up with a number.

There's no Excel formula for predicting CLV for each customer in a way that considers their purchasing rhythms and sensitivities to price individually.

The reason the process is so complicated is that calculating CLV perfectly requires complex information, like the amount spent by each customer and the frequency and duration of their purchases. How much data you have available matters too. Models cannot be trained accurately without a large enough sample size to get an understanding of variation in the individual personalities that interact with your brand.

That's why to calculate individual-level CLV without the aid of a software solution, you'll need:

- A clear definition of what it should represent for your business
- A data scientist trained in machine learning
- Sophisticated models developed by researchers
- A lot of customers with over a year of recorded purchasing activities
- Comprehensive, valid data about your customers and their behavior
- A lot of computation power
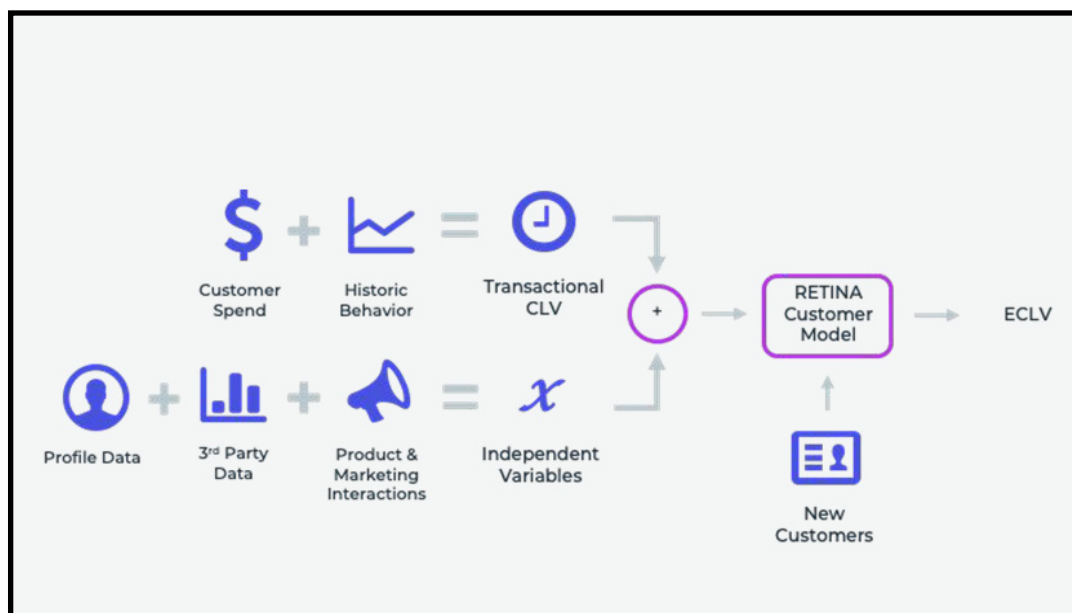
What's the Process?

Now that you've got everything you need, let's jump into a high-level exploration of the steps to building and operationalizing a CLV model in your organization.

Step 1 - Locate necessary data: There are five types of relevant datasets in a database—transaction history, customer demographics, profile data, marketing actions, and product, website, and app interactions.
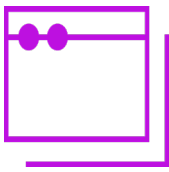
Step 2 - Forecast every existing customer's behavior: Using the data you've just gathered. The goal is to predict the number of future transactions, predicted churn date, and expected spend amount. Time series analysis is one way for your data team to accomplish this step.

Step 3 - Split the data: Split into two sets for training and testing. The training data will form the basis of your model and reserving test data allows you to benchmark how well the model performs on new information before pushing it live. A common standard is to use 80% of your data for training and 20% for testing.

Step 4 - Add characteristics and attributes: Make sure to use as many characteristics you can for your customer data sets. Once you push it into production, your model should ideally be able to make predictions about brand new customers. That's an additional challenge because new customers don't have much, if any, previous data to measure. It's important to identify some features or attributes with a high degree of impact on CLV predictions that occur early in the customer lifecycle. It's even possible to predict CLV before a person becomes a customer by building segments and lookalike audiences.

Step 5 - Train the model: Now the dataset is ready for a combination of modern machine learning methods and the best components of Pareto-NBD data models. Choosing the model that best captures assumptions about the cause/effect relationship between variables and metrics is key. An algorithm then applies the rules of the model to your data to produce outputs like the final CLV calculation.

Step 6 - Validate the model: By comparing actual recent customer purchases to the predictions the model would have made in the past by withholding a few weeks or months of recent purchasing data in model tests. This back-testing process is a critical step because it's possible for models to produce highly accurate predictions out of the gate that do not actually generalize well to other samples.

Step 7 - Put the model into production: At this stage, insights from the model actually become accessible and actionable for decision makers. Make sure the total CLV metric and other helpful insights gleaned in the process of creating the model are delivered to key stakeholders in a fashion that's aligned with how they already like to work—whether that's linking to their favorite dashboard via an API or some other more custom solution.
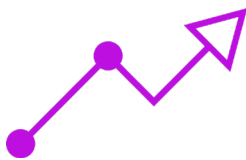
Step 8 - Finally, watch for model drift: Or other lapses in accuracy or usefulness over time. CLV changes as new data about each customer is observed, especially when marketers take effective action based on their CLV metrics to increase any combination of their customers' purchase value, purchase frequency, and customer lifespan. Make sure to re-calculate CLV on a regular basis and compile your findings in a customer migration analysis to better understand how the value of each of your customers changes over time.

Remember, model drift can be indicative of very positive evolutions in the impact of your marketing campaigns once your team puts CLV metrics into production. Customer migration analysis is not only a good data science practice, it can also show you how dramatically CLV has made an impact on your unique business.

How to Calculate CLV Summation

CLV is the single most important metric because it estimates how much value any given customer will bring to your business over the course of time and demonstrates who your best customers are — and what they have in common. It gives you powerful business intelligence and insights to make data-driven decisions for better business performance.

It's important to calculate CLV at the individual-customer level, and not the aggregate level. Your customers are individuals and have different attributes. These unique attributes can be leveraged to discover which characteristics indicate a higher or lower customer value and which customers are your best-fit customers for profitability.

Once you get into early and predictive CLV modeling, coupled with AI/ML, calculating CLV can be a challenge. The process is complicated and requires complex information and clean data. By following the steps mapped out in this white paper, you can get started on discovering how CLV can impact your business. If you do not have the time or resources for complex modeling, testing, and remodeling, finding a software or services solution that will help you increase speed-to-insights (that utilizes industry-proven early and predictive CLV models) would be the best course of action to bolster your analytics and business intelligence functions.

A few of our customers include:

Schedule a demo to see what Retina can do for you.

Retina empowers brands to make data-driven business decisions that boost revenue, all while maximizing customer-level profitability. Schedule a demo to speak with one of our experts.

To learn more go to: www.retina.ai