THE DEFINITIVE GUIDE TO

Predictive Marketing

the 8 predictions marketers need today

to transform the customer experience and optimize customer lifetime value

SAILTHRU

intro

Predictive intelligence is one of the top capabilities that marketers are focused on adopting this year. With good reason: the ability to know the next action of an individual customer is gold, quite literally, as this knowledge translates directly into revenue.

Predictions are not new, but what's now possible is revolutionary. Significant advancements allow today's marketer to use data science to generate and apply true predictions at the individual user level, rather than reporting to forecast at the segment level.

It's the difference between knowing who recently purchased, how many times and for how much versus knowing that a specific customer is 99% likely to purchase. Imagine how your marketing strategy, messaging and budget allocations would change if you had this information at your fingertips?

But with any major advancement, technology messaging landmines threaten a marketer's ability to identify the technologies that can actually deliver on the predictive promise.

This guide details a clear path forward by explaining how predictive intelligence must be applied to be useful, efficient and cost effective; and identifying the eight predictions, out of the myriad possible, that are the most telling, actionable and powerful to marketers.

the Science of Modern Predictive Technology

Historically, marketers have used multidimensional data cubes to power predictions. The most effective and frequently used approach is RFM modeling, where marketers create segments based on purchase Recency, Frequency and Monetary value to determine their best customers. By combining this data with linear regression analysis, marketers could then pinpoint the propensity for a segment or cohort of customers to take a specific action (most likely to purchase).

It sounds like a lot of work because it is. Worse, it's incredibly inefficient in terms of cost, time and accuracy. Stitching together the data means exporting and normalizing, which can be expensive depending on the resources needed, and incredibly time consuming, which can mean the results are already nullified by new data coming in. And in the end, you'll have a best guess at what that segment of customers may or may not do, which is a bit of a fool's errand considering individual actions will vary tremendously as human action is quite unpredictable when based on previous behaviors alone.

Modern predictive technology differs greatly.

The raw materials used by today's predictive intelligence tools to create an accurate prediction are multiple streams of time series data – data points consisting of successive

measurements made over a specific time interval. Think sunrises, ocean tides and even the daily close of the S&P 500. Using time series data, as opposed to individual or disparate data points alone (i.e. purchase recency, frequency and value), allows data to exist within context and cadence, enabling a better understanding of customer actions. A simple example: a cohort of customers new to your brand may behave very differently than a cohort of high-spending, loyal customers.

Depending on the predictions you are looking to make, these time series may be attached to any of a variety of variables, such as the number of messages opened per day, clicks, and even sign-up times. To be accurate, a predictive algorithm needs to look at dozens of these variables – again, over time – amassing hundreds of data points per user.

Next come the guts of a software's predictive strength – models. The most advanced predictive tools automatically build millions of models every day and then test those models for accuracy using recent data and actual behaviors to prove efficacy. For example, Sailthru's predictive intelligence tool, Sightlines, excludes the most recent month of customer data, and then examines how well a given model, if fed data up until the most recent month, predicts customer behaviors. When successful, that model is used to make predictions on that given day. In the end, one model per client per prediction is selected as the most accurate—for that particular day. Every 24 hours, the entire process is repeated.

Why does the model have such a short lifespan? Because the world changes. In an extreme case, the model that best predicts user behavior on Thanksgiving will probably not be the one that has the most explanatory power on Black Friday. Other, more subtle forms of seasonality may also come into play. Or brands may change strategy or push out discounts. For a prediction to remain relevant, the model needs to adapt.

A data scientist would tell you that in reality it's much more complicated. While that is true, the basics remain unchanged. Marketers would still do well to start by asking any predictive intelligence provider the following key questions:

1. **What data is used and what is the process to incorporate new data collected over time?**
2. **How are rich, cross channel time series of user behavior incorporated?**
3. **How does their predictive tool improve over traditional segmentation methodologies, like RFM modeling?**
4. **What kinds of models are used, how are they tested for accuracy and how often are they refreshed?**

Putting Predictions into Action

The most significant potential pitfall for any marketer using predictive intelligence is when it comes to actually using that insight to engage customers. The barrier to success comes from having multiple data flows and disparate systems for housing customer data, making predictions and messaging/engagement.

Individual point solutions for predictions can be configured to make any number of predictions; however, the only way those predictions are ever brought to life is through segment level exporting, which means they are not made at the individual level and the time it takes to move data from one solution to the next dramatically decreases accuracy. Marketing Clouds suffer from the same challenge. While they are sold under one brand name, they are built through acquisition and have multiple, separate technologies, each with their own database that comprise the cloud.

In order to make predictive intelligence effective, marketers must have access to a single platform for data collection, marketing automation and predictions. By minimizing data flows, predicting behavior at the user level and automating optimization, a marketer can impact consumer behavior on a daily basis and dramatically increase revenue-bearing KPIs.

Enterprise Marketing Clouds and Predictive Point Solutions

Enterprise Marketing Clouds and predictive point solutions rely on multiple data flows, which means tremendous costs in time, money and accuracy. Any marketer goaling to improve key metrics based on predictions will see these solutions fall short as the resulting predictions are only available at the segment level.

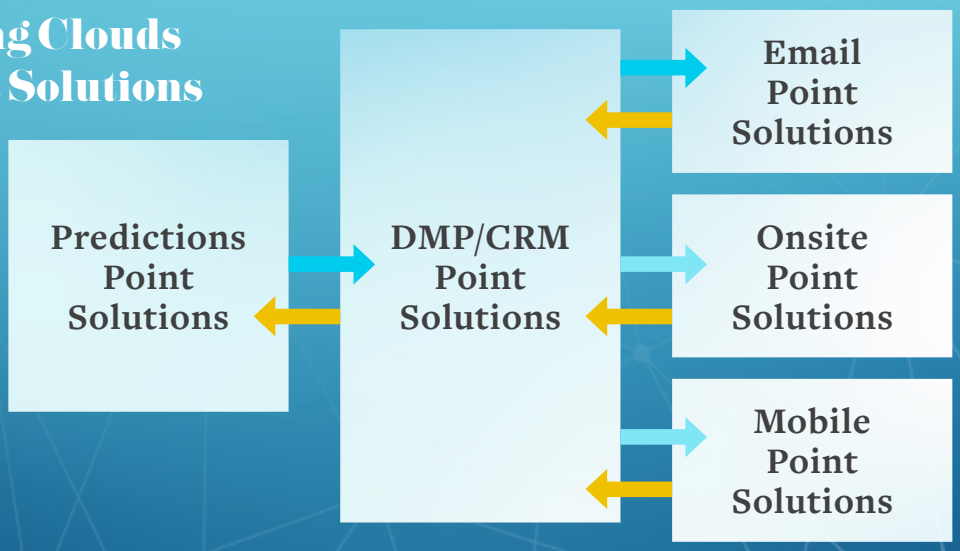

Predictions with a Single Platform

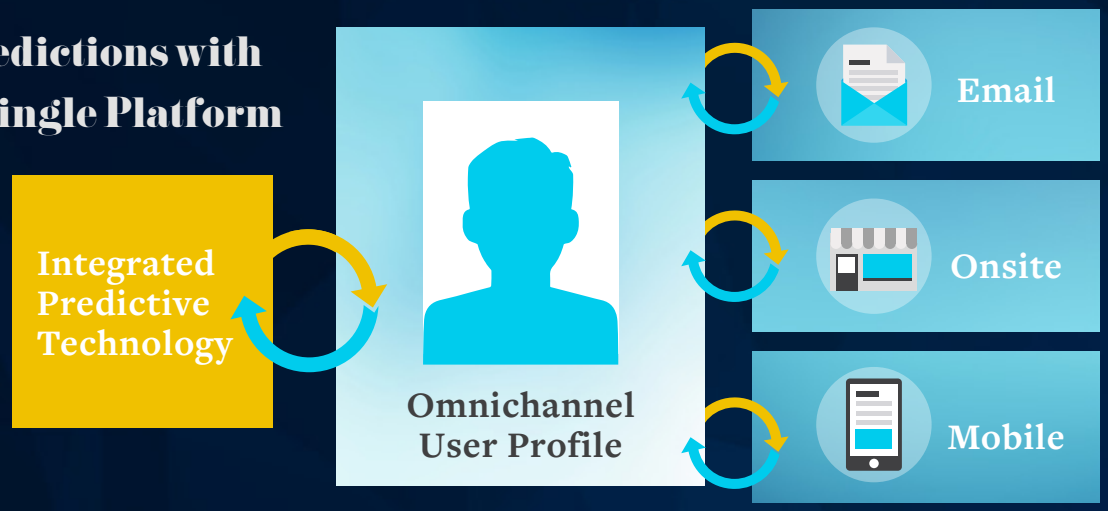

The most effective way to drive transformative revenue gains from predictions is through the use of a single platform that is natively built for multichannel data collection, cross-channel engagement and predictions all at the individual user level.

Evaluating predictive technologies? Be sure to ask these questions:

1. How long before new user behavior is reflected in recommendations and predictions?

2. Are predictions actionable at the individual user level or only as a coarse segmentation?

3. How many data flows (exports and ingests) are needed for me to leverage predictions in specific channels like email?

Defining Predictions VS. Recommendations

Everyone is familiar with online product recommendations. Amazon was one of the first big websites to use them; now the phrase, "You might also like...." is nearly ubiquitous across the commercial Web. So, they're nothing new.

What is new is that some marketing technology companies are referring to their recommendations as predictions. But recommendations and predictions are very different things, with hugely differing potential, and before any marketer makes a decision to incorporate predictive intelligence into their marketing mix they must understand this difference. Otherwise, they'll risk making decisions based on false marketing claims.

Think about how a recommendation is developed. Generally, a recommendation is drawn from previous behavioral data in a customer's history, which results in the individual being grouped into a single or multiple segments. Software technologies then use click-stream data and previous purchases (or abandonment) to recommend products that fit for that buyer based on look-a-like modeling. So that "we recommend" message is based on what customers with other similar behaviors have gone on to click or purchase. Think of it this way: recommendations rely on a basic business logic, "if bought X, recommend Y", using simply "if then" logic statements. Smart marketers, however, go far beyond this and use personalization to make recommendations based on the individual customer's explicit behaviors and implied interests and they see meaningful lift in response metrics and revenue.

In either case, saying "We predict that if we recommend this book that the user will like it," doesn't change the fact that you're making a recommendation not a prediction. The "prediction" is that the recommendation will be more effective than randomized content or editor/merchandiser selects.

Predictions are fundamentally different. Rather than relying on the type of "collaborative filtering" recommendations do, predictions make an explicit claim at a behavior or product being relevant based on a data model, not simple logic statements that incorporate 1 or 2 data points at most. At the technical level, it takes an entirely different tech stack in order to manipulate data and build the models necessary to derive predictive intelligence. At the functional level, predictions go far beyond just recommendations in both application and flexibility: by leveraging predictive intelligence a marketer can control marketing cadence, channel, discounting, product price range, content and messaging.

Combining predictions with recommendations is truly powerful, but again, don't fall for "predictive recommendations" unless a platform is truly able to communicate the difference between their predictions and recommendations.

Not sure if your technology partner is making recommendations or predictions? Ask them:

1. How specifically are your recommendations and predictions different?
2. What is the lift seen when using your recommendations, how is that lift augmented when predictions are also used?
3. How do predictions and recommendations learn from users engagement with content or messages?

Product Recommendations

Andrea

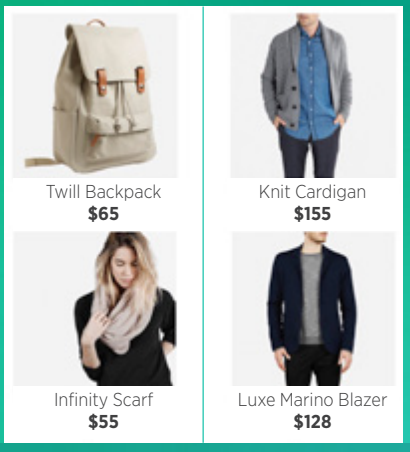

John

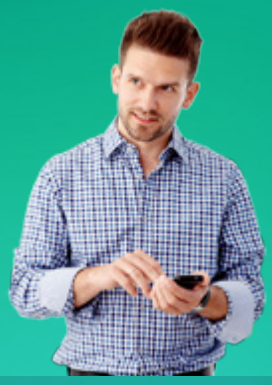

Behavioral recommendations drive response lift by serving products based on the user segment. Personalized recommendations are far more advanced and optimize conversion by serving content designed to appeal specifically to the individual.

Predicted AOV is $55

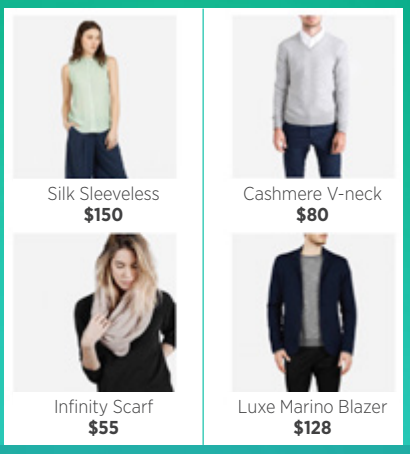

Predicted AOV is $123

Recommendations alone will demonstrate lift over randomized content; however, there is significant margin for lost revenue when not combined with specific predictions. Take "John" who is predicted to spend $123. By serving a sweater prized at $80, potential revenue is lost.

Predicted AOV is $55

Serve products that cost $60-75

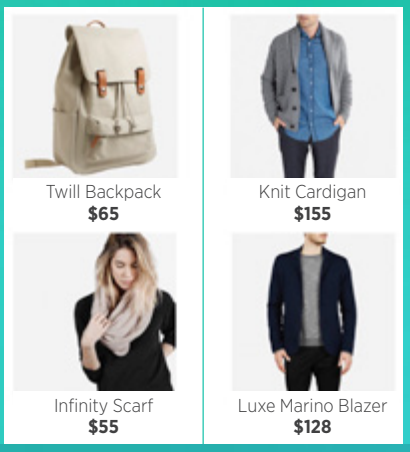

Predicted AOV is $123

Serve products that cost $130-160

By combining recommendations and predictions, a marketer can optimize content served to individual users so that you impact both conversion and revenue. This is the only approach that effectively optimizes customer lifetime value.

the 8 predictions every marketer needs

There are eight predictions that are the most telling, the most actionable, and the most powerful to marketers. While predictive marketing platforms have the power to predict any number of behaviors and results, this group of 8 are the predictions that marketers need to optimize customer lifetime value.

prediction 1

Purchase Probability

The first prediction, purchase probability, should be a major component of any marketer's strategy. If a marketer has no idea who's most likely to buy, or only a flawed idea based on segmentation, he or she can easily misappropriate marketing dollars and inadvertently hurt an organization's bottom line and brand integrity.

If you're about to start an email campaign that includes discounts, for instance, you want to offer those who are already very likely to make a purchase a different treatment than those who are less likely to buy. If someone is already willing to make a purchase, you probably don't need to incentivize him or her with free shipping or 20% off. For this group, a content-based treatment may be just as effective, and kinder to margins.

Purchase probability can, and should, be available for different time horizons, from 24 hours to 30 days. Brands that mostly see high-value, low-frequency purchases will be most interested in a customer's purchase probability over a relatively long period of time, say, 30 days.

A 7 day time frame works well for the vast majority of brands. A week is enough time to reach a customer 2 or 3 times (via email, social or other channels) and to build to a sale rather than just push one. But a flash-sale site is most interested in the immediate, and more likely to look at a customer's propensity to buy in just the next 24 hours.

Average Order Value

Knowing a customer's future predicted average order value is likely to be the linchpin of a marketer's efforts to increase that number. An average order value prediction can be combined with a customer's likelihood to purchase within a given time frame, and also with a recommendation for a specific product, increasing the power of each.

Say that the model predicts that a particular customer's next purchase is likely to be worth about $85. This is not the time to show her a recommendation for a $75 item. Instead, in order to entice her to spend more, you want to make sure that the next product recommendation is for an item that costs a bit more – perhaps between between $100 and $125.

Knowing a customer's average order value also allows brands to separate out their high-value users and to offer them different treatments within larger campaigns.

prediction 3

Total Revenue

Propensity to purchase, average order value and total revenue are all interrelated and build upon each other. Among the three, the total revenue prediction is especially important. It gets at the heart of what all marketers should be trying to optimize: lifetime customer value.

Total revenue should ideally be forecast along a longer time horizon than the previous two commerce-focused metrics (probability of making a purchase and average order value). Savvy marketers play the long game, after all, and this metric can indicate if, over the long run, their actions are enhancing or eroding the strength of the brand and their customer engagement. A total revenue number that's calculated for the next 30 days is useful because it maps nicely to the likelihood-to-purchase metric. But a total revenue prediction going out 365 days can be even more valuable.

The total revenue prediction is also an irreplaceable tool for helping marketers identify their most valuable customers. That's especially important given that 80% of a brand's future revenue is going to come from just 20% of its existing customers, according to research from Gartner Group. However, experience with predictive technology suggests that yardstick could be far too generous. Instead, it may be more accurate to say that 95% of a brand's future revenue is going to come from a mere 5% of its existing customers.

This extreme concentration in spend makes it even more important to identify the small cadre of customers that are responsible for the future of your brand. And predictive intelligence proves that traditional RFM (recency, frequency, monetary value) analysis is sorely lacking when it comes to identifying high-value customers.

When compared to Sailthru Sightlines in a specific test, RFM modeling identified only 40% of the top 10% of the brand's users. Within the top 0.3% of users, RFM identified only 60%. Overall, RFM modeling identified only about half of the customers a brand most needs to engage with, which means a massive volume of revenue would have been lost for this specific company.

prediction

4

Pageviews

The vast majority of media websites still monetize based almost exclusively on pageviews—this metric is the media equivalent to the revenue prediction that is so important to ecommerce players.

There are two ways in which this prediction must be expressed. The first is as a sorted order breakdown of users, based on their predicted total number of pageviews. The second is as an aggregated number, which can help sites better predict their advertising revenue and their ability to meet advertiser commitments.

The early results of predictive technology show that most valuable users to online publishers are even more select than those of ecommerce sites. While oftentimes only about 5% of users account for about 95% of the revenue of ecommerce sites, for publishers the concentration of users that drive pageviews can be even more dramatic. That makes correctly identifying them, and appropriately messaging them, all the more mission critical.

This prediction is also increasingly valuable to ecommerce marketers. As brands become publishers and content and commerce continue to converge, the ability for a commerce marketer to predict the volume of user pageviews will allow them to better understand the ongoing top of funnel behaviors that their most valuable customers demonstrate.

Probability of Opting Out

Every message sent by a marketer increases the probability that some user, somewhere, will opt out. But what does the profile of a user who is about to opt out look like? Only predictive intelligence has the power to answer this question. Understanding the probability that a particular user will opt out lets marketers determine if they should suppress email and mobile messages to that user until he or she is less likely to abandon the brand.

The probability-of-opting-out prediction is particularly important because the risk/reward interplay of each incremental email is more fraught than many marketers initially comprehend. When a user opts out, a brand doesn't just lose revenue for that day or for that campaign. It's much worse: A brand loses all the revenue that person would have spent in the future, as well as all the money spent getting her to convert for the first time.

With a predictive tool that pinpoints a customer's odds of opting out, marketers have, for the first time, a flag that says, "Maybe we shouldn't push to this person today."

Using a single platform for predictions and marketing automation makes acting on this prediction incredibly easy. The marketer never has to manually suppress or unsuppress a user, or debate over how long users should be suppressed for. All of this activity happens dynamically, so if the 20% of a list that is most likely to opt out is being suppressed, the probability of anyone in that group opening an email will, over time, rise on its own and the marketer will automatically engage with each user when the time is right.

Probability of Opening Email

Knowing the chances that a user will open a marketing email, or see your text message, can help marketers decide how that particular person is best approached. A customer's probability of opening her email from a brand is in some ways a proxy for their opt out prediction, because customers who don't read a brand's emails are at greater risk of opting out. A brand that isn't engaging with a specific customer is likely to lose her, period.

Customers with low probabilities of opening email can also be temporarily suppressed from campaigns, just as those with high risks of opting out can be. The greatest cause of disengagement with a brand is overstimulation. The less messaging sent, the more likely a user is to engage when they are contacted.

In addition, the more messages a brand sends that are never opened, the less impressed their ISP is going to be with the value of those emails. That's going to impact your deliverability to all users, even the rabid fans.

Strategic use of this prediction can do more than just limit losses, especially when combined with other predictions. What does it mean if you have high-value users that aren't opening your emails who also have a high probability of making a purchase within the next 7 days? That's a strong signal that you need to try to reach them through other channels. You might also try targeting just these customers with more aggressive subject lines, to see if that will entice them to open.

prediction 7

Message Volume

This number tells marketers how many messages an individual user is likely to receive from their organization, including campaigns, transactional emails and triggered messages.

This metric can be used as a frequency cap to help prevent message fatigue. No one wants to bombard their users with messages (we hope!), but many rules-based platforms don't provide transparency into the various events that can trigger an email. Marketers need to make sure a customer can't inadvertently trigger a dozen emails at once. If a customer is at or near the cap for the number of predicted messages they're likely to receive, it may make sense to set up a rule that temporarily suppresses additional campaign or non-critical emails until the number of predicted messages falls to a pre-determined number.

The rationale is simple. If a user is predicted to purchase and you're drowning them with messaging, the likelihood of that user opening a specific message might dramatically decrease while their likelihood to opt out may dramatically increase. It's why these predictions are not just valuable on their own, but even more valuable when used in tandem.

prediction 8

Predicted Click Rate

Conversion starts with a click. Every marketer knows this, it's why response metrics are often times a marketer's go-to when measuring the effectiveness of their email program, as well as onsite and mobile tests.

While we always recommend that marketers evaluate tests over lengthy time horizons to determine impact on lifetime value, opens and clicks will always be imperative because they evaluate top of funnel activity.

When it comes to email marketing and mobile texts/push notifications, marketers who can predict the click rate of individual users will understand how message cadence may need to be further adjusted and how to automate optimization of calls to action. And by predicting click rate alongside pageviews, purchases and total revenue a marketer can effectively predict and optimize messaging at every step of the funnel.

Conclusion

Predictive intelligence is a crucial tool for the modern marketer. Rather than seeking to optimize response rates or conversion alone, predictions allow marketers to reach further downstream to impact revenues and lifetime customer value. By choosing the most meaningful predictions, and combining them strategically, marketers can better prevent at-risk customers from opting out, encourage others to spend more, and build the long-term value of their brands.

Not all predictive technologies are made equal, so when considering a potential partner, use this as your vendor evaluation cheat sheet:

1	What data is used and what is the process to incorporate new data collected over time?
2	How are rich, cross channel time series of user behavior incorporated?
3	How does their predictive tool improve over traditional segmentation methodologies, like RFM modeling?
4	What kinds of models are used, how are they tested for accuracy and how often are they refreshed?
5	How long before new user behavior is reflected in recommendations and predictions?
6	Are predictions actionable at the individual user level or only as a coarse segmentation?
7	How many data flows (exports and ingests) are needed for me to leverage predictions in specific channels like email?
8	How specifically are your recommendations and predictions different?
9	What is the lift seen when using your recommendations, how is that lift augmented when predictions are also used?
10	How do predictions and recommendations learn from users engagement with content or messages?

About Sailthru

Sailthru is the only single solution that combines omnichannel data collection, automated personalization and predictive intelligence. With more than 400 customers, including commerce and media leaders such as Business Insider, Everlane, The Economist, Alex and Ani, and others, Sailthru is powering personal connections between brands and more than 1 billion consumers worldwide, every day.

Mashable

Aol.

FRANK & OAK

THE HUFFINGTON POST

ALEX AND ANI

BUSINESS INSIDER

BIRCHBOX

delivery.com
Your neighborhood. Delivered.

The Economist

EVERLANE

COUNTRY OUTFITTER
WE KNOW COUNTRY

VARIETY

SAILTHRU

www.sailthru.com
sales@sailthru.com
817.812.8689
@sailthru

NYC HQ
160 Varick St., 12th Floor
New York, NY 10013

San Francisco
360 Clementina St.
San Francisco, CA 94103

Los Angeles
7083 Hollywood Blvd
Los Angeles, CA 90028

London
18 Soho Square
London, UK, W1D 3QL

To See Sailthru in Action, Contact Us at 877.812.8689 or **CLICK HERE!**