

DATA-DRIVEN SALES

HOW THE BEST B2B SALES LEADERS USE DATA TO GROW FASTER

Copyright © 2018 by Clearbit
All rights reserved. Printed in Berkeley, California.

No part of this book may be used or reproduced in any format without written permission except in the case of brief quotations used in articles or reviews.

Producer: Matt Sornson
Editor: Nathalie Arbel
Copyeditors: Lucy Spencer, Paul Fairbairn, Megan Van Vlack
Design: Staff Only, Jay McNair
Illustrations: Nick Jarvis and Waveman Studio

Visit www.clearbit.com/books/data-driven-sales to read these chapters on the web, and contact matt@clearbit.com for more information.

CONTENTS

Acknowledgements 4

Introduction 5

Chapter 1: The Key to SaaS Pricing 6
Kyle Poyar

Chapter 2: Inbound Lead Qualification 18
Francis Brero

Chapter 3: Outbound Sales Fundamentals 29
Kyle Williams and Ben Salzman

Chapter 4: Automated Outbound Sales 41
Michael Thomas

Chapter 5: Sales Forecasting 51
Neil Ryland

Chapter 6: Finding Sales Triggers 62
Ed Fry

Chapter 7: Sales Compensation 72
David Skok

Chapter 8: Sales Operations 82
Don Otvos

Chapter 9: Sales Coaching 93
Taft Love

Chapter 10: Sales Hiring 103
Mark Roberge

About Clearbit 122

ACKNOWLEDGEMENTS

Thank you to our guest authors, Kyle Poyar, Francis Brero, Ben Salzman, Kyle Williams, Neil Ryland, Ed Fry, David Skok, Don Otvos, Taft Love, and Mark Roberge.

We're also very grateful to the folks who supported the project and provided their expertise, including Andy Wilson, Simon Whittick, Alex Bates, Guillaume Cabane, Mark Miller, Andy Jiang, David Doorman, Emmanuelle Skala, and Lauren Detro.

We'd like to acknowledge the following companies for their help and willingness to share: OpenView Venture Partners, Logikcull, MadKudu, Geckoboard, Dogpatch Advisors, Google, Kahuna, Segment, Peakon, Hull, DigitalOcean, Matrix Partners, HubSpot, SalesLoft, Yammer, App Annie, Datahug, and SmartRecruiters.

INTRODUCTION

> *"Every chapter shows exactly how a company overcame their challenges and tactical ways you can do the same. Most sales advice in SaaS is generic. This isn't."* - Hiten Shah

Data-Driven Sales offers an exclusive look inside some of the fastest growing sales teams in the technology industry and how they are using data to transform their companies. Experts from HubSpot, DigitalOcean, Google, and other industry leaders cover topics ranging from how to manage salespeople to the exact email templates that companies like Segment are using to generate millions of dollars in revenue every month.

This book covers topics including:

- Pricing your SaaS product
- Scoring and prioritizing inbound leads
- Scaling your outbound sales process
- Automating outbound email sequences
- Forecasting pipeline and revenue
- Closing enterprise deals with sales triggers
- Compensating salespeople
- Automating workflows to maximize rep productivity
- Coaching sales reps and measuring performance
- Hiring and training top sales talent

Data-Driven Sales is a blueprint for sales leaders and entrepreneurs that want to use the latest best practices to grow revenue and optimize their sales machine.

CHAPTER 1

THE KEY TO SAAS PRICING

Choosing the right
pricing for your customers



Kyle Poyar

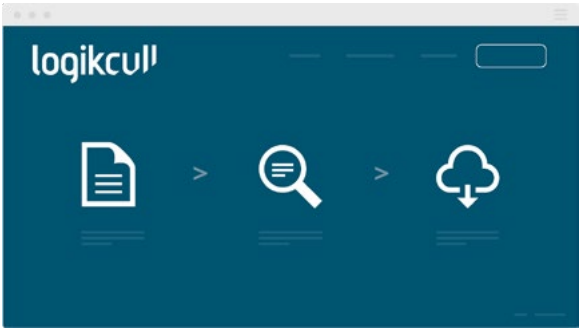
Director of Market Strategy at OpenView

In September 2009, Andy Wilson learned that *Inc. Magazine* had named his company, Logik, #181 on their list of the fastest growing businesses in the US. He had grown the e-discovery services company 1,067% over the prior three years to $4.4 million in revenue. It was profitable, too. Logik had just eight employees and brought in a $3 million profit in 2008.

That same year, the recession hit and Andy made an unconventional decision. He decided to shut down Logik's core business and transform the company from a services business model to Software-as-a-Service (SaaS). Even though his business had been highly successful, Andy realized it was ripe for disruption. "We saw the future and it was not humans managing hard drives and creating databases. That's going to be done by software. Software can do a better job, and it's going to do a better job," he says.

Andy and his cofounder, Sheng Yang, set out to rebuild Logik from the bottom up, leveraging automation and software to streamline the process of managing digital information so companies could do it themselves without needing an expensive third-party vendor. Andy counted 2,500 steps that were required to fulfill a client request as a services business. His goal was to cut that down to three steps or less with their SaaS offering. This would make it fast and easy for corporate legal teams, law firms, and investigators to organize, search through, and review a lot of unstructured data, a process they refer to as e-discovery.

After several years of development, Andy was ready to release the SaaS offering to the market, now branded as Logikcull. There was only one hitch: pricing.



He knew he didn't want to price like the services businesses of yesteryear, but beyond that he had no idea how to structure his pricing. The e-discovery market had been dominated by old-school, legacy services companies that charged exorbitant fees. "There's a data fee for everything. It's really unpredictable and can explode on you in a number of ways, like if you had 100 GB of data you needed to sift through, if you went with a legacy e-discovery provider that would probably cost well over $100,000 just to get it indexed and posted for review. And it's probably going to take a week or more of time," he tells me.

Compared to this old-school process, Andy's new technology saved time, kept data more secure, and opened up new use cases beyond e-discovery. Customers would be able to speed up the e-discovery process so it takes minutes or hours rather than weeks, cutting manual labor and associated costs. In other words, the product did a lot more than legacy systems, and he wanted to capture the additional value.

Andy's first step was to list the problems that he wanted his new pricing model to solve for his business. Two were top of mind: lumpy revenue and the commoditization of data. "One of the biggest [problems] was the lumpy revenue problem that you see in episodic types of business models ... You're really in a weird state of unpredictable business growth. So, we wanted to solve that by creating a subscription [model]. We also wanted to move away from commodity types of pricing models and the biggest commodity being quantity of data."

Andy had one other rule of thumb. He thought that in order to get lawyers to be willing to change their habits, the company would need to provide orders of magnitude more value than existing solutions.

To overcome the inertia of the market and lawyers' reluctance to buy new products, in other words, Andy couldn't just rest on having a far superior product. He also thought he needed to launch with a cut-rate initial price point.

He decided to start simple: he launched a straightforward subscription model similar to that of MailChimp. The Atlanta-based email marketing provider ties its subscription pricing to the number of marketing contacts a customer uploads and sends emails to, rather than the amount of storage taken up by those contacts. Andy decided to try a similar model and charge customers based on the number of documents they wanted to host inside Logikcull. He set an extremely low price of $0.13 per document per month, which worked out to a much more affordable rate compared to the hefty gigabyte (GB) storage fees customers were paying legacy providers. He liked that this pricing model created a steady subscription revenue stream and shifted away from the amorphous gigabytes of data stored toward a metric that he thought would be easier for lawyers to wrap their heads around.

MAILCHIMP PRICING CHART

SUBSCRIBERS	EMAILS PER MONTH	MONTHLY COST
0 – 2,000	12,000	FREE
0 – 500	UNLIMITED	$10
501 – 1,000		$15
1,001 – 1,050		$20
1,051 – 1,150		$25

Unfortunately, customers hated it. Unlike MailChimp customers, Andy's customers had no idea how many documents they were going to have, so document-based pricing wasn't any easier for customers to predict than gigabytes of storage. Even worse, Andy couldn't tell a customer how many documents they would need "because the job of Logikcull is to unpack all this information," counting up and de-duplicating all of the files that a customer puts into the system. Only after a customer used Andy's product would they truly know their document count and hence the price they would have to pay. Customers' document counts fluctuated over time, too, which made it difficult for them to commit to a given subscription tier. Andy knew this would ultimately hurt his business, since many finance teams won't sign off on a deal unless they know the exact price they'll have to pay at the end of the year.

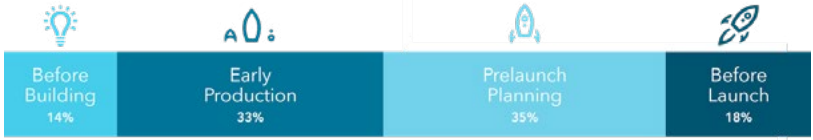
The solution: Start listening to your customers

In a survey of 1,000 SaaS companies conducted by OpenView, we learned that—like Andy—50% of companies procrastinate on pricing decisions. The survey showed that the majority of companies consider pricing either right before or as their product is getting ready for launch, rather than toward the beginning of the development process. Often, pricing decisions are made with little or no data. At most early-stage companies, pricing isn't a part of anyone's job description and gets left to a CEO who's already constrained for time.

x.ai, the AI personal assistant, took a different approach when they went to market with their initial paid offerings to great fanfare in October 2016. They carefully conducted in-depth research on pricing from beta customers, using surveys as well as in-person roundtables to gather input. Similarly, it's how Meetup built demand for and monetized their successful enterprise product[1] after trying and failing several times to crack the code on how to sell to the B2B market. Meetup asked businesses what new features they would find most valuable and how much of a premium they would pay for these enhancements.

WHEN SAAS COMPANIES CONSIDERED NEW PRODUCT PRICING

OpenView 2017 SaaS Pricing Survey (N=1,010)



CUSTOMER DISCOVERY RESEARCH INVOLVED



DO YOU HAVE PEOPLE IN YOUR ORGANIZATION DEDICATED TO PRICING?



WHO OWNS PRICING AT THE COMPANY DURING EXPANSION STAGE



HOW WELL DOES YOUR PRICING METRIC CORRESPOND TO THE VALUE DELIVERED BY YOUR PRODUCT



After Andy's initial pricing misfortune, he vowed to take a more data-driven approach to pricing that put the customer front and center. He went out and talked to dozens of Logikcull's customers and prospects and learned that customers wanted predictability more than anything else. The challenge was making that work for his business. "You need to come up with a model that satisfies that and also aligns with value in the best way possible, which is not necessarily easily predicted at the beginning." But rather than dwelling on generating a predictable revenue stream for his own business, Andy chose to create predictability for his customers.

This insight led Andy to shift their pricing model from per-document pricing, which had caused friction, toward pricing based on a combination of users and data. For example, a law firm could subscribe to a package that offered full access to Logikcull for five attorneys at a time, along with a generous 100 GB of data storage. If they ever needed more data later on, they could easily top up their subscription with a low per-GB fee for the overage, since data would not be the sole driver of the pricing. This new approach made it much easier for customers to budget for Logikcull and reinforced the value of the software independent of commodity cost of storage. It also appealed to Logikcull's wide array of customers, including law firms, corporations, and governments.

Andy's new pricing model is an example of what economists refer to as a three-part tariff. Rather than charging a customer for each incremental unit of usage, in a three-part tariff there is a base platform fee that includes an initial amount of usage. That means a customer can predict with near certainty how much they will be billed for. But a sliding scale is built in so that if they use more than their allocated amount of data, a business like Andy's isn't stuck with the bill. In other words, the customer gets 90% predictability, but the service provider (Andy) can model the pricing to ensure the maximum value is captured from each customer.

Linear Pricing (LP) - Each hosted document costs $0.10.

2-Part Tariff (2PT) - The software has a base platform fee of $10,000 and each hosted document costs $0.10 more.

3-Part Tariff (3PT) - Again, the software has a base platform fee, but the fee is $25,000 because the first 150K hosted documents are free. Each additional hosted document costs $0.15.

See Three Part Tariffs[2] *by Tom Tunguz for a full explanation*

Research[3] indicates that this type of pricing leads to higher customer satisfaction because customers don't experience the same pain that they do when they have to pay for each amount of incremental usage. For finance teams, this is especially valuable because they can accurately model all of their costs into projections. That in turn removes psychological and financial barriers and leads to faster sales cycles for companies like Logickull.

In talking to customers about his new pricing model, Andy started to pick up on other valuable insights as well, such as new use cases for the product. He trained everybody in the company to ask simple, open-ended questions like "What are you using Logikcull for?" Then, to make sure it wasn't a one-off use case, he would talk to potential customers in the same segment and see if they had the same problem.

As Andy built a longer and longer list of use cases beyond traditional e-discovery, he started training the sales and account teams to accelerate product adoption within existing accounts by expanding the number of use cases. More use cases meant more users would want access to Logikcull and more data would be brought into the software, giving Andy two separate levers to upsell customers. In one example, Andy sold a corporate legal team on the traditional e-discovery use case and expanded to six departments afterward. "We just said 'Hey, could you introduce us to the tax team? We talked with customer X and they have this problem, and apparently Logikcull is really good for that issue.' Before you know it, it starts working its way up."

Once they get a customer in the door, Andy now knows that Logikcull has a great shot at growing that customer's spend by 2x, 3x, or even up to 20x within a single year. Seeing that rapid expansion opportunity led Andy to formalize their process for identifying top candidates for an upsell. He recently instituted an account-based metric in Salesforce indicating the total available subscription they believe is possible. Now he pushes his team to achieve that goal over time rather than all at once. "I probably say this every day with the sales team," Andy said. "Just get them in the door. Stop trying to maximize the initial sale because that's just the first bite. You're going to get multiple bites."

Customer discovery further helped Andy hone the way he measured and communicated ROI to customers by accelerating their already impressive land-and-expand business model. Logikcull's main ROI is efficiency, so Andy's team keeps track of a customer's time spent before and after adopting the product. "We do before-and-after studies whenever we're kicking off a new customer. We'll ask things like 'How long does it take you today, on average?' and 'How many times do you do that?' and we'll record that. Then we'll ask, 'Now, in Logikcull, how long does it take you?' All of a sudden, we have a much more compelling story. We can arm our internal champion with really clear ROI data so that when they walk up to the person that has budget, it's a no-brainer." Depending on the use case, Andy captures different data points so that the ROI story is as compelling and personal as possible.

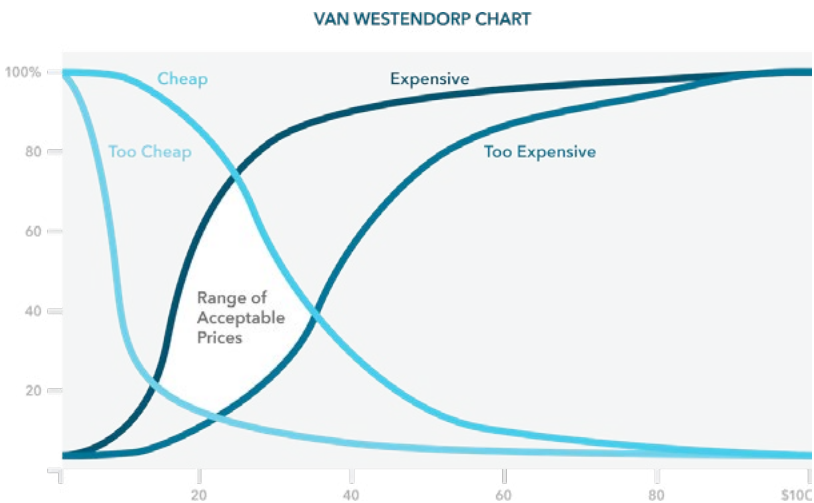
How to implement customer development into your own pricing

When Andy advises early stage companies on pricing today, he suggests that they rapidly get started with their own customer development for pricing purposes. He says, "It starts with talking to customers and asking simple questions."

What Andy's describing is the Van Westendorp Price Sensitivity Meter, a pricing technique named after its creator, economist Peter van

Westendorp, who was the first to suggest that directly asking people how much they would pay for a product would result in junk data. His method solves that problem by measuring buyers' preconceived notions of different price points. He asks at what price point a product would be (1) too cheap to have faith in, (2) a bargain worth paying for, (3) expensive enough to seem reliable, or (4) too expensive to purchase. The goal is to put the respondent in a different state of mind when they answer the questions, away from how much they have to spend and toward how they perceive the price and value of the product. By asking a series of questions, van Westendorp collects a richer set of data that reveals the price sensitivity of a buyer across a range of prices instead of just one data point.

The results of each of the four price points are then plotted together on a graph, and the intersection of the lines reveals a target price range for the product. Pricing toward the lower end of the range would enable a penetration strategy, whereby a company pursues the most adoption possible while keeping a credible price point so buyers don't doubt the product quality. Pricing toward the higher end of the range, meanwhile, enables a skimming or profit-maximization strategy.



The Van Westendorp graphs reveal key psychological thresholds in the minds of buyers, too. For instance, if a large contingent says that at $50 per month the product is getting expensive, the company would likely find a large drop-off in demand from pricing at, say, $55 versus $49. That drop-off in demand would be far larger than pricing changes further away from psychological thresholds; for instance, a price of $49 versus $43.

Unlike more complex pricing sensitivity tests, the Van Westendorp technique can be done with a relatively small sample size; the data can come from one-on-one conversations or an inexpensive online survey of buyers. But one risk of the Van Westendorp technique is that savvy buyers may lowball the price points they offer up so as to minimize the resulting price of the product. To offset that risk, economists recommend asking indirect pricing questions to get more context on buyers' thought processes and switching costs.

Direct questions

- What do you think is too expensive of a price that you wouldn't buy it?
- What's too cheap that you wouldn't trust it?

Indirect questions

- How do buyers think about the value and ROI of the solution?
- How would they pitch the business case to their boss?
- How does the buying process work, and what's the threshold for a deal to go to procurement?
- What would be the risks of no longer using the solution?

Contextual questions helped Andy realize that regardless of how much value his product delivers, when selling into government, it's best to avoid procurement at all costs. Through customer discovery, Andy found out that one government buyer had a budget threshold of $10,000 per year before a purchase had to go through a formal RFP process. This would include the dreaded security audit and lots of bureaucratic red tape. Not only would procurement slow down a deal, it could kill a deal entirely or bring a new competitor into the mix.

Andy decided that he would come up with a pricing model that accommodated this customer need, even though it was much lower than his typical deal sizes. Lo and behold, that customer "went from $10,000 a year to $200,000 a year for three years because we came in at a price point they could actually start with," Andy said. Once he got through the initial purchase hurdle, Andy knew that subsequent product expansion would be subject to much less scrutiny and procurement meddling compared to the first sale.

Looking back, Andy regrets not adopting customer discovery sooner. "If you can have enough of these conversations, I think you're going to get a much better view about the world. That would be incredibly helpful for any entrepreneur. My experience is similar, except I did it after the fact."

Despite arriving late to customer discovery, Andy points out that it's had a powerful impact on his business, especially when he adds up the dozens of pricing experiments he's conducted over the years.

Andy's SaaS business has far eclipsed his old services company, and now, armed with a new $10 million funding round, the future of Logikcull looks brighter than ever.

CHAPTER 2

INBOUND LEAD QUALIFICATION

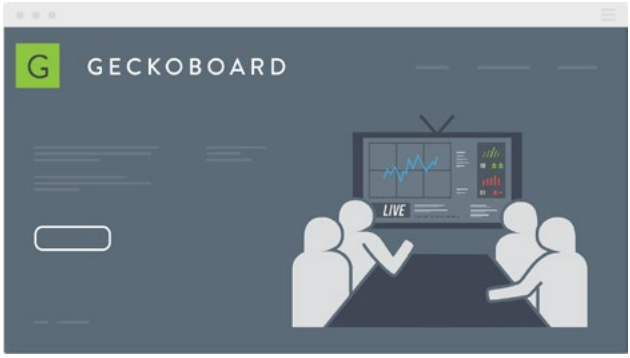
How one marketer used data and machine learning to solve his lead quality problem



Francis Brero

Co-founder and Chief Revenue Officer at MadKudu

When Simon Whittick joined Geckoboard as its first VP of Marketing, he took all the standard steps to attract more visitors to their site, convert them, and grow the SaaS company's revenue. He and his team wrote content for their popular blog, ran paid advertising campaigns, and set up email nurture campaigns. At the end of his first year, he was as successful as almost any other marketing executive in the industry. The site was attracting hundreds of thousands of visitors every month, and the business was booking millions in annual recurring revenue. But unknowingly, his success was driving one of his coworkers crazy.

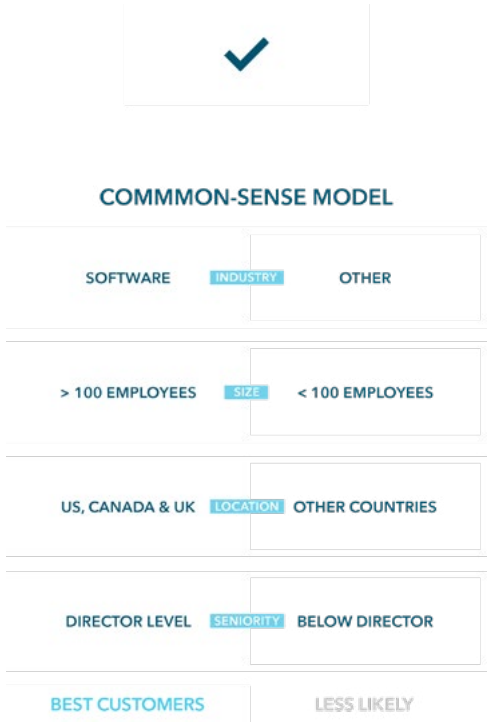


While 10,000 leads a month earned Whittick applause at the company's weekly all-hands meeting, it was keeping Geckoboard's only sales development rep (SDR), Alex Bates, at the office on nights and weekends. Many of the inbound leads were self-serve customers who required no conversation with sales, or tire kickers who were not ready to buy. This left Alex manually qualifying leads and wasting tons of his time.

As a result, Geckoboard's sales efficiency—one of the most critical metrics for any company—was slumping. In other words, Whittick wasn't only driving a junior sales rep crazy; he was leaving money on the table.

Over the course of the next year, Whittick built a data-backed machine learning process to solve his company's lead-qualification problems. In the process, he turned Bates into not only an adoring fan of his, but a one-man sales team as efficient as a typical ten-person SDR team. Without any technical background, Whittick figured out a way to change the shape of his company using data and a bit of machine learning.

One day toward the end of last year, Bates and Whittick sat down to discuss how they could solve their lead-quality problem. They had close to 10,000 leads coming in each month, but they needed to figure out which of those leads to send to sales. Their first instinct was to refine their ideal customer profile. They'd both read all the sales and marketing blogs preaching its importance. They started with a Ideal Customer Profile based on some simple audience rules.



On paper, Geckoboard's ideal customer was a software company with more than 100 employees; they typically sold to a director or VP. But the truth was that a lot of companies outside that explicit profile would be great customers. For example, their initial model excluded a company with 95 employees even if it looked almost identical to one of their best customers. When they looked at their past data, they learned that leads in what they believed to be their ideal customer profile converted at twice the rate. But they only accounted for 0.7% of the conversions. They needed a more nuanced and flexible inbound model.

BASIC ICP LEAD SCORING RESULTS

CUSTOMER FIT SCORE	% OF LEADS	% OF CONVERSIONS	RATE OF CONVERSIONS
YES	0.3%	0.7%	+2.1X
NO	99.7%	99.3%	✕

Prior to joining the Geckoboard team, Whittick had worked for Marin Software. While he was there, he began to notice a shift in the way companies approached marketing. The most forward-thinking companies had begun to hire technical employees with degrees in mathematics instead of business. He heard stories of companies that were replacing entire teams and doubling revenue by using publicly (or privately) available information and crunching it to their advantage. As time went on, many of those employees left their jobs to provide the same level of automation to smaller companies without the budget to hire data scientists.

Between his time at Marin Software and Geckoboard, dozens of start-ups popped up to help larger companies embrace the data revolution. Companies like Clearbit mined the web for demographic and firmographic data that could be used to better filter leads. My own company, MadKudu, makes it possible to pull insights from that data without

having a PhD in computer science. By 2016, the entire marketing technology landscape had shifted. With an executive team that embraced innovation and big bets, Whittick decided to make it Geckoboard's competitive advantage.

The first step Whittick took was to develop his own flexible point-based scoring system. Previously a lead was either given a 1 or a 0. A lead was either a software company with 100 or more employees or it wasn't. It was binary. The obvious problem of this model was that a company with 95 employees would be excluded. In addition, a company with 100 employees was given the same score as a company with 1,000 employees, even though the latter was twice as valuable.

In his new model, Whittick gave leads a score based on multiple criteria. For example, he'd give a lead 50 points for having 500 employees or more, and negative 2 points if it had less than 10 employees. A director-level job title would receive 10 points, whereas a manager would only receive 3. This created an exponentially larger score range, which meant that Bates could prioritize leads. If he called the top score leads, he'd have the option to call B-tier leads. The model was weighted toward the large accounts Geckoboard knew could significantly impact revenue. For example, a US-based real estate company with 500 employees and a director buyer would be routed to the top of Alex's lead list, even though it didn't fit the software industry criteria.

This new model was similar to the way SDRs have scored leads for over a decade,[4] only more efficient. Prior to automated lead scoring, sales reps were told by their managers to prioritize leads based on four criteria: budget, authority, need, and timing (or as it's commonly referred to, BANT). This method is more flexible than a rigid ideal customer profile, but it is only as strong as the rep behind it. Human error, irrational judgment, and varying levels of experience lead to a process with little rhyme or reason. That's why Whittick chose to automate the task and take humans out of the process entirely.

ADVANCED POINT BASED LEAD SCORING

CRITERIA		POINTS
NUMBER OF EMPLOYEES	> 500	+50
NUMBER OF EMPLOYEES	100 ~ 500	+40
NUMBER OF EMPLOYEES	50 ~ 100	+10
NUMBER OF EMPLOYEES	10 ~ 50	+3
NUMBER OF EMPLOYEES	< 10	-2
JOB TITLE	FOUNDER, OWNER, OR DIRECTOR	+10
INDUSTRY	SOFTWARE	+10
INDUSTRY	REAL ESTATE	-10

POINTS BASED LEAD SCORING RESULTS

CUSTOMER FIT SCORE	% OF LEADS	% OF CONVERSIONS	RATE OF CONVERSIONS
VERY GOOD	8%	11%	+1.4X
GOOD	7%	13%	+1.7X
MEDIUM	32%	52%	+1.6X
LOW	52%	24%	-2.2X

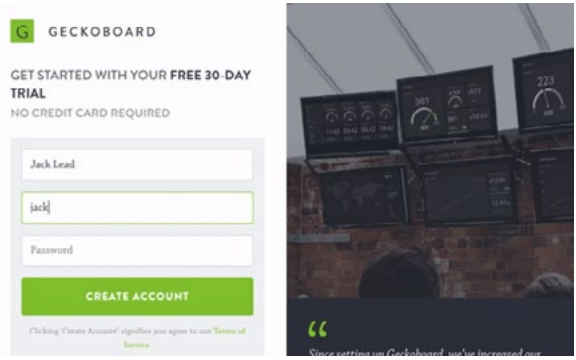
Immediately the company began to see results from their lead-scoring investment. Within the first month, leads were converting at twice the rate. As a result, Bates was spending less time to close more deals. Sales efficiency—revenue collected divided by the time and resources to earn

it—rose significantly. Still, Whittick knew he could improve the results and save Bates even more time.

One of the biggest shifts that Whittick saw in the technology industry was the speed at which data could be put to use as a result of new tools. In the old world that he inhabited, a lead couldn't be scored until it hit a company's CRM. Enrichment software took hours to append valuable data to a lead.

That information could be sent to the CRM and the lead scored accordingly before the visitor began typing in the next text box.

After his first lead scoring success, Whittick decided to make another bet. Bates frequently complained about leads that were obviously bad fits—the type of conversation that takes 30 seconds to know there isn't a mutual fit. Many of the companies were too small to need sophisticated dashboards yet. Whittick enlisted one of the company's front-end developers to help him solve the problem. They built logic into the product demo request page that would ask for a visitor's email address and then, before sending them to the next page, score the lead. On the back end, additional information would be appended to the lead using Clearbit, and it would be run through MadKudu's scoring algorithm. If it received a high-enough score, the next page would ask for the lead's phone number and tell them to expect a call shortly; if the score was low, they'd be routed through a low-touch email cadence. It was radically successful.



Before implementing their real-time lead scoring solution, only about 15% of Bates' conversations were meaningful. The new website logic meant that he could cut 85% of the calls he took every day and focus on higher quality accounts. Once again, sales efficiency increased significantly.

In addition to the speed at which information could be appended, processed, and acted on, Whittick saw another change in the marketing technology world: there was suddenly more data than most companies knew what to do with. Marketers could know what CRM, email server, and chat service a company used. They could know when a company was hiring a new employee, when they were written about by a major news outlet, and how much money they'd raised. It was overwhelming. But thanks to tools like Segment, marketers could pipe all that data into a CRM or marketing automation system and act on it. Then they could combine it with information like how frequently someone visited their own site, how often they emailed sales or support, and when they went through a specific part of the onboarding process. For a data-driven marketer like Whittick, this new world was utopia.

In conversations with Bates, Whittick learned that the best leads were ones that went through their onboarding process pre-sales conversation. During the Geckoboard free trial, users were prompted to build their first dashboard, connect a data source, and upload their company's logo and color palette. As is the case with many SaaS solutions, most users dropped off before completing all the steps. Those users weren't ready for a conversation with sales. But when Bates was looking at his lead list, he had no way of knowing whether or not a free trial user had completed onboarding. As a result, he was spending at least half of his time with people that weren't ready to talk or buy.

Combining usage data from the website and their app, Whittick set out to refine the lead scoring model even further. Each time a free-trial user completed a step in the onboarding process, it was recorded and sent back to the CRM using Segment. The model would then give that lead a couple of additional points. If the user completed all of the steps, bonus

points would be added and the lead would be raised to the top of Bates' lead queue in Salesforce. Again, Bates began spending less time talking to customers prematurely and more time having high-quality conversations that led to revenue. Whittick had figured out how to save the sales team time and increase sales efficiency further.



But while Whittick and Bates were celebrating their improved conversion rate success, a new problem was emerging. By the summer of 2016, they had enlisted my team at MadKudu to automate their lead scoring. Rather than manually analyzing conversion data and adjusting their lead scoring model accordingly, our machine learning tool was built to do all the work for them. There was a small problem. Today, machine learning algorithms are only as strong as the humans instructing them. In other words, they are incredibly efficient at analyzing huge sets of data and optimizing toward an end result, but a human is responsible for setting that end result. Early on, Whittick set up the model so that

it would optimize for the shortest possible sales cycle and the highest account value. He didn't, however, instruct it to account for churn, an essential metric for any SaaS company. As a result, the model was sending Bates leads that closed quickly, but dropped the service fast too. Fortunately, the solution was simple.

After learning about the problems with his model, Whittick instructed MadKudu's algorithm to analyze customers by lifetime value (LTV) and adjust the model to optimize for that. He also instructed it to analyze the accounts that churned quickly and score leads that looked like this negatively.

> **Example:** For Geckoboard, Digital Agencies were very likely to convert and the old scoring algorithm scored them highly. However, agencies had a 5x chance of churning after three months when the project they were working on ended.

At this point, the leads being sent to Bates were significantly better in aggregate than the leads he had previously been receiving. However, there were still false positives that would throw him off. While the overall stats on scored leads were looking great, the mistakes the model made hurt sales and marketing trust and were hard to accept. To combat this and make the qualification model close to perfect, Whittick had Bates start flagging any highly scored leads that made it through.

Through this process, they found that many of the bad leads that made it through were students (student@devbootcamp.com), fake signups (steve@apple.com), or more traditional companies that did not have the technology profile of a company who would likely use Geckoboard (tractors@acmefarmequipment.com). Whittick was then able to add specific, derived features to their scoring system to effectively filter these leads out and yet again improve the leads making it to Bates.

At this point, Geckoboard can predict 80% of their conversions from just 12% of their signups. By increasing sales efficiency with machine

learning, Whittick found a way to enable Bates to do the work an average sales team of five could typically handle.

From self-driving trucks to food delivery robots, this is the story of twenty-first-century business. Companies like Geckoboard are employing fewer people and creating more economic value than enterprises ten times their size. Leaders like Whittick are center stage in this revolutionary tale, figuring out how to optimize sales efficiency or conversion rates or any other metric given to them, just like the artificial intelligence they now employ. But of course, this has been happening over many years, even decades. The difference—and this cannot be overstated—is that Whittick doesn't have a PhD in applied math or computer science. The technology available to marketers today enables companies to generate twice the revenue with half the people.

CHAPTER 3

OUTBOUND SALES FUNDAMENTALS

Using data and Outbound Ops to scale





Kyle Williams and Ben Salzman

Co-founders of Dogpatch Advisors

In August 2009, I received an email that announced the launch of Google App Script, a product that introduced automation into the Google Apps product suite. Immediately, I sent a chat message to my coworker, Kyle Williams, a sales rep based out of Google's Austin office.

"What does this mean for us?" I asked him, pasting in a link to the product announcement.

Seconds later he responded, "Already on it."

Twenty years ago, Kyle wouldn't have been a sales rep. He's far too geeky (he's one of my closest friends, so I reserve the right to say these things). But my technical friend represents a new breed of salespeople that are changing the technology industry, boosting profits, and slashing jobs en masse.

Our chat conversation, and the subsequent tools that we built, opened my eyes to this reality.

Days after Google App Scripts launched, Kyle built a script that allowed the rows and columns from his spreadsheet to feed into Gmail. "We can automate our email campaigns now," he told me. At first I didn't get it; marketing automation tools like Marketo had introduced scalable email years ago, and our entire sales team at Google could send email from Salesforce. By the end of our conversation, though, I realized that Kyle had stumbled upon a small nuance that might have a huge impact on our sales funnel.

The problem with automation, he explained, was that it was immediately recognizable to the recipient as a marketing email. The lack of per-

sonalization meant that prospects felt no incentive to respond, let alone agree to a meeting. But automation has its obvious benefits: compared to the single-digit-per-hour emails that a human can send, automation enabled scale. With his new script, Kyle wanted to have his cake and eat it too. The first version made it possible to load data into a Google Sheet—for example, snippets like "I noticed that you use Lotus Notes"—and then pull that data into Gmail to send a semi-personalized email.

That fall, we began shopping the idea of a personalized outbound sales process around Google. While management was quietly supportive of the idea of scaling outreach, there was zero formal process or strategy. Most people didn't want to be associated with the idea, seeing cold email as spam, and pointing us to marketing leadership for approval.

The feedback we got from management was, "We like this. But we don't want to send generic, low-value messages and risk hurting our brand."

Some of our earliest campaigns were sent to contacts (inbound leads) already in our database and centered around events, webinars, and other marketing- and sales-driven activities. These tests were wildly successful. We emailed hundreds of potential customers at once, swapping in words, phrases, and sentences relevant to the prospect, all driven by data. At that time, our benchmark for replies from one-off outbound email hovered around 30%. These emails were performing at around 30%, but the volume was 10–100x our previous outbound efforts. All of this without the involvement of customer relationship management (CRM), marketing, or fluffy HTML templates.

But we ran into problems quickly. As we turned our targets from classic field marketing to new deal creation and cold outbound prospecting, we quickly identified two major roadblocks to ramping this up. First, most sales emails were focused solely on products or features. But any attempt we made to personalize them and make them less "salesy" were difficult to scale. In other words, we didn't have a repeatable framework for creating valuable content that we could scale without human input.

The second problem was related to data. At the time (and this is still largely the case), data providers like Jigsaw and Hoovers only provided contact data and rudimentary attribute data. In other words, we could buy massive lists of emails and their associated geography or industry, but nothing more.

What we really wanted was "proxy data."

We wanted to know useful information, like what types of technology a company used, so that we could change our messaging accordingly. For example, whether a company uses Microsoft Office or Lotus Notes is an indicator of their culture and the types of messages they might respond to.

Later that year, Kyle and I attended a sales training workshop at Google's Austin office. The outside trainer broke down the fundamentals of creating value through cold emails. With his method, he said, we could see reply rates at or above 50%. The only problem? He suggested writing the emails by hand.

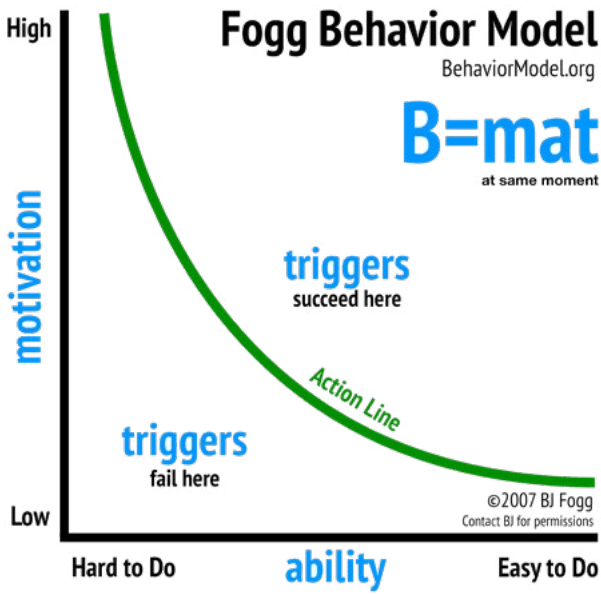
At the lunch break, I pulled Kyle aside. We immediately agreed this training was not only valuable, but transformative. If we could combine the sales trainer's strategy for personalized hooks and content with the sending tool, we could far exceed our quotas, and build a powerful tool in the process.

Later that afternoon, Kyle and I spent two hours in a conference room furiously whiteboarding ideas for personalized outbound at scale. The outcome of this meeting was simple agreement on the three things necessary for high-performing outbound at scale in this new world:

1. Large volumes of data about companies and people
2. High quality copy/content relevant for the prospect
3. Product that solves a real problem in a large market

At the end of our meeting, we realized we were missing a key piece. While we understood our messaging strategy, we didn't have a strategic way of thinking about the behavior we were aiming to drive. We needed a behavior framework to drive our messaging. Fortunately, right as we were looking for a framework, a scientist at Stanford's Persuasive Technology Lab was developing exactly what we needed.

In 2009, BJ Fogg published a paper that proved for any behavior to happen, sufficient motivation and ability must exist with the presence of a trigger. And those triggers would be different based on someone's situation. For example, a lower level manager might have high motivation (lots of pain/cost implications) but limited ability. A senior vice president (SVP) might have the ability, but not realize there's a problem. After reading this paper, Kyle and I asked ourselves: What if we could map motivation and ability for each prospect and allow that to inform our call-to-action (a.k.a. trigger type)?



In September 2009, with a new goal in mind, Kyle built a script that could ping email servers to identify which email security vendor they used. This allowed us to target messaging based on the type of IT team that was likely to purchase that technology. For example, IronPort customers tended to have concerns about reliability, while open source users tended to have a management issue and be concerned about total cost of ownership.

As we ramped up our campaigns, opportunities came quickly. We broke into new accounts, unstuck deals, and filled seats at events faster and easier than ever before. Some campaigns produced 70% open rates and 20% first meeting rates.

As a result, our small team was beating teams two and three times our size in terms of sales efficiency and pipeline creation.

But with each success came a new problem to solve. As we began scaling the program, we realized we were spending most of our time preparing lists, cleaning data, and manipulating data. We agreed that in order to scale, individual sales reps couldn't be stuck with these data janitor tasks. And thus, we invented the function of Outbound Operations to solve this problem. It wasn't until our next company that we formalized the processes and structures and found a way to create $30 million in pipeline without Google's cash and brand behind us.

How we invented the Outbound Operations function

Over the five years since we left Google, Kyle and I have helped dozens of companies build Outbound Operations into their organizations. In the rest of this chapter, we'll share what we learned and how you can build this function into your company.

In July 2013, I joined an enterprise software startup, Kahuna, as the fifth employee. I was brought in to manage one salesperson, and tasked

with building the team and helping close deals along the way. Our CEO also set a target of generating $1 million in new, qualified opportunities per quarter. My job was to replicate the process I helped pioneer at Google—with a fraction of the resources.

While confident that some version of Outbound Ops would fit, I was nervous about promising fast results without testing this outside of Google. When I broke down this process to our CEO, initially he wasn't sold on the idea of "email blasts." He was convinced that one-off, handcrafted emails would outperform any scalable solution. As we debated, we realized we had the same criteria for success: we wanted extreme personalization and we wanted lots of new opportunities. While most sales organizations achieve one of these two, few achieve both. Outbound Ops, I explained, is designed to avoid the false choice between personalization and scale.

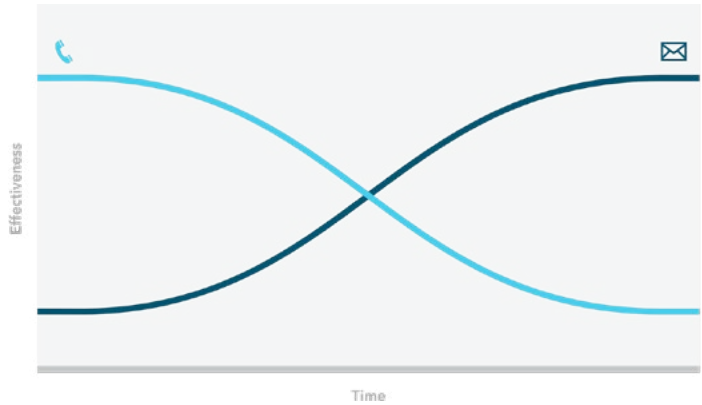
Our early results were phenomenal. We scheduled calls with Kayak, Yahoo, the *New York Times*, and other premier brands immediately, despite having no brand recognition. But after a couple of months, we decided that things were moving too slowly. Our sales team was not technical, and often got bogged down with the mundane task of cleaning data. We needed someone to run the entire outbound sales strategy, from data collection to opportunity creation. Without someone to manage the planning process, shepherd these data workstreams, and implement a rigorous testing and iteration framework, we'd never see the results I'd seen at Google. Fortunately, Kyle came to the rescue when he joined as the company's 10th employee.

After seeing the success of leveraging "proxy data," Kyle suggested that we identify which technologies our ideal customers were using in their mobile apps. If we could understand which apps had competitive software development kits (SDKs) installed, we could personalize messaging based on weaknesses for each of the incumbent vendors. To address this, he wrote a program with Charles Proxy to sniff his own Wi-Fi traffic as each app downloaded to his iPad. Four hours later, we had the data on which apps had installed which SDKs for the top 1,000 apps on our target list.

This data allowed us to run large outbound campaigns with precision, personalizing content to all 1,000 accounts, specifically referencing shortfalls in their existing solution, and customizing our content based on seniority and vertical. Our average reply rate was 50% and our average first meeting rate was 20%. That first quarter, we exceeded the $1 million target our CEO set by a factor of three.

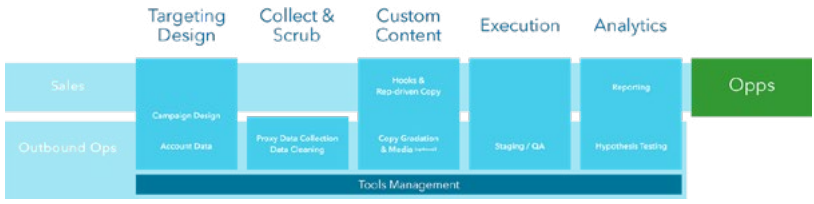
With a couple of wins behind us, the executive team gave us free reign to scale the team. At the time, the conventional wisdom said we should hire a handful of college graduates as sales development representatives (SDRs). By 2013, everyone in Silicon Valley was praising Aaron Ross's *Predictable Revenue*, which first introduced the concept a couple of years before. But something about this strategy didn't sit well with us. Why hire five people to do work that Kyle could automate in a few Red Bull-fueled nights? Unsure of our bold idea, we decided to take a hybrid approach. Kyle would build automation systems to supercharge the workflow of one or two reps.

When we initially wrote the job descriptions, we agreed to look for two traits in candidates. First, we would look for people with strong written communication skills. Whereas phone skills (cold calling and sales scripts) used to be the only prerequisite for an inside sales representative, we felt that the ability to inspire behavior through emails was far more important. Second, we wanted sales professionals who thought strategically about how to leverage data in campaigns—more Nate Silver-style quant, less country club schmoozer.



As we interviewed potential SDRs, we found plenty of good writers, but very few candidates knew their way around a spreadsheet. We'd ask them to rate their Excel skills; they'd respond "above average," and then we'd test them and watch as they struggled to use the filter function. It wasn't until we took the problem to our operations lead, Amy, that we realized we were fishing in the wrong waters.

After we explained our criteria to her, she asked something novel: what if we broke down the sales roles and had reps only focus on copy and strategy, and never touching raw data? She offered to help run data work streams so that reps could focus on customized copy and hooks in the emails. Whereas the average SDR might spend an entire day building a list of 100 sales leads, a strong operations hire can leverage scripts and other forms of automation to build a list ten times the size in half the time. Just as specialization of labor built our complex society, we were using it to generate deals like never before.



After a couple of quarters of strong results, we learned that the Outbound Operations team should be responsible for four activities: campaign planning, data cleansing, data and campaign review, and rigorous analytics.

Campaign planning

We learned very early on that the highest performing campaigns were the ones that evolved from carefully planned and discussed strategies. Rather than haphazardly trying to plan and execute each campaign on the same day, we set up a weekly campaign planning meeting, where

we'd tap the brains of our team's smartest marketers, salespeople, and executives.

While it wasn't evident initially, we learned that some of the behavioral triggers we so desperately wanted were locked up in the heads of our co-workers. Over lunch one day, a product marketer told me they had just finished A/B testing which website copy performs best by persona. She then asked why this data wouldn't directly inform our outbound messaging. She was right. After that, I decided that we needed to formalize this knowledge sharing.

Data cleansing

Another early insight was how often an email would go out that looked something like this: "Hi JONES, I noticed you work at Twitter, INC." These types of face-palm moments that outed our emails as impersonal were a result of sloppy data collection and cleansing.

After one of these emails went to our CEO's best friend, we decided that we needed a formal process to solve this problem. We decided that sales reps would no longer touch data. Instead, a dedicated Outbound Operations team member would clean each spreadsheet before it went to the next stage.

Data and campaign review

Even after we formalized the data cleansing process, we still had a few face-palm moments. In one instance, an SDR emailed a nonexecutive board member at Priceline and misspelled his name. From that, we learned that a formal campaign audit was necessary.

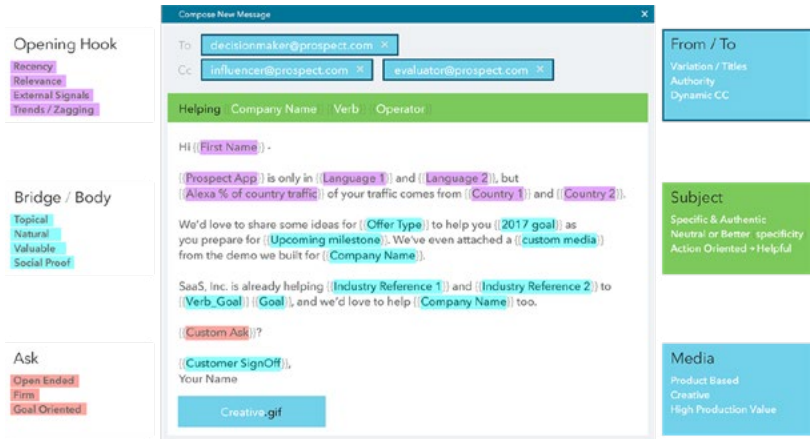
Before each message (email copy) went into an active workflow, we sent it to our marketing department's leadership for approval. This prevented silly mistakes and off-brand messaging by overeager salespeople.

Rigorous analytics

The most important thing we added to the Outbound Operations process was a formalized analytics system. After talking to some of our

salespeople, we learned that many of them lacked the statistical knowledge to run proper A/B tests. Many of them would try a subject line on ten emails and then ditch it for a proven one. In other cases, representatives would put all their eggs in one basket and fail to produce any pipeline, thanks to a bad test.

To achieve constantly improving results, we gave analytics responsibility to one employee instead of leaving salespeople to their own devices. Housing analytics in one place helped us smooth out revenue growth. It also gave us the ability to store our learnings in a central repository, as opposed to a half-dozen people's brains.



The results from our new Outbound Operations method were astounding. At the end of the first official year, the outbound sales team had generated $30 million in pipeline. That year, 75% of the revenue we booked came from our program.

In the time since we set up the Outbound Operations team at Kahuna, Kyle and I both left to start our own company. In doing so, we've had conversations with dozens of sales leaders about what we've learned. The reception ranges from skeptical to exuberant. But one thing is certain. Most teams—inspired by Aaron Ross's *Predictable Revenue*—have

implemented a sales process that is almost ten years old now. Meanwhile, the explosion of publicly available data and technology to make sense of it all has opened up new opportunities that the old playbook isn't fit to leverage. Real-time APIs, scripts, and SQL databases are the types of concepts that will make a salesperson's head spin. But they are the language of the day.

While we couldn't have predicted it back in 2009 at Google, what Kyle stumbled upon that August day foreshadowed a shift that will have huge implications. Whereas the specialization first taught in *Predictable Revenue* was an important predecessor of this shift, most sales leaders have missed the point: technology isn't changing the role of salespeople; it's making them increasingly obsolete. While the salespeople of my parent's generation schmoozed and boozed their clients over country club dinners, the quants replacing them are writing scripts and building widgets that generate 10 times the new business (revenue) with half the people.

CHAPTER 4

AUTOMATED OUTBOUND SALES

How Segment increased outbound email conversion by 200%



Michael Thomas

Founder of SimpleData

When Segment's VP of Growth, Guillaume Cabane, first began automating his company's outbound email campaigns, he expected two things to happen. First, because humans would no longer be a bottleneck, he expected the volume of activity to skyrocket. Second, because humans wouldn't be actively personalizing messages, he expected the conversion rates to go down.

For Cabane, who calls himself Segment's mad scientist, this was a simple optimization to measure. As long as the increase in volume was significant enough to outweigh the decreased conversion rate, he'd consider the automation a success. Four months later, when he ran his first report, he was astounded. Automation had risen into the top five sources of qualified opportunities at Segment. But he saw something else in the data that surprised him.

Rather than transition entirely to the new automated system, Cabane had decided to run an A/B test. He asked his coworker, Mark Miller, one of the company's SDRs, to continue emailing prospects as he had for the last six months so he could compare the results. When he looked at the conversion rates, he realized that the automated email system was actually outperforming Miller in volume and efficiency.

Not only was he able to email four times as many people as the manual process could in a week, but he was achieving a 55% open rate and 13% response rate from a cold email.

At SimpleData, many of our customers asked whether or not automation was worth the investment. We'd describe the trade-off most

companies make: volume usually goes up, but conversion rates go down; generally the output is higher in aggregate—the classic "quantity vs. quality" debate.





That's why Segment's results surprised me so much. They are the first company I've encountered that has actually been able to use automation to optimize both variables. Last month I sat down with the Segment team to take a look inside their most successful automation campaigns to date.

In terms of numbers, what's interesting in the case of Segment—and it's the case for many other B2B companies—is "while traffic is good, our conversion rate is pretty average," Cabane told us. "In the range of a few percent, as it is for many SaaS infrastructure companies."

One day, Cabane asked a fairly simple question: "What can I do to make them convert?" In each weekly growth meeting, he asks his team to suggest ideas, no matter how crazy they are. During that meeting, everyone tries to ignore limitations of cost or technological constraints. Then at the end of the meeting, everyone explores the possibility of the best ideas. One idea proposed in a meeting last year, by coworker Zeb Hermann, was the ability to email visitors that hadn't converted. A couple of weeks later, Cabane built the following workflow.

1. Call the Clearbit Reveal API and request the IP address of the visitor, which is then converted into the domain of the visiting company.

2. De-duplicate results against customer and active lead lists to suppress accounts that should not be emailed.

3. Call the MadKudu API to add a lead score to the domain to know whether it is a viable candidate for outreach.

4. Call Clearbit's Prospector API, which finds the key people at that domain, their roles, and their email addresses.

5. Call the Clearbit Enrichment API, which adds enriched person data to each prospect for email personalization.

6. Send those new prospects through the lead scoring process again.

7. Send highly personalized emails through your email platform of choice.

After running the workflow a few times, Cabane realized that a high percentage of site visitors were bad prospecting targets.

He broke them down into the following groups:

Crawlers and Bots - Search engines and other spiders constantly index sites looking for changes. This is great for your SEO, but useless for outbound emailing.

Employees - Employees visit their own site daily.

Partners - Again, they are likely to be visiting extensively.

Your Customers - Even if they have already converted, they'll still use your home page as a starting point to get to your dashboard.

The last group was one that Cabane wanted to avoid at all costs. "When I do outbound, I'd rather be safe than sorry. It's usually okay to have some margin of error in growth experiments; in outbound, it is not okay to be wrong," Cabane told me.

In order to avoid emailing customers, Cabane built extensive filtering into the workflow. Instead of wiring together five APIs, Segment used a customer data platform called Hull. This collected data from their SQL data warehouse (with all their signups), Salesforce (with all their customers), and the list of companies visiting the website (from Reveal) and merged it into one unified customer profile: a "master profile."

With everything on one profile, Cabane could then spot and flag profiles from employees and partners to exclude. But to err on the side of caution, the job would also flag a domain from the list if anyone from that company had at any time signed up for Segment's free product. In other words, if an engineer at IBM had signed up for an account two years ago, no one from IBM would ever be emailed through this automation workflow. In total, between 60% and 80% of the domains were filtered out. The remaining flag-free prospects were then added to a segment to sync to Customer.io into their automated outbound email cadence.

Writing programmatic emails

Once he had a new list of leads—a list that, unlike most, was refreshed each day—Cabane enlisted one of the company's SDRs, Mark Miller, to help him find a way to write email templates that felt personal and delivered value. (See a video of the Segment team in action.[5])

One year prior, Miller had joined the company as one of its first sales reps. Being the youngest member of the team, he was given the unglamorous task of building lists and emailing prospects. Many of his days were spent searching Google and entering data into spreadsheets. "I vividly remember, before I switched over to Customer.io, there was an afternoon where I was trying to send my own outbound emails. I was

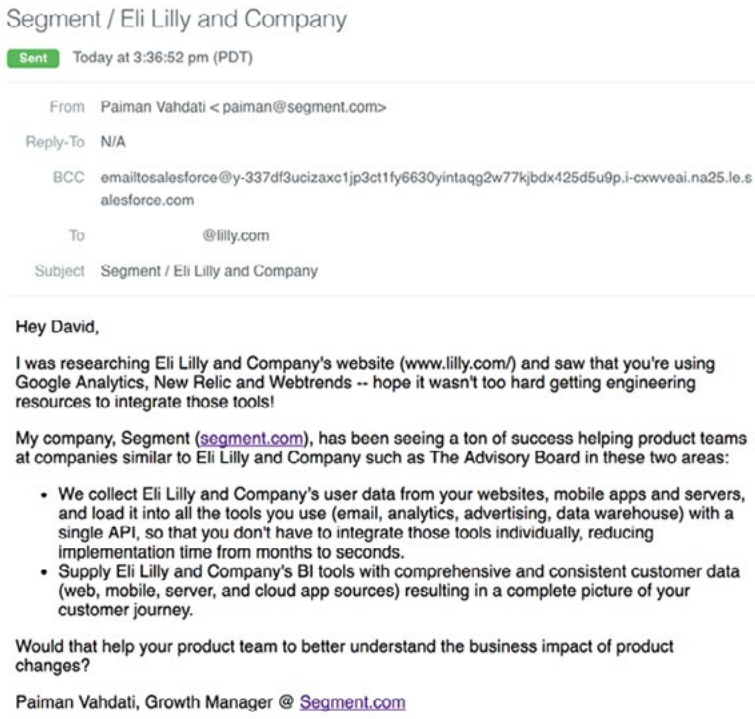
sitting on a couch, and I had a list of companies, and I went to each one and would check out Ghostery to find out what technologies they were using. I would then go into the email that I was trying to send in Outreach, I would type in those technologies, and then I would press Send."

It was an entire day's work.

He ran into the same dilemma all SDRs face. "There are two huge challenges of doing outbound as an SDR, as a single person," he told us. "The first is doing personalizations at a fast rate, and two, sending enough emails that you'll get enough replies to then allow you to set up enough opportunities to then hit your quota." Cabane's new automation workflows promised to solve both of those problems.

Cabane told Miller that if implemented correctly, automation could help speed things up exponentially, but this type would also add personalization. By adding data such as a lead's role or industry, the Reveal loop could take advantage of a wealth of data and pull it into Customer.io's liquid logic. Then they could use if-then logic to programmatically build emails.

This allowed him to do personalization at scale. Personalization is how you get conversion. If the email doesn't help the lead directly, then they have no need to respond. Therefore, Miller built the automated email just as he would as an SDR reaching out to a single individual he had researched thoroughly. Here it is, in its entirety, broken down by section below.



Segment / Eli Lilly and Company

Sent Today at 3:36:52 pm (PDT)

From Paiman Vahdati < paiman@segment.com>

Reply-To N/A

BCC emailtosalesforce@y-337df3ucizaxc1jp3ct1fy6630yintaqg2w77kjbdx425d5u9p.i-cxwveai.na25.le.salesforce.com

To @lilly.com

Subject Segment / Eli Lilly and Company

Hey David,

I was researching Eli Lilly and Company's website (www.lilly.com/) and saw that you're using Google Analytics, New Relic and Webtrends -- hope it wasn't too hard getting engineering resources to integrate those tools!

My company, Segment (segment.com), has been seeing a ton of success helping product teams at companies similar to Eli Lilly and Company such as The Advisory Board in these two areas:

- We collect Eli Lilly and Company's user data from your websites, mobile apps and servers, and load it into all the tools you use (email, analytics, advertising, data warehouse) with a single API, so that you don't have to integrate those tools individually, reducing implementation time from months to seconds.
- Supply Eli Lilly and Company's BI tools with comprehensive and consistent customer data (web, mobile, server, and cloud app sources) resulting in a complete picture of your customer journey.

Would that help your product team to better understand the business impact of product changes?

Paiman Vahdati, Growth Manager @ Segment.com

Introduction: Why are we contacting you?

> I was researching Eli Lilly and Company's website and saw that you're using Google Analytics, New Relic, and Webtrends -- hope it wasn't too hard getting engineering resources to integrate those tools!

There are three elements of personalization in this email. First, and most obviously, they reference the company's name. This is standard practice among the best sales development teams, but it's still important to note.

Second, they use enriched data to reference what technologies the company is using. This shows that the SDR has done research on their company.

And third, they use the lead's role (job title) to reference a pain point they may have experienced recently. In this example, they email a product manager and reference the pain of getting engineering resources to help with data infrastructure problems.

Value proposition: Targeted message

> My company, Segment, has been seeing a ton of success helping product teams at companies similar to Eli Lilly and Company such as The Advisory Board in these two areas:
>
> 1. We collect Eli Lilly and Company user data from your websites, mobile apps and servers, and load it into all the tools you use (email, analytics, advertising, data warehouse) with a single API, so that you don't have to integrate those tools individually, reducing implementation time from months to seconds.
>
> 2. Supply Eli Lilly and Company's BI tools with comprehensive and consistent customer data (web, mobile, server, and cloud app sources) resulting in a complete picture of your customer journey.

Miller calls this part "the pitch." It's the point in the email that he tells the prospect about his company's value proposition. Once again, automation allows for personalization at scale.

First, by enriching the data to include a lead's industry, Segment is able to reference similar companies in the industry. This social proof helps drive conversions, as the prospect is likely to recognize the company.

Second, by using the role data again, he can reference a pain point that a product manager has likely faced. This helps Miller pitch the right features to the right prospects.

Call-to-action: Making an ask

> Would that help your product team to better understand the business impact of product changes?

By asking a question at the end of the email, the SDR has started a conversation. From there, it is their job to communicate with the prospect and try to convince them to see a demo.

There isn't a long drip campaign to each lead, either. The sequence is just two emails. If there is no reply, they send one more email that includes a targeted case study based on their industry, determined automatically with their enriched data.

> Wanted to follow up with an example of how a customer of ours uses Segment, which may come across better than an email from a stranger :) Here's how LogMeIn uses Segment to unify data across the diverse toolsets that each team uses: https://segment.com/customers/logmein/story. Does Segment seem useful to you or anyone else on your team?

In this case, they would send a case study with a product manager at a technology company. Once again, this personalization ensures that Segment shows the most relevant value proposition to a lead. If the lead doesn't reply, they aren't contacted again.

So far the emails have been radically successful. Compared to industry averages of 30%, Miller's automated emails have a 55% open rate. "It's extremely efficient," says Cabane.

After only four months of running the program, 22% of all qualified opportunities come to Segment through the Reveal loop. In addition, the revenue from this channel is predictable because it is correlated to their consistent website traffic.

By being first to market with this technique, Cabane also reaps the benefit of novelty. “When you’re first to market, you have higher conversion rates, better conversion rates, and better open rates. You’re doing something new. In the case of the Reveal loop, it’s exactly that. I was among the first people to do something, and so my performance is amazing.”

CHAPTER 5

SALES FORECASTING

The art and science
of predictive sales forecasting



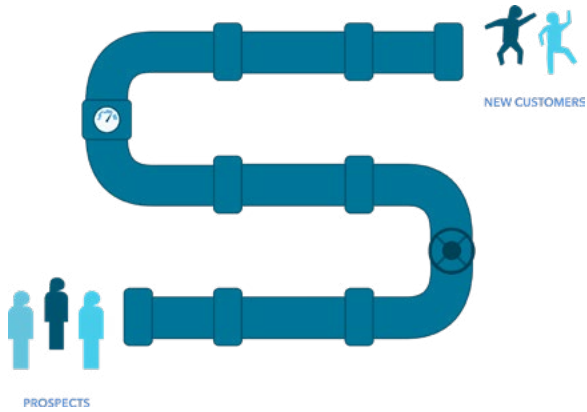
Neil Ryland

CRO at Peakon

In the fall of 2010, the CEO of Huddle—my boss at the time—gathered our senior team in a conference room. The sun was beating down through the windows at 8 a.m. in our Regus office in San Francisco. A couple of minutes after we sat down, our UK office peers connected via video conference. The gloom of the London weather, and the atmosphere in that conference room, soon filled our office as well. "We missed our projected revenue numbers by 40%," our CEO said. Of course, many of us already knew this. We'd been feeling the stress of being behind our target for the last month. The question on all of our minds was, how could we get our sales forecast so wrong?

I've been a part of enough sales teams over the last 10 years of my career to know what it feels like to miss a target. The truth is that sales forecasting with perfect accuracy is impossible. Forecasting is by its very nature an exercise in predicting the future. It's always a best guess. But this quarter was by far the worst. Not only had we missed the target by 40%, we'd also made a number of important business planning decisions based on hitting our number and the resulting new business. We'd opened offices in the United States, hired four senior salespeople, and invested in a large headquarters office in London. The list went on. In that meeting, we all had to own up to our forecasting mistakes and determine how we'd improve going forward.

Anyone who's been in sales long enough will relate to this story. They'll relate to the balance between optimism and accuracy—that is, the tendency for nearly everyone in a sales organization to inflate numbers to please their higher-ups. They'll know how difficult it is to combine both art and science into a single number that affects million-dollar deci-



sions. But as hard as predicting the future may be, it's possible to make accurate forecasts using data and a disciplined process. Below, I've tried to share my story—including both the stumbles and successes—in hopes of helping you develop a sales forecasting technique you can rely on.

Introduction to the Deal Forecasting methodology

I remember three months into my first job in sales, my sales director from Capscan, Dave Mead, stopped me in the kitchen area. As I picked up my coffee to head back to my desk, he said, "Ryland, how much revenue are you going to book this month?" My heart racing a bit, I proceeded to reel out everything I was working on and planning to work on. He cut me off and said, "A number will do just fine."

I ended up taking a rough guess and walked off feeling like an idiot. It felt like such a simple question and yet I struggled so much to answer it.

From that point on, I have always taken a keen interest in the art of forecasting and its importance.

Later that week I sat down with Dave and talked through how I could improve my forecasts. I wanted to be consistent in my revenue numbers

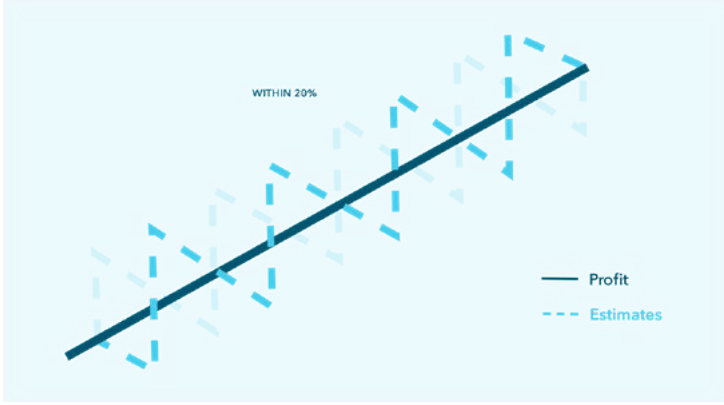
and forecast with accuracy. That is where I first learned about the simple but highly skilled Deal Forecasting methodology.

The Deal Forecasting methodology is as simple as listing the deals you expect to close and tallying the revenue you expect to earn from them. It requires minimal data analysis and is more about trusting the sales qualification process than any complicated science. In a smaller organization (20 reps or less), it can be very effective because it doesn't add much overhead or require a sales operations hire to implement.

New+Addon+Uplift Forecast (Open + Closed Won)

	Current Month
AE Forecast	Total New+Addon+Uplift
Commit	€217K
Upside	€145K
Pipeline	€20K
Total	**€382K**

The Deal Forecasting methodology can also work well when deal volumes are low. At one of the companies I advise, IncuBus London, this model works well. They sell a B2B service to larger organizations and only a handful of sales reps, so the volume is low enough that they've been able to stay within the magic 20% number (the accuracy that most CFOs and boards ask for). In other words, they hit their forecasts each month.



Having a conversation with reps each week about their pipeline also helps them "pick their battles." The sales manager sits down with each rep to review deals that have the best chance of closing that month based on product-market fit, decision-maker access, and geography. As a result, everyone is on the same page as to which opportunities they should be focusing on.

However, this model no longer works when you want to forecast beyond your average sales cycle. For example, if your sales cycle is 30 days and you are using a named deal forecast, it is impossible to forecast the following month, as the opportunities will not have even been created.

In my Capscan days, as a short-term, focused sales rep, my priority was always the last working day of the month, but as I moved into more senior roles, I realized that having a 30-day vision was not going to give me the information I needed to be successful in my job. How could I plan how many people to hire, set sales quotas, or tell the marketing team how many leads we required, let alone convince venture capitalists to invest millions of dollars into our business, if I couldn't see further than a month ahead? Capscan went on to be acquired by GB Group for multiple millions and that decision was based on more than the next 30 days of deals, I assure you.

Deal Forecasting Method

Pros

- Provides managers with deep insights into the deals that matter
- Doesn't require historical data

Cons

- High risk due to dependency of sales reps' qualification diligence
- Impossible to predict beyond your average sales cycle time

Improving Sales Forecasting as Sales Cycles Grow

When I moved on from Capscan to join Huddle at the beginning of their journey, we initially used the Deal Forecasting technique. But as our sales cycle time grew and we began to close bigger deals, we quickly ran into its limitations. I learned within the first few months that this model was only accurate at predicting deals at the bottom of the pipeline.

Our first step to improvement was working with our Commercial Director, Charlie Blake-Thomas, to improve our deal forecasting through a much stronger qualification process. Charlie had a keen eye for detail and I quickly noticed a pattern that he followed in each sales meeting. Even when the most amiable of CMOs began speaking during a meeting, he had a way of cutting through the fluff and getting them to agree to exact numbers (like the leads they would deliver our sales team).

Charlie taught me in-depth sales techniques using the BANT qualification method and the CustomerCentric approach to deal management. By following a standardized process across the team, we were able to develop our first weighted forecasts, or an Opportunity Stage Forecasting technique, thus introducing me to the next evolution of sales forecasting I'd use in my career.

We used past sales data from how many leads had turned into closed-won opportunities and collaborated with the entire sales team to understand the typical steps required to close a deal. We then built a clear and concise pipeline process with a prediction on revenue we would close from each stage. Our model looked something like the chart below, which is fairly standard for a SaaS SMB/mid-market process.

STAGE	REQUIREMENT	PERCENTAGE
LEAD	CORRECT DETAILS AND DEMO BOOKED	10%
QUALIFIED	B. A. N. T. QUALIFIED	25%
VALIDATED	PRODUCT DEMO COMPLETED	50%
VETTED	PILOT COMPLETE AND ALL DECISION MAKERS ENGAGED	75%
PROCUREMENT	FINAL NEGOTIATIONS	90%

Based on the example above, our reps could have up to to 10 working deals a month. At the time, we had an average order value of $10,000 and knew that our win rate from opportunities being created, where a clear need and timeframe had been expressed, was 25%. Therefore, if we ensured that reps qualified consistently to that criteria, we would be able to forecast revenue from opportunities created. The real-life maths for each rep was such: we knew that we could hire more reps to hit a $25,000 monthly target based on marketing delivering the agreed level of leads to create four qualified opportunities a month. In addition, we were then able to make informed decisions by monitoring conversion rates at each stage of the pipeline and adjusting the weights accordingly.

A couple of months after developing this forecasting method, we used it to inform a decision to invest in an inside sales team in the US. When I moved to San Francisco to help build that team with Lesley Young—one of the brightest minds in the Valley—I saw firsthand how weekly

one-on-one meetings and an Excel spreadsheet could lead to accurate forecasts.

The model that we developed at Huddle was perfect early on because it enabled us to predict with relative accuracy without developing a complex predictive model. However, as we scaled, it became difficult and time intensive to update the probabilities associated with each stage. This led to decreasing forecast accuracy and increasing administrative cost. As we hired more reps, it also became difficult to enforce proper Salesforce usage. Close dates were wildly inaccurate, or missing altogether in some cases.

As we started to move up market to larger deals with longer sales cycles, the maths behind our 25% close rate was no longer standing up over a 30-day period. I learned that for all the allure that larger deals and bigger customer logos bring, they make it difficult to forecast with accuracy. In addition, we started to see a bigger proportion of our business coming from government. Larger enterprises and government organizations have far more complex buying cycles, and therefore we had a number of deals significantly above the AOV that would get stuck in our sales process. So although at a local-manager level this could be managed through deal reviews, at a more senior level, this stage of forecasting (weighted forecasting process) was causing a lot of inaccuracies for board reports.

Weighted Forecast Model

Pros

- Great for getting a process in place for forecasting when you have reps working 8–12 deals a month
- Creates focus and structure for pipeline management
- Drives accuracy in forecasting, as reps are forced to follow a process in order to progress a deal

Cons

- High admin required to constantly check that sales stage lines are accurate
- At VP or director level, the challenges of sales cycle and ancient opportunities can skew this when you are not close to every deal in the business

Category Forecasting and Moving Toward a Predictive Model

Toward the end of my time at Huddle, we implemented what is known as the "Category Forecasting technique." This method was simple enough that we didn't need to have a full-time sales operations hire dedicated to managing it, but it was more nuanced than the Weighted Forecast because it accounted for how a rep was feeling about a given deal. In other words, it considered the art as much as the science.

We added a new field in Salesforce, which was mandated, and all reps had to mark the status of each deal.

Commit
Verbal confirmation to sign by the close date the sales rep has marked in the CRM application

Upside
The sales opportunity has the propensity to close but not all decision makers have been engaged

Pipeline
No clear understanding of the decision-making process or understanding of nonaction (e.g., RFP process)

By overlaying this data over the weighted forecast call, we were able to see what the science was telling us (based on historical conversion

rate data) and combine it with the art from the reps' forecast category (based on how a rep was feeling about a deal).

Because reps had to follow a strict process each time they created or updated an opportunity, we found that the pipeline stages were more accurate. In addition, it helped us to spot reps who would keep deals at an early stage of the pipeline (e.g., commit and close deals from 25% to close in a day) and coach them on using pipeline stages correctly.

When I got to Peakon, where I am currently the Chief Revenue Officer, I immediately found a team of people excited to use data to improve the Category Forecasting method even more. No person was more excited about using data in the sales process than our co-founder Dan Rogers (who I also learned can drink like a rugby player).

One afternoon last year, we sat down with a six-pack of beer to determine how we could improve our forecasting before going out for our next round of funding. As a company that sells an analytics product, we knew that investors would want to see a more nuanced forecasting model. So we looked at the flow of leads to determine what makes a lead "good" and how we might automate part of the qualification process to start projecting revenue based on how many leads we had.

We quickly spotted trends based on geography, company size, and job title. Then we updated our forecast model to account for the changes. We no longer had a simple probability attached to each lead stage; instead, each lead was given a unique probability based on historical conversion data of similar leads. Rather than saying that a website signup had a 10% chance of closing, we evaluated where the lead had visited from (lead source) and what country their IP address was located in (geography) and who from the company had signed up (job title).

If a lead looked exactly like 10 others that had closed quickly, then it might get a 25% probability.

Data points that influence our forecast

1. Employee's personal win rates
2. Lead score
3. Opportunity age
4. Vertical
5. Company size

Of course, no company can perfectly predict how much revenue they will make in a given month, quarter, or fiscal year, just like no meteorologist can perfectly predict the weather. There are simply too many variables to consider. Every business is affected by both macroeconomic and microeconomic factors. Interest rates in Japan affect the price of oil in Kuwait, which in turn affects how much money your would-be perfect customer can invest in analytics software. In the age of Software-as-a-Service (SaaS), your competitors' products can, and do, change every day. So perfection is not the goal. Instead, forecasting is about getting close enough to make informed decisions.

Yet, "as close as possible" is no exact figure either. Ultimately every business must determine how much time and resources they can devote to forecasting, as opposed to the hundreds of other activities they could be doing. And in my experience, the best approach is the same as it is in just about every other facet of sales. Start swimming. Learn as you go. Get better every day, and remember that if you don't track it, you can't ever measure it, let alone predict it!

CHAPTER 6

FINDING SALES TRIGGERS

How DigitalOcean closes more enterprise deals with automated sales triggers



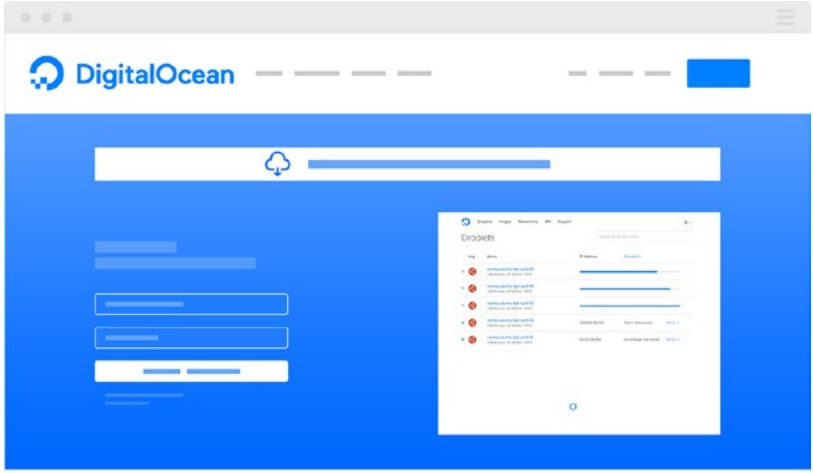
Ed Fry

Growth at Hull

When Emmanuelle Skala started her new job as the VP of Sales and Customer Success at DigitalOcean, she couldn't help but remark at what a rare company she had joined. For more than a decade she'd been a part of fast-growth companies. But all of them had relied on large sales teams and expensive marketing budgets to fuel that growth. The five-year-old cloud computing company was an exception. They'd landed more than 500,000 customers, hired hundreds of employees, and raised $300 million in funding—all without a sales team.

Spend enough time in the technology industry and you'll realize what an echo chamber it is. The same buzzwords are thrown around day after day at conferences hosted in the same event centers that are populated by people that all look more or less the same. Naturally, companies employ many of the same strategies for growth, and this makes DigitalOcean's success all the more surprising. But it didn't prevent the team from eventually having to hire a sales team to "go bring in new logos." At some point, even companies that do manage to build a Field of Dreams have to go out and find new customers on the path to IPO.

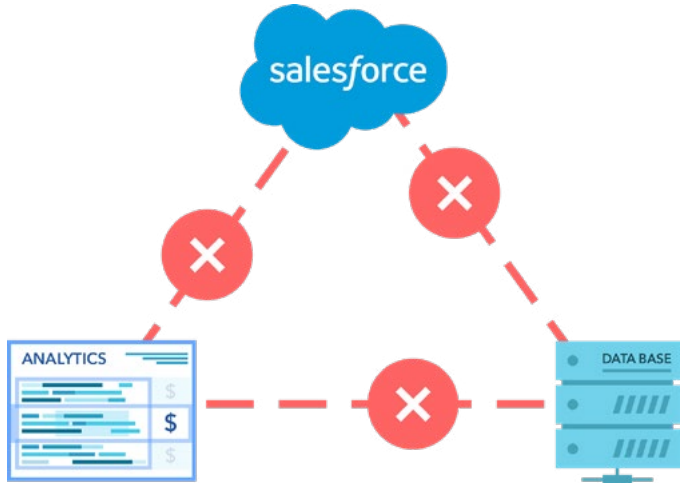
With three previous VP of Sales positions under her belt, Skala was well acquainted with the challenges of building and scaling a team. Unlike the previous companies where she'd worked, many of DigitalOcean's customers were self-service, and primarily engineers. It became immediately clear that hiring a team of SDRs and setting them loose was going to hurt their brand and their relationship with potential high-value prospects.



However, DigitalOcean did have a unique advantage. When Skala joined, there were millions of members in their community and more than 500,000 active customers. Instead of building out an SDR team and executing an outbound strategy, Emmanuelle decided to look at capturing more value from the existing signups and potential customers.

“We said we shouldn’t be going outbound. We have incredible brand equity among developers and just need to educate them a bit more,” she said.

She told the sales team to switch their focus from cold-calling prospects to reaching out to existing customers who trusted DigitalOcean. They ran into problems with this too, though. Like most B2B companies we work with at Hull, their customer data was siloed and scattered. The sales team wasn’t sure who at a given company they should be reaching out to and what made a qualified lead. And all the valuable data was stuck in analytics tools and backend databases, not packaged up and synced to the CRMs that the sales team used.



It was at this point that Emmanuelle linked up with DigitalOcean's Growth and Demand Generation Manager, David Dorman. To determine which customers and leads to contact first, Emmanuelle suggested they take a two-pronged approach. First, she tasked DigitalOcean's internal team of data scientists with the job of doing some sales analysis and finding correlations and actions among existing customers that could indicate growth. At the same time, she also tasked an outside team of consultants with the same data and challenge. Her goal was to verify and test different models against each other to find the winning combination. Naturally, the two teams developed different hypotheses as to what might drive growth within accounts.

The external team of consultants focused on team size. Similar to Dropbox and their move into enterprise, the external consultants used account data together with third-party enrichment data (like Clearbit) to look for customers who worked at large companies with many employees. Thanks to the self-service product, there would sometimes be multiple accounts from the same company, too.

DigitalOcean's internal team started with a similar approach. Using enrichment data, they could append business data to each email and

account (even with many Gmail addresses) to understand whether they were business users. For those that appeared to be organizations, they took it further and focused on product usage.

"We used Segment, and a suite of connected tools, to better understand patterns of usage. A lot of it was very specific to DigitalOcean-type usage," she told me.

Their team could look at product usage, including what products were being used, their capabilities, any team features, and billing settings. Furthermore, they could see the volume of usage, including any significant spikes, which tended to indicate testing.

By using firmographic data, to ensure that they were talking to real businesses, together with the product usage data, which indicated a "real-life" intent to build something with high value (like team features and testing a load balancer), they could cue in their sales team to the right accounts. Combining the data led to a much closer correlation with higher value accounts than the external consultants' model.

However, statistical models don't close deals. To turn the model into a useful tool for sales, they needed to identify sales triggers in product usage that were tracked and recorded in their analytics software. Then they needed to package up those events to sync reliably in a format that was native and useful to their CRM. Without this, the data would be inaccessible to sales.

It's not good enough to say "let's have a call"

Skala knew from previous experience that content was the key to closing deals. "Once you have all this rich, contextual data, your outreach is really only as good as your content. It's not okay just to say, 'We noticed X, let's have a call,' especially with the audience we're serving," she said.

DigitalOcean's ideal customer is a developer, and that audience tends to be sales-averse. "Because we sell to developers and because we're

self-service, there's always been a high reluctance to communicate with customers other than about things that are absolutely necessary," she said.

The challenge was aligning messaging and outreach with the customer's journey and what they wanted to hear. The generic, "send-to-all" sales and marketing campaigns were not tolerated by their technical audience. That's why they turned to technical how-to content, which was found to be useful when the outreach happened at the right time.

"It's generally pretty technical how-to content that works well with our audience. It's more like, 'Hey, we noticed ABC about your usage, here's a great piece of content on how to better implement a certain technology or integrate with something or take advantage or set up something.' It's not marketing content," she said.

By using the sales triggers that their data science team uncovered—and working with Dorman's team to sync them to the CRM—the sales team was able to send the right content at the perfect time. As a result, it positioned the sales team and DigitalOcean as a helpful resource, not an annoying burden in their customer's inbox. They knew who to talk to, what to say, and when to say it.

"Now that we have been mining all this data and are starting to correlate it, the appetite among the team to actually do this has changed dramatically because there's a realization that the messages can be very targeted and can be very helpful and don't have to appear to be marketing-y or sales-y," she said.

Testing the model with a sales A/B test

As was the case when they first created their model, Emmanuelle wanted to learn whether the new sales outreach was truly effective. To do this, she created an A/B test. Of the eligible users, 80% were divided up among the sales team for outreach and 20% were held back as a control

group. The inside sales team then used these pre-prepared lists to send their messages.

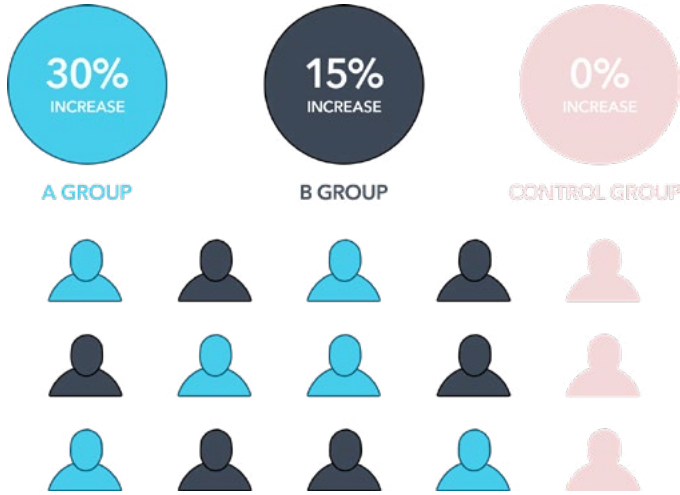


At the end of the first quarter of 2017, she and the data team compared the three cohorts:

1. Outreach, with responses to the outreach
2. Outreach, without responses
3. Control group, without any outreach

What they found was that the accounts who received proactive outreach to offer a consultative conversation were twice as likely to expand their usage above a specified threshold. In turn, setting these customers up for success led to a 30% increase in account value compared to those who received no outreach.

Curiously, the middle cohort who received outreach, but did not respond, also saw an increase in usage. 50% of those accounts found an increase in product usage too.



This put a dollar value on DigitalOcean's "serendipity," and also proved value beyond just booking meetings. Sales activity could drive significant revenue growth outside of sales calls.

"Sometimes salespeople are too focused on the ask being a meeting," she said. Later she told me that even the leads who didn't respond or take a call saw a significant increase in spend. This encouraged their sales team to be more and more engaged with the outbound sales effort. By using specific trigger events from product usage as a conversation opener—paired with the right content—they were given the perfect opportunity to do that.

"You don't have to just always go for the meeting, as long as you have really interesting, relevant content to send, and then it will improve results. Partner with the content team and make sure it's an integrative strategy."

With this strategy firing, they implemented a similar strategy with their account management team. Previously, DigitalOcean only contacted an account if there was an open ticket. Seeing the success of their

proactive outreach strategy on the sales team, they decided to run the same test. Once again, account managers were assigned to customers to proactively reach out, with a proportion held back as a control group.

DigitalOcean found that revenue churn among accounts was half that of the test group (with proactive outreach) compared to no outreach at all.

"It just put the nail in the coffin and everybody realized that, 'Okay, yeah, this works' as long as you're contextual and as long as you're relevant. That spurred us on into diving deeper and finding more and more triggers," she said.

Sales can continue to grow by finding new triggers

With more specific triggers, sales had more conversation openers with which to build relationships and grow accounts.

If a lead looked exactly like 10 others that had closed quickly, then it might get a 25% probability.

"Since then, we've launched, I think, three more PQL types and even just yesterday at our Hack-A-Thon, some of the data science team has been working on this and just hacked together a whole new series of data that now we can start using," she said.

More importantly, DigitalOcean's team is looking at using the triggers and highly relevant messages elsewhere to deliver helpful content to the right people at the right time. They know WHO they're talking to, WHAT to say, WHEN to say it, and now WHERE they can message. "At this point, everyone's pretty convinced that it works and we're even finding other use cases like in-app messaging and chat to create new sales opportunities," she said.

Compared to the traditional spray and pray method, DigitalOcean found a way to increase revenue without annoying the community that helped them get to where they are. Rather than employing more people and seeing diminishing returns with each hire, they are orchestrating their existing data to work with a small sales team to speed up their revenue growth.

“The appetite for even doing this kind of stuff has changed. Now there’s no longer any pushback to doing this. Overall, the company is much more aligned toward even allowing for these things to happen because of the fact that they’re contextual and relevant.”

CHAPTER 7

SALES COMPENSATION

How HubSpot accelerated growth by optimizing sales compensation



David Skok

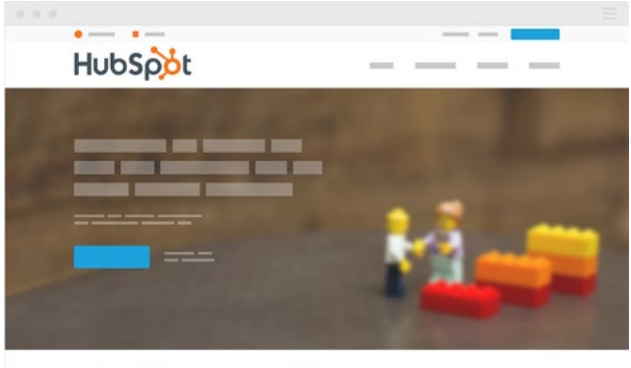
VC at Matrix Partners

and the author of forEntrepreneurs

When I started my first company many years ago, software sales was radically different than the SaaS landscape that so many readers of this book know. Sales cycles were long, and whether we liked it or not, the customer was stuck with what they had as soon as the papers were signed. The SaaS model changed everything. For founders and sales managers, this shift has required changes to everything from the way a company sells to the manner in which it delivers the product. But, one of the most important and overlooked changes has been the way sales-people are motivated and compensated.

Unlike traditional software sales, the job of a sales team no longer ends when a new customer signs a contract. Instead, it is crucial to retain customers over many years, as that is how you maximize your revenues.

In this chapter, I'll share the story of how HubSpot, one of our portfolio companies, experimented with their compensation plan to accelerate growth by reducing churn. I've pulled out three compensation plans at three different phases of the company in hopes of sharing some insight into how a plan should change over time.



The hunting plan: $300K to $3M ARR

When HubSpot's VP of Sales, Mark Roberge, developed the company's first compensation plan, he had one goal in mind: simplicity. Over the course of the first year, he and the company's founders had landed 100 customers. In doing so, they grew to $300,000 per year in annual recurring revenue (ARR). After that first year, Roberge decided it was time to step out of the closing role and into a sales management role. He knew that it was essential that his first plan was easily understood by reps to avoid confusion. But he also wanted the plan to incentivize rapid customer acquisition so the company could reach product-market fit as quickly as possible.

As Roberge writes in his book, The Sales Acceleration Formula,[6] "In our first plan, we paid each salesperson $2 upfront for every $1 of monthly recurring revenue they brought on. For example, if a salesperson closed a customer on a $500 per month subscription, the salesperson earned a one-time $1,000 sales commission for that deal."

To protect the company, reps were told that if a customer cancelled within four months, they would not receive the commission. In practice, this "claw-back," as it's known in the industry, was very simple. If the customer churned within the first four months, HubSpot took the entire commission back (deducted from the salesperson's next month of earned commissions).



This hunting-oriented plan, as Roberge likes to call it, was incredibly effective. Within six months, the company acquired 900 customers and

scaled from $300K to $3M in ARR. But of course, nothing about scaling a sales team is easy, and sales compensation was no exception.

Immediately after implementing the plan, customer churn exploded. Reps were aggressively signing customers without properly evaluating whether or not they were a good fit. In addition, they frequently over-sold the product and set the wrong expectations. In order to solve this problem, Roberge and the rest of the executive team sat down to develop a more robust plan.

Reducing churn by 70% by incentivizing good behavior

One of the things that makes Roberge so effective at what he does is his discipline and love of data. In order to understand why churn had increased so much, he and his team made no assumptions. Instead they looked into the numbers.

At the time, every HubSpot customer was assigned a post-sale consultant to help them understand the product and get the most out of their investment. In Roberge's mind, customer success was tied to how well the product was implemented. So the first theory was that the post-sale consultants were the problem. They thought solving the churn problem was as simple as cutting the bad consultants and learning from the good ones. But when they analyzed the churn level by consultant, they found something surprising: there was no difference between them.

Next, they analyzed the churn rate by sales rep. If the results from the first analysis were surprising, these were shocking. They found that across the 10-person sales organization, there was a tenfold difference between the salesperson with the lowest churn rate and the salesperson with the highest. HubSpot didn't have an onboarding problem. They had a sales behavior problem.

After analyzing the data, Roberge shared the results with the entire sales team. In a conference room one day, he educated each of the reps on the importance of retaining customers. He showed them what revenue projections (and the value of their equity) looked like with 3% churn, 5%, 10%, and so on. Then he told them that he would be changing the comp plan next quarter to align churn performance with the sales goals.

The next quarter, Roberge stack-ranked each salesperson based on their churn performance. In a team all-hands meeting, he shared the results and announced the new compensation plan.

In the new compensation plan, the top-performing quartile of reps would receive $4 for every $1 of monthly revenue booked. He told these reps "congratulations" and told them to keep up the good work. The goal here was to "reward the winners," a key strategy in any compensation plan.

The next quartile would receive $3 for every dollar of recurring revenue they booked. Still he was rewarding good performance and incentivizing good behavior, since each of these reps would be receiving a 50% bump in their commission rate.

For the salespeople in the third quartile, he told them that they would be receiving the same $2 per dollar of monthly revenue.

For salespeople in the last quartile, he cut their commission to $1 per dollar of monthly revenue booked. He told them that they were bringing the company unprofitable customers. But more importantly, they were wasting their customers' money. The idea of the docked pay was that it would "penalize the losers" and guide salespeople toward better behavior going forward.

Immediately the company solved its churn problem. Within two quarters, HubSpot's churn had dropped 70%. But for Roberge, a quant VP of Sales with a love for optimization, the experiments weren't over.

Signing more annual contracts

For any business, cash is king. For early-stage companies, it's the difference between life and death—or at least IPO and slow growth. As I've written on my blog, when a company is able to collect cash up front from customers, it can accelerate revenue growth (assuming LTV is higher than CAC). That's why it's so important to consider average prepayment commitment when building a compensation plan.



After about 18 months of running his "Customer Success Plan," Roberge decided to make a few changes to align with the company's goal of accelerating growth through cash management. By this point, HubSpot had product-market fit and was aggressively expanding. The company knew from historical data that if they invested a dollar in marketing and sales, they'd earn more than a dollar in revenue over the course of a customer's lifetime. Therefore every dollar, and the speed at which it was collected, mattered.

In order to increase the average contract length and collect more cash up front, Roberge created the following compensation model:

1. Salespeople would earn $2 per $1 of monthly recurring revenue
2. The commission would be paid out as follows:

a. 50 percent on the first month's payment
b. 25 percent on the sixth month's payment
c. 25 percent on the twelfth month's payment

Under the new plan, if a salesperson signed a month-to-month customer, they would have to wait an entire year to receive their full commission. Conversely, if they signed an annual contract and collected all the money up front, they would get the entire check in the first month.

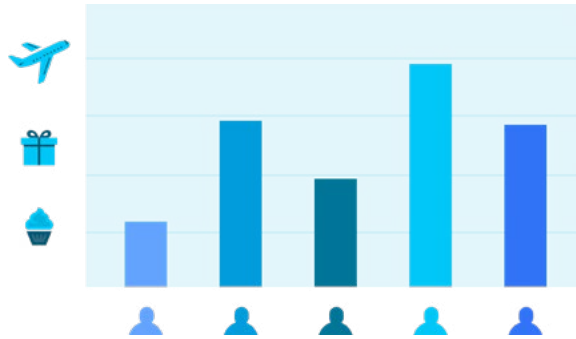
Once again, the new plan was a huge success. For the 18 months that the "Customer Success Plan" was in place, the average prepayment commitment was 2.5 months. When they implemented this new plan, that number jumped to 7 months.

The new plan aligned with the company's goals of customer success. They knew from experience that when a customer signed up for an annual contract, they invested more time and resources in setting up HubSpot's software. The plan also put more control in the salesperson's hands, making them masters of their own destiny. And of course, it gave HubSpot the cash to invest in further growth. It was a win for everyone.

Influencing short-term behavior with non-cash rewards

One effective tactic I've seen our portfolio companies use to incentivize the right sales behaviors is the non-cash reward. This is an effective way to influence something in the short term without making a change to the sales compensation plan. It's also fun and can help build a positive team culture.

Roberge likes to joke that the answer to every question in sales behavior is "run a contest." He says that if the sales comp plan is Batman, the sales contest is Robin.



In his book *The Sales Acceleration Formula,* Roberge offers the following tips for sales managers running contests.

Align the contest with a short-term behavior change desired for the majority of the team
Like sales commission plans, sales contests are a great way to drive home desired behavior. For example, you may fear a summer slump and want to boost activity in June. This desire would be difficult to pull off through the commission plan. However, an activity-based contest in June would do the trick.

Make the contest team-based
If there are twelve people on the sales team, form four teams of three salespeople and have the teams compete rather than have every man fight for himself. This approach has a remarkable impact on team culture, especially in the early phases of team building. For the first three years at HubSpot, every contest run was a team contest. The positive impact on team culture was remarkable. High-performing salespeople helped out their teammates who were lagging behind. The salespeople who were lagging behind worked late in order to avoid letting their teams down. After three years of team-based contests, a contest based on individual performance was tested. For the first time, there were accusations of cheating and backstabbing behavior on the floor. HubSpot immediately returned to team contests.

Make the prize team-based
In addition to making the contest team-based, choose a reward that the team experiences together. Rent a limo to take them to the casino, buy them a golf outing, or send them sailing for a day. Making the prize team-based maximizes the positive impact on culture. Not only does the team win together, they experience the reward together. They return to the office with photos of the great time they had ... together. People feel good about their colleagues, and teams feel motivated to win the following month.

Send out updated contest standings every night
At least once per day, the contest standings should be published to the entire sales team, if not to the entire company! This a critical execution point. Without daily updates, contest effectiveness will drop precipitously. Even if it means compiling and posting the results manually, publish the results every day.

Choose the time frame wisely
The time frame needs to be long enough to drive home the desired behavior change but short enough that salespeople stay engaged. A daily time frame is too short. Weekly contests are on the briefer end of acceptable. A quarterly time frame is probably too long. Monthly contests are ideal.

Avoid contest fever
Don't read this section and implement five simultaneous contests.

Conclusion

A common opinion is that sales should be responsible for signing up new accounts; many companies let other teams in their organization worry about how to actually retain the customer. But this way of thinking has negative consequences on any company's churn rate and revenue retention.

Instead, the key is to keep things simple and to align sales behaviors with the company's top-level business goals. And, across our portfolio, the companies that have grown fastest are the ones that experiment with every part of their business, like Roberge and the team at HubSpot have done.

CHAPTER 8

SALES OPERATIONS

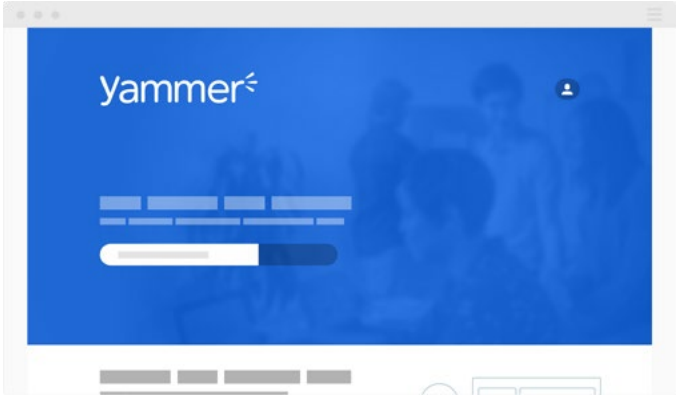
How to enable your team with great data and a repeatable sales process



Don Otvos

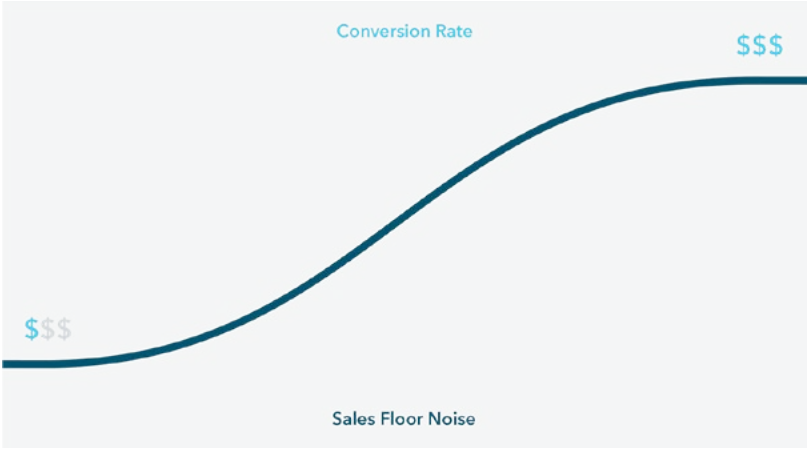
Solutions Engineering at SalesLoft

In the winter of 2010, I walked into Yammer's headquarters in San Francisco to interview for a sales operations job. At the time, Yammer was one of the hottest companies in Silicon Valley. They had just moved to San Francisco after spinning off from Geni and were on the cusp of a hiring spree. The company was featured in tech press nearly every week after recently winning TechCrunch Disrupt, and its founder, David Sacks, had just raised millions of dollars from his fellow PayPal mafiosos. But when the VP of Sales walked me through the sales floor that first day, I thought I was walking through a college library during finals week. It was dead quiet. By that point, I had been in sales for a decade and I knew that nothing was worse for a sales team than a floor with no noise.



It's been seven years since my early days at Yammer—a lifetime in Silicon Valley. Since then, the sales technology landscape has been flipped on its head more than once. Billions of dollars have been invested in making the modern salesperson's job more efficient through various

apps and workflows. Often, it is up to the sales operations team to implement those technologies. But for all the change the industry has seen, I believe that the role of sales operations is still the same: it is and always will be about enabling salespeople to spend as much of their day as possible selling and as little time as possible doing tedious work. In other words, our job is to keep the sales floor as noisy as possible.



Until robots replace the salesperson entirely (don't hold your breath), noise level and revenue growth will have a near-perfect correlation. When a salesperson isn't talking to a prospect, they are usually doing busy work like entering data in the CRM or researching companies—work that shouldn't be done by a rep making north of six figures.

That first day in 2010, I told David Satterwhite, Yammer's recently hired VP of Sales, "It's so quiet in here, I think I heard a pin drop back there."

"Why do you think we're interviewing you?" he asked with a laugh.

A couple of weeks after my first interview, I joined the company. Not everyone was so welcoming of my critique in those early days, though. At lunch during the first week, I told one of the salespeople that it was strangely quiet. He jumped into a list of reasons: "We're working

smarter, not harder on this team ... We use the latest technology, which means we don't have to spend as much time on the phone."

"Doesn't the latest technology help you spend more time talking to prospects?" I asked rhetorically.

Needless to say, I wasn't the most popular person at first. But it didn't take long until they came to appreciate the value of sales operations.

When I joined Yammer, the company had acquired hundreds of thousands of "freemium" users. But the only information we collected at sign-up was an email address. We kept this low barrier to entry in order to increase sign-ups. Of course this was no help when a sales rep wanted to sell the paid version of our product. They had no context for which industry the company was in, or who they were emailing at the company, or even where that person was located. Therefore, reps were spending hours researching prospects before they reached out to sell them on the premium version of Yammer. That first summer, I set out to fix that.

One day I sent a message on our Yammer network asking coworkers if any of their kids were home for the summer and needed internships. Within a couple of weeks, I took over a conference room and six interns were spending eight hours a day researching companies and updating records in Salesforce. We were able to identify our most active users and most active networks to work on first, and then worked down the list.

Not only was I able to use this information to design our sales territories, but with more context sales reps were able to send a targeted email to the right buyer at a company based on our win/loss analyses. When a prospect in the manufacturing industry answered the phone, one of our reps could now share a relevant use case from a similar company in that first conversation. But more importantly, our added precision—and subsequent increase in conversion rate—cost less money since we had divided the labor into high- and low-skills tasks.

At the end of that summer, we expanded our office from an old traditional cube layout to a new honeycomb style with an open floor. I couldn't help but smile when complaints started coming in that the noise on our floor was getting too loud. Sure enough, revenue was growing like a weed, and sales performance was at an all-time high.

Lesson 1: A quiet sales floor is your canary in the coal mine

Later that year, Peter Fishman, Yammer's Head of Engineering at the time, came to me with an idea. He and his team had been building out reports and databases to track two of our most important metrics: engagement and virality by network. We knew anecdotally that we had the most success selling our product into accounts where the freemium version of Yammer was "going viral" and being adopted quickly. We also knew that if Yammer grew too fast in an organization, the IT team was more apt to shut it down. So it was essential that the sales team prioritize their outreach and target the IT team for a call. But simply giving reps a list of the most active networks wouldn't work, since some of those companies were slow to adopt industries or small businesses. Peter's idea was to marry the database that our army of interns had compiled with the backend database.

Together with Peter's team, we created a data warehouse where we stored the firmographic account data from the CRM with the product usage data from our backend database. The idea was to leverage this information to run reports to learn where to focus our sales efforts.

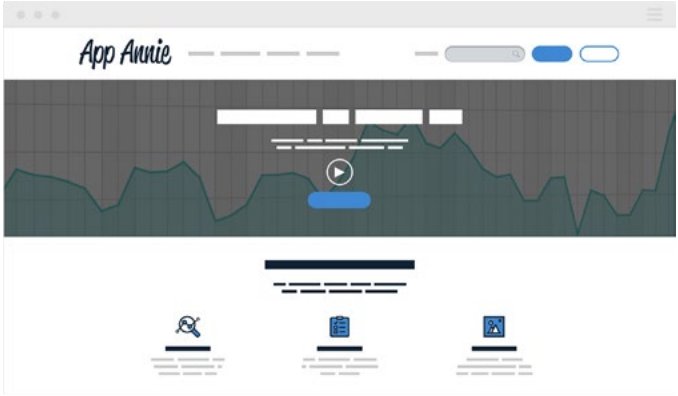
One of the first reports we ran was the geographic location of our most active networks. The results were surprising. Adoption was strongest in the Netherlands, of all places. When we called one of our biggest customers in that country, they explained why it was so popular with the Dutch. An executive at the company, Wolters Kluwer, told us, "In the Netherlands, companies have a very flat organizational structure.

It's not uncommon for an average employee to share an idea with their CEO. Yammer enables this communication in a very elegant way."

Of course, this was important to know because it informed where we invested sales and marketing dollars. We quickly began to focus efforts on freemium users in the Netherlands—a country that wasn't on our radar months earlier—and other adoption hot spots.

Over the years, I've learned that the sales operations role is more quarterback than kicker. We're not the janitors at a company tasked with cleaning up messy data; we're the people whose work informs business strategy. Without accurate (or any) data, it's difficult for an executive team to make key decisions about where to invest or forecast what revenue will look like.

At App Annie, where I went to work after Yammer, I became more convinced than ever of sales operations' importance. At the time, the mobile app store analytics company had just reached product-market fit. We were successfully selling our product to mobile gaming studios that wanted to understand who was searching for their games and how they were reaching the app. But that market was limited. We had raised venture money and knew that we had to expand to new horizons if we were going to join the unicorn club.



Shortly before joining App Annie, I had re-read *Crossing the Chasm*, a book all about expanding to new markets. In the book, author Geoffrey Moore breaks buyers down into four groups: early adopters, early majority, late majority, and laggards. The gaming studios were our early adopters. They had high appetites for risk and wanted to use the latest technology. The challenge was identifying what our early majority market would be.

One of the things that I noticed looking at our early customer data was the prevalence of "sexy" brands. We didn't have many large customers, but the ones we did were some of the most well-known brands in the world. In a meeting with Oliver Lo, who was our VP of Marketing, we decided that the reason was intuitive: companies that care about their brand were investing in mobile apps. The next quarter, we decided to create marketing and sales campaigns targeting these companies. Once again, I hired an army of interns to build a list of Directors of Brand and VPs of Digital Marketing at Fortune 500 companies. Then our SDRs started contacting them.

From the response and close data, we were able to validate the hypothesis further. We noticed that companies like Nestlé and McDonald's not only responded but cruised through our sales cycle quickly.

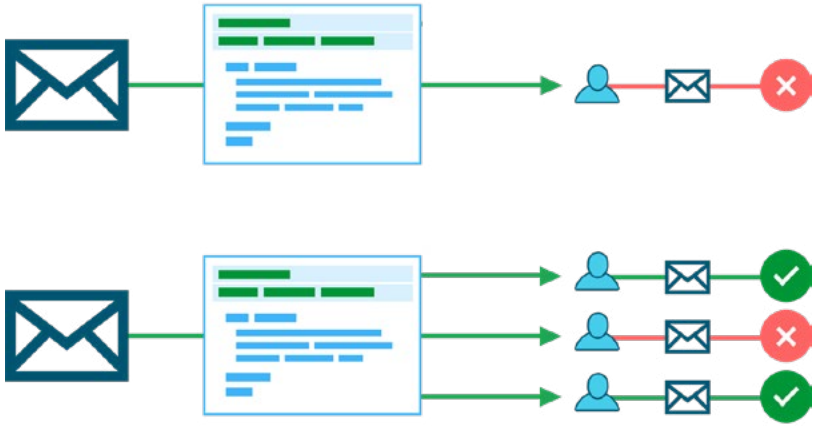
Of course, today there are products that automate all of this. Products like Infer and EverString help sales teams identify the common attributes among their customers and then find more prospects like them. But as complicated as the technology behind those apps is, the underlying premise is simple business strategy: figure out what is working and then try to replicate it. At every sales organization I've ever worked at that task falls on sales operations.

Lesson 2: Use data to inform business strategy

One of the things that I had always heard from my colleagues at App Annie and Yammer was that a sales rep should always avoid being

single-threaded. That is, during the sales cycle, a rep should include multiple contacts at a company in the conversation. If they first speak with a VP of Sales, they should ask, "Who else should we include in this conversation?" The reason is pretty simple: prospects often go dark on you. A VP of Sales might leave the company or go on a six-month sabbatical, and you don't want all your eggs—or that six-figure commission check—in one basket. At Yammer, we required a discussion with the company's IT team in order to pass their security requirements.

But I've been in sales long enough to know that there are a lot of myths and anecdotes that get passed along without any data to support their accuracy. So when I arrived at Datahug in 2014, I asked someone on the product team to check into the "Single-Threaded Theory." One of the benefits of joining the company was that we were selling to sales teams, so for the first time in my career I was the end user of the product I was pushing.



Shortly after I put in the request, Ray Smith, Datahug's CEO, had the engineering team anonymize our user data, crunch it, and build a report. In his findings, we learned that if you paired the engagement data—the frequency of emails between a rep and prospect—with the

salesperson's forecast, you could almost perfectly predict whether or not a deal would close based on that engagement. We also learned that there was a strong correlation between engagement and the number of contacts in an opportunity being communicated with. More engaged contacts meant more conversations throughout the sales cycle, which meant higher close rates. In other words, the Single-Threaded Theory was proven true.

At around that same time, LinkedIn launched a new feature where on each user's profile they recommended similar profiles. This "People Also Viewed" feature was a goldmine for sales reps, especially with the Single-Threaded Theory in mind. The data team at LinkedIn was doing my job for me by analyzing similarities in profiles and handing them over. They organized people that worked on the same teams, communicated frequently, and reported to one another. So on a VP of Sales' profile, we could see the reps that reported to them and the Director of Sales at that company, as they were almost always the most frequent profiles that were also viewed. That summer, I hired my third army of interns to go through each of our opportunities in Salesforce and go to their LinkedIn profile to add the names and titles from each People Also Viewed profile.

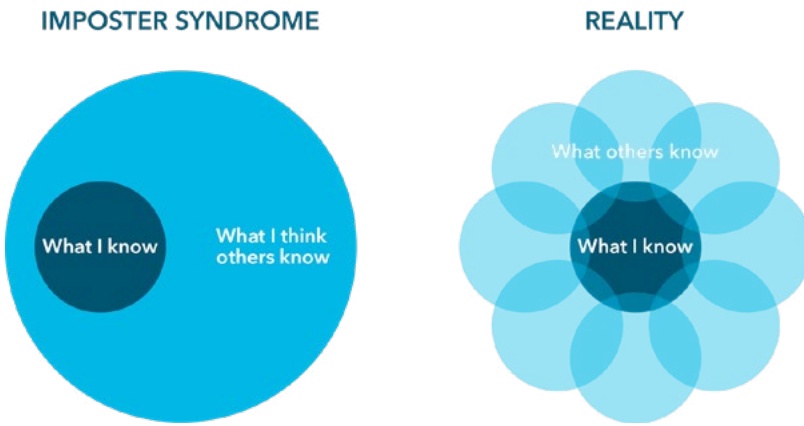


But once we had more than one contact per account, we started running into another problem. Salespeople always look for the path of least resistance and our team was no exception. They were given accounts with 10 contacts in some cases, and rather than tailor emails to each individual, they pasted in the same email copy and fired away. The results were less than ideal. To solve this, we worked with our marketing team

to write copy and create cadences for each customer persona. Sure enough, the response rate went up and face-palm moments went down.

Lesson 3: Don't let your reps get single-threaded

When I first started working in sales operations, I vividly remember coming home each day thinking two things: first, I thought I was in the wrong role because of the amount of mistakes I made, and second, I thought the problem immediately in front of me had no solution. My imposter syndrome—something that nearly every professional feels early in their career—has since gone away, and now whenever I run up against a hard problem, I have faith that my team and I will be able to solve it. But those evolutions in mindset didn't happen naturally or easily.



I've been lucky to spend most of my career in a place with a high density of sales operations leaders. It seems that each week there is an event in San Francisco that I can attend and learn something new. It was at Meetups and happy hours that I began to see that everyone in sales operations feels more or less the same way. We're all stumbling our way through problems and inevitably feeling stressed when we can't solve

them. Paradoxically, the smartest people I've met in the industry are the ones who are the first to admit what they don't know. They ask questions at these events, listen more than they talk, and are honest about the challenges facing their business.

So I guess that's my last piece of advice for the sales operations folks reading this. No one really knows what they're doing, but we're all in this together. So don't stress, don't be afraid to ask for help, and attend a happy hour (or two) to meet folks dealing with the same challenges you face.

CHAPTER 9

SALES COACHING

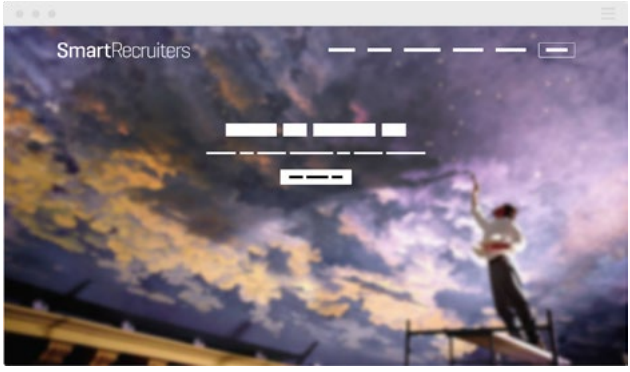
Using data to measure, motivate, and improve your sales team's performance



Taft Love

Global Head of Sales Development and Sales Operations at SmartRecruiters

On my first day as Director of Sales Development at SmartRecruiters, someone walked up to me and said, “Man, you really have the team fired up. The reps are calling like crazy.” At the time, the company didn’t have any formal sales development reporting in place, but there was enough infrastructure for me to run a quick report on how many calls the team was making per week. It was a paltry sum and far from crazy: 120 calls per week across a team of four reps.



The situation I walked into at SmartRecruiters is one that many people in sales operations and management find on their first day. There was little to no process; Salesforce was a mess, and when you asked a rep how many calls they were making per day or what their connect rate was, they froze up. In this chapter, I’ll explain how we changed that.

Over the last year, my team and I have built what I would argue is the most important asset to any sales team: a coaching and training framework. Data has always been at the core of it. From the beginning, our team has followed Peter Drucker’s mantra, “If you can’t measure it, you

can't manage it." It wasn't always pretty, and we made our fair share of mistakes, but eventually we built robust call reporting, a systematic way to diagnose problems our reps were facing, and a data-driven coaching framework.

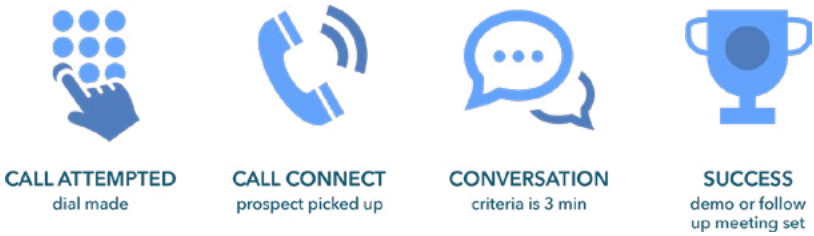
Step 1: Set up your reporting infrastructure

After that conversation I had on my first day, I knew that the first step toward increasing pipeline was to increase our call volume. I began by making an important hire: a former colleague and metrics-focused manager, Chris Bryson. We worked together to define and measure success in terms of activity volume. Once that was in place, I called a team meeting and shared the "Calls per week" report with all the SDRs. I've never been a fan of yelling or intense cheerleading to motivate a team, so I simply shared the report and walked everyone through the concept of an activity funnel in order to explain why increasing call volume was important. For most reps, this was their first experience seeing a sales and marketing funnel. Then I asked everyone how we could increase volume. Someone suggested using the dialer feature in our email automation tool and another rep suggested evaluating new sources of lead data for outbound, but the truth is that just being aware of their low performance was enough motivation to boost the numbers.

After a few months of calling, we learned two things: the team was willing to call if properly motivated, and the team's results weren't as good as what we had achieved at PandaDoc, our previous company. Fortunately, we had a wealth of data from which we hoped to tease out some conclusions about our performance. Chris and I began by coming up with a list of questions that we wanted to answer using the new data. Some of the questions we asked were:

- Are the reps' complaints about low connect rate valid?
- When they do connect, how often are they successful?
- When they connect and are unsuccessful, what is going wrong?

The Task record for every call in our data set included a disposition (outcome of the call) and call duration. We bucketed all of our dispositions into one of the following categories:

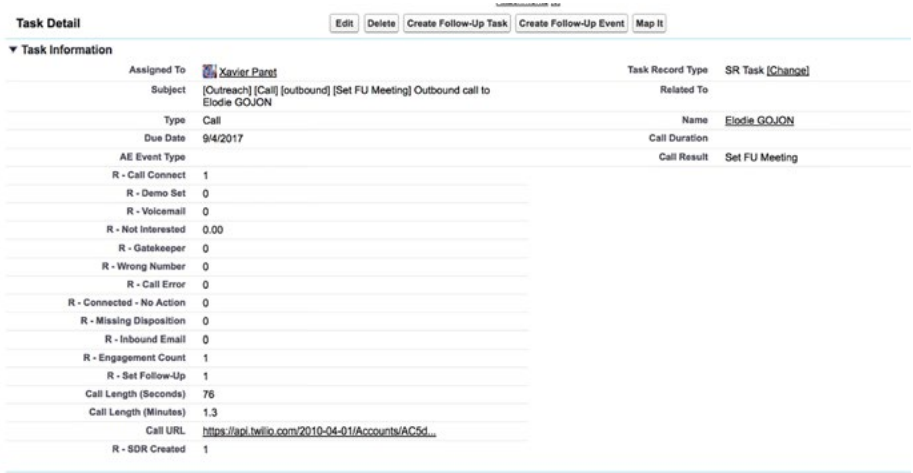


Unfortunately, when we started this process, our dialer, Outreach, sent the data we needed to Salesforce in a format that made reporting difficult. All of the important data was combined in text format within the task's subject line. This format is easy for a human to understand, but impossible for a computer to use in a meaningful way. For example, if a rep called someone, connected, and set a follow-up meeting, the subject field was populated with the following data: "[Outreach] [Outbound call] [Connect] [Set FU meeting]".

In order to turn this into useful data for reporting, I wrote a series of simple custom (formula) fields that read the subject line of each task and translated the text string into something that Salesforce's reporting module could read. For each disposition, I wrote a formula that looked like this:

```
IF(
CONTAINS( Subject, "[Set FU Meeting]"),
1,
0
)
```

In the above example, the formula would then push either a 1 or a 0 to a custom field named “R - Set FU Meeting?” (The R makes it easy to quickly pull up all custom fields that relate to reporting). Suddenly we could draw conclusions from the aggregate call data using these simple binary fields.



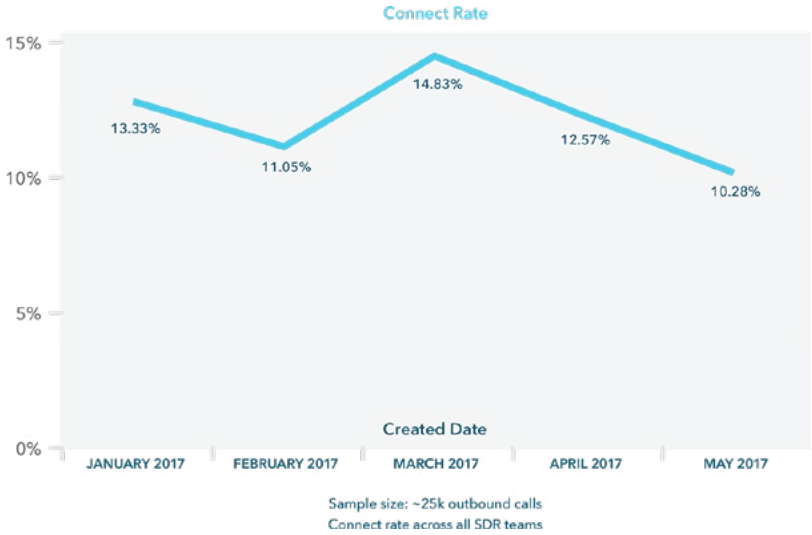
Task Detail Edit | Delete | Create Follow-Up Task | Create Follow-Up Event | Map It

Task Information

Assigned To	Xavier Paret	Task Record Type	SR Task [Change]
Subject	[Outreach] [Call] [outbound] [Set FU Meeting] Outbound call to Elodie GOJON	Related To	
Type	Call	Name	Elodie GOJON
Due Date	9/4/2017	Call Duration	
AE Event Type		Call Result	Set FU Meeting
R - Call Connect	1		
R - Demo Set	0		
R - Voicemail	0		
R - Not Interested	0.00		
R - Gatekeeper	0		
R - Wrong Number	0		
R - Call Error	0		
R - Connected - No Action	0		
R - Missing Disposition	0		
R - Inbound Email	0		
R - Engagement Count	1		
R - Set Follow-Up	1		
Call Length (Seconds)	76		
Call Length (Minutes)	1.3		
Call URL	https://api.twilio.com/2010-04-01/Accounts/AC5d...		
R - SDR Created	1		

Step 2: Diagnose the problem

Once we had our reporting infrastructure in place, we started answering the aforementioned questions. The first metric we looked at was our connect rate. How often did someone pick up the phone when we called? We found that the connect rate was low, but not far below what we expected. However, as our calling increased, the connect rate seemed to dip. It ranged from a high of almost 15% to a low of 11%, depending on the month.



Of course, the 15% connect rate was a team average, and if you know anything about sales management, you'll know that performance across the team varies widely. We had some reps that were connecting on 10% of their calls and others that were connecting on 20%. With that data, we looked into things like what time the low-performing reps were calling their prospects. In one instance, we had a rep on the West Coast that didn't realize that they were making a lot of their calls to reps on the East Coast at lunch time.

In addition to the connect rate, we looked into the success rate (the calls that connected and led to a demo or follow-up meeting). But once again, we needed to look deeper than one single metric to diagnose the problem. For example, results were choppy. Some reps had a success rate north of 30%, while others only succeeded on 1 out of every 10 connects.

In order to break this down further, we looked at each rep's average call duration when they connected. We learned that some reps were struggling with their opener, as evidenced by short average call durations. Conversely, we had reps who were obviously struggling to convince

prospects to take a meeting, since they had average call durations of 10 or more minutes and didn't result in a demo or follow-up meeting.

Step 3: Create a training framework

All of these reports helped managers use data to inform their 1:1 coaching, but we also started to see trends that required training the entire team. Over a whiteboard session, Chris and I began mapping out the training framework that is in place today. We knew that it was essential to get our team's buy-in, so our first action was to get everyone in a conference room to explain our plan.

In the meeting, we explained to the reps that as a group we had all been struggling with a few things (for example, our outbound email response rates were too low to achieve our team goals by focusing on email).

We reminded them that a spot on the team required one of two things: delivering results, or working hard and following a process (though preferably both). Then we explained the new training framework. For each training topic, we would do the following:

Host a presentation
One of the sales managers or high-performing reps would give a short presentation on a skill or new tactic that we wanted the team to try.

Team Q&A
After the presentation, we would discuss the content as a team. This was a chance to ask questions and to suggest tactical changes they could make to their workflow.

Homework
Each rep would be assigned homework after the presentation. For example, if the presentation was on probing for pain on a qualification call, the assignment would be to try the new tactic and then listen to a

recording of the call and provide a self-assessment for how well they did.

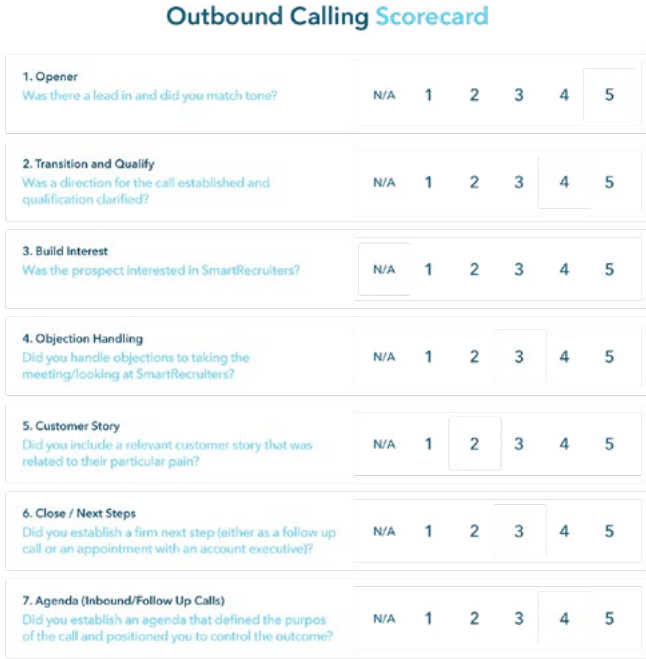
Manager review and 1:1
The sales manager reviewed the homework assignment and then had a 1:1 to discuss. In the pain-probing example, they would listen to the same call and provide their own feedback.

Two examples of training we provided were helping reps open calls and helping them close calls. As I mentioned before, we ran an analysis on the failed calls for each rep to see how long they lasted. About half of them weren't getting past the first minute or two with a prospect, and the other half were spending 10 or 15 minutes on the phone and still failing to set up another meeting. So Chris, the wizard of call-scripts-that-don't-feel-scripted, developed a series of trainings. We met with the team, Chris presented and explained the scripts, and then reps started implementing the new openers and closers. Because Chris left the scripts as more of a skeleton than a rigid word-for-word script, reps were able to deliver them in their unique voice. Then if that wasn't working, our 1:1 coaching sessions could get them back on track. Immediately we saw our success rates climb.



Here are the scripts we use today.[7]

As soon as we started delivering feedback to our reps, Chris and I knew that we needed a way to track the feedback and house all the data. That led us to start trying out products like ExecVision, Gong.io, and Chorus. After a month of testing each product, we chose ExecVision and started using their scorecard feature to deliver consistent feedback at scale. But we didn't want to inundate reps with a bunch of post-call questions and a heavier workload, so we introduced each area of self-assessment with the trainings every two weeks. For example, when Chris gave his presentation on opening scripts, he asked reps to start assessing their calls with a simple 1–5 score. We asked them to answer a simple question: "Was there a lead in and did you match tone?" After a couple weeks, we introduced a second question. And so on. Today the scorecard has seven questions that reps answer after each call.



When I first read sales management books on training, I thought giving reps homework like this and asking them to assess their skills so frequently would be condescending. But interestingly, I found that reps actually welcomed the feedback. Many of them had no professional experience or had never worked in a company with good management, so they were starving for this type of tactical feedback. The structure we put around our coaching also made the feedback feel less arbitrary. Every rep knew that they would get feedback each week—it was just part of the process.

On that first day about a year ago, when I sat down with the team and explained the funnel, I told them that our team is a machine. Our activity is the fuel and it's our job to make the machine run as efficiently as possible. After all, that's the way to take home bigger and bigger commissions.

Often, sales teams look at one part of the machine and build process around it. They may systematically A/B test outbound emails or write out call scripts and iterate on them. Rarely do teams focus on what I'd argue is the most important part: the people. While it seems almost dehumanizing to see the team as a machine to be tinkered with, like a car, it's what everyone wants and needs. Steve Jobs used to say that if you change the direction of a vector at the beginning of its path, the change is dramatic a few miles into space. I think of sales coaching in a similar way: each improvement to the process nudges the vector a little bit. But the impact on culture, revenue, and the reps' professional development is huge.

CHAPTER 10

SALES HIRING

Using data to interview and hire sales rep candidates



Mark Roberge

Best-selling author and former CRO of HubSpot

It's been two and half years since I published my book, *The Sales Acceleration Formula*, and over a year since I left my job as Chief Revenue Officer at HubSpot. A lot has changed in that short time period, but two things are still true for sales teams: hiring great salespeople continues to be the most important activity, and there is still no single formula that every company can apply to the process.

When I joined HubSpot in 2007, the company was a small startup with three people and a small office in Cambridge, Massachusetts. Within seven years, the company scaled to more than $100 million in revenue. One of the most essential parts of that journey was figuring out the process by which we hired salespeople. Every time I get up on stage to speak or meet with founders, they ask me the same question: What should I look for in a salesperson? As it turns out, this is a dangerous question, and I had to learn it the hard way.

Some of the first reps that I hired at HubSpot were top performers at their last company. I wined and dined them, took them to ball games, and even convinced a few to join our small startup. Yet, some of those "sure bets" didn't work out. So what happened?

I learned early on that each salesperson is gifted in what I refer to as a "sales context." There are reps who can close dozens of deals a week in a transactional SMB environment, and then struggle to close a single enterprise deal. Some reps were born to sell to marketers, and others can hardly speak the same language as the VP of Marketing on the other end of the line. Similarly, every company has its own sales context. Some have sales cycles that last multiple quarters. Others close deals on the first call. When the unique strengths of the salesperson align with the company's sales context, it is a beautiful thing. When they do not, it becomes an uphill battle.

Fortunately for sales managers, success (and failure) in sales is arguably the most quantifiable compared to other functions (e.g., product, HR, marketing, etc.). Key performance indicators (KPIs) vary by company, but ultimately customer lifetime value is what determines success for every rep and sales team. The question is how to hire the people that can accelerate healthy revenue and beat their quotas.

At HubSpot, I spent many restless nights thinking about this question. Over the years I developed a process that helped us identify good candidates in the hiring process, evaluate them, and then feed that data back into the model. It was simple, but at times the insights that we gathered were counterintuitive. Below, I've shared that process in hopes of helping you build your own unique sales hiring process.

Step 1: Establish a theory of the ideal sales characteristics

First, I listed the characteristics I thought would correlate with sales success. For each characteristic, I documented a clear definition. What did I mean by "intelligence?" What did it mean to be "aggressive?" My intention was to score each candidate on a scale of 1 to 10 for each characteristic. Therefore I needed to define what a score of "1" versus a score of "5" versus a score of "10" represented for each characteristic. For each candidate, I summarized the results on an Interview Scorecard.

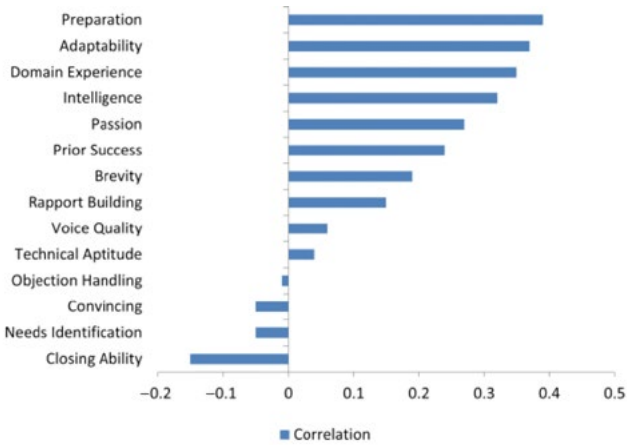
Step 2: Define an evaluation strategy for each characteristic

Once I defined the characteristics I was looking for, I needed a plan to evaluate candidates on each characteristic. What behavioral questions could I ask? Would I use role plays? Should there be an exercise for

the candidate prior to the interview? How could I leverage reference checks?

Step 3: Score candidates against the ideal sales characteristics

Back in the early days of HubSpot, I simply filled out the Interview Scorecard after each interview. The process was not overly sophisticated. I used Microsoft Excel. (We were a startup—I needed to be "hacky.") The key to the process was discipline, not sophisticated technology. I documented my findings and learnings as I went, and used them to constantly tweak my approach.



Here are the results from our first model.

Step 4: Learn and iterate on the model while engineering the Sales Hiring Formula

A few months in, I had a handful of salespeople on board. Many were doing great. A few were progressing more slowly than others. By remaining disciplined in the process described in Step 3, I was in an

optimal position to learn from these first hires and begin to understand our ideal hiring criteria. I was ready to engineer my company's sales hiring formula. I simply went back to the Interview Scorecards for the top performers and asked myself the following questions:

- Which characteristics do these top performers have in common? Are these characteristics predictors of success here at HubSpot? Once I identified them, I increased the weight of these characteristics.

- Which characteristics do not seem to matter? Which characteristics do not predict success? I needed to decrease the weight of these characteristics or eliminate them altogether.

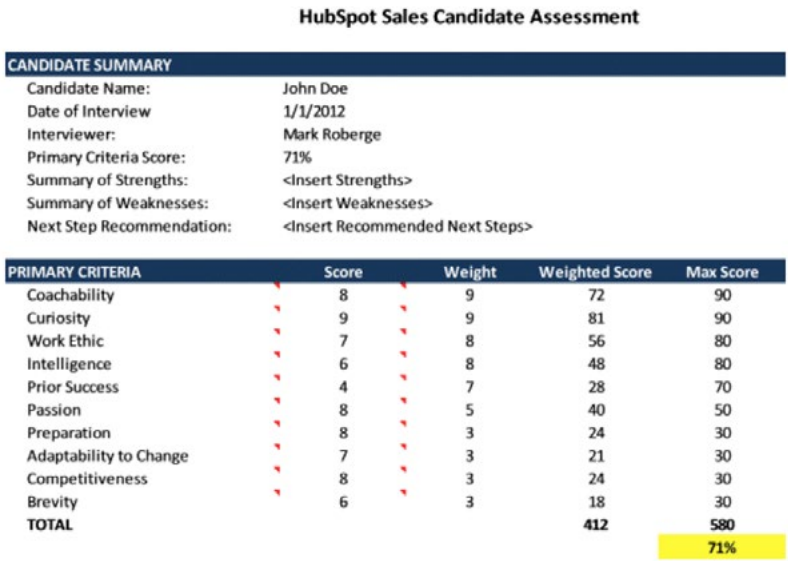
- What am I missing? I had to think beyond the Scorecard and reflect on these top performers. Was there another consistent, meaningful characteristic to be found among them? If so, I had to add the characteristic to the Interview Scorecard and start rating candidates on it.

While we designed this process with scale in mind, it can work for almost any company no matter its stage or size. Recently I met with two founders of a SaaS company based in the northeast United States who told me that they wanted to hire a new breed of salesperson. Their idea was to only hire a very specific profile: someone who had graduated from a Top 5 MBA program (Harvard, Columbia, etc.) and then went to work for a top-tier consulting firm (McKinsey, Bain, etc.). The idea was that a high business acumen and intelligence level would lead to higher sales efficiency. But while the early results were promising, they eventually struggled with this profile. The reps weren't accustomed to 50% of their compensation being variable, they hadn't developed the same tactics that seasoned reps do, etc. Fortunately they were tracking their hires with a model similar to the one we used at HubSpot and were able to pivot their hiring strategy quickly.

Given the fact that every company has its own unique sales context, it's natural to be surprised by what characteristics work out the best. After the first year running our model, we came to some counterintui-

tive findings: the characteristics that are traditionally associated with salespeople, such as aggression and strong objection-handling ability, had the worst correlation with success.

Once we had a large-enough sample size, we were able to start running regression analyses against important metrics like "revenue booked by rep." We took the score each rep received during their interviews and compared them to their quota attainment and bookings later. Then we fed that data back into the model, which enabled us to weight certain characteristics more strongly.



HubSpot Sales Candidate Assessment

CANDIDATE SUMMARY	
Candidate Name:	John Doe
Date of Interview	1/1/2012
Interviewer:	Mark Roberge
Primary Criteria Score:	71%
Summary of Strengths:	<Insert Strengths>
Summary of Weaknesses:	<Insert Weaknesses>
Next Step Recommendation:	<Insert Recommended Next Steps>

PRIMARY CRITERIA	Score	Weight	Weighted Score	Max Score
Coachability	8	9	72	90
Curiosity	9	9	81	90
Work Ethic	7	8	56	80
Intelligence	6	8	48	80
Prior Success	4	7	28	70
Passion	8	5	40	50
Preparation	8	3	24	30
Adaptability to Change	7	3	21	30
Competitiveness	8	3	24	30
Brevity	6	3	18	30
TOTAL			**412**	**580**
				71%

HubSpot sales candidate assessment

But we quickly learned that revenue and quota attainment weren't the only metrics that mattered. Sales reps can potentially close deals by selling to the wrong customers or by using sales pressure tactics. However, those deals will usually churn fast. So an important metric for success was whether or not the customers had a long lifetime. LTV

also takes into consideration the size of the monthly payment, which can indicate how well a rep prioritizes larger deals and manages their overall funnel to get the best revenue out of a given set of small and large opportunities.

Unfortunately, LTV wasn't the perfect metric to measure either. One of the first things we realized was that it didn't consider changes in rep behavior fast enough. For example, if a rep who was overselling deals—which leads to high churn—starts finding better customers, their LTV won't start to go up until one or two years later. That's a problem because the model will always lag reality.

This realization led us to develop leading indicators that could help predict LTV and future bookings by rep. Rather than look at a rep's LTV, we began looking at things like the engagement of their customers within the product. We broke our product into 25 different features and then queried if the customer had used at least five of those features within the first 90 days. If they did, they were seeing the value of our product. If not, they were more likely to churn. This alone helped us get ahead of the problem by three or four quarters.

Once we had a model we were confident in, we learned that HubSpot's best reps shared five characteristics:

1. Coachability
2. Curiosity
3. Prior success
4. Intelligence
5. Work ethic

Here's how we interviewed for each of those.

Coachability

Coachability is the ability to absorb and apply coaching.

Coachability was the most significant influencer of my hiring decision. As I think back to most of the rock stars we hired, their coachability was the personality trait that really stood out in their interviews. Evaluating this characteristic consumed the majority of my interview. Here is the three-step process I employed to evaluate this characteristic.

Step 1: Set up a role-playing exercise that models your buyer context
After some rapport-building questions at the outset of the interview, I would verbally set up a role play with the candidate.

> [Hiring Manager] "Jess, let's do a role play. I am going to play the role of VP of Marketing at a security software startup here in Boston. The company has about 20 employees. The marketing team is small—only two people. As the VP of Marketing, I ended up as a lead in the CRM system, and I was assigned to you. As you reviewed the lead details, you saw that I had visited the HubSpot website last night and downloaded the company's eBook on inbound marketing. We will role play your opening call with me. Your goal is to do some light discovery and set an appointment to discuss my needs further. Do you have any questions? If not, please begin when you are ready."

Step 2: Evaluate the candidate's ability to self-diagnose
Once the role play was complete, I would ask the candidate to self-assess.

> [Hiring Manager] "Great work, Jess. How do you think you did?" Jess's response to this question represented the first insight about her coachability that I collected. I wanted to see how reflective and analytical the candidate was about her performance. If the candidate simply stated, "I did great," that was a bad sign. I wanted to see the candidate reflect on and analyze her performance. I wanted to hear specifics about what she thought she did well and what she

thought she could have improved. Next, I would build on some of her observations. [Hiring Manager] "Great reflection, Jess. I agree with many of your points. You mentioned that you could have done a better job handling my question on SEO. If we could rewind to that section of the role play, what would you do differently?"

A candidate with a high degree of coachability is able to reflect, self-diagnose, and propose improvements to their weak areas. At this point, I would provide the candidate with the opportunity to demonstrate these abilities.

Step 3: Evaluate the candidate's ability to absorb and apply coaching
At this point, I would begin some proactive coaching to see how she would absorb and apply the feedback. Absorb and apply: these two actions represent the essence of strong coachability. Some people struggle to even absorb the coaching, perhaps because they are poor listeners or simply don't recognize the importance of feedback. Others absorb the information but struggle to apply it, perhaps because they are less adaptable or less skilled at thinking on their feet. I want to hire candidates who can both absorb and apply coaching.

[Hiring Manager] "Okay, Jess, in every interview I provide one area of positive feedback and one area of improvement."

Both components of this statement are important. If I offer only opportunities for improvement, the candidate might think she is bombing the interview. I run the risk of her freezing up on me, preventing me from evaluating her true abilities. By leading with a bit of positive feedback, I strike a warmer tone. After hearing a bit of praise, the candidate is more likely to feel comfortable and behave normally.

[Hiring Manager] "I thought your opening rapport building was great, Jess. I liked how you broke the ice and created an immediate connection when you talked about your visit to Wrigley Field as a child. The area in which I would like to see improvement is the

depth at which you seek to understand the prospect's goal. Let me teach you how we deepen goal discovery here at HubSpot ..."

I would then begin to coach the candidate. By this point, I would usually be up on the whiteboard, coaching her and also closely observing the candidate during this process. Is she glassy-eyed or is she taking notes and asking good follow-up questions?

After a few minutes, I would ask if the process made sense. I would request that she redo the role play, this time attempting to apply some of the coaching I had just provided her.

Now, most people really mess up the second pass. Their heads are spinning. They know the job is on the line. They are sitting with the VP of Sales. They've just received my feedback and must immediately apply it. In this situation, I am looking for effort, not perfection.

I will say that I have probably conducted well over 1,000 interviews during my six years in the head-of-sales seat at HubSpot. Across the full population of candidates I've screened, perhaps only five people absolutely crushed the second role-play attempt. Those who did so became absolute rock stars in our funnel. What's the takeaway? Don't expect perfection, but rather look for effort. If you witness perfection, hire that candidate at all costs! You've just spent 10 minutes with a candidate and witnessed meaningful improvement over that short time. Imagine how much progress you could make in a day, a week, a month!

Curiosity

Curiosity is the ability to understand a potential customer's context through effective questioning and listening.

I have taught several classes on the subject of sales at MIT, Harvard, and other top universities across the United States. One of my favorite ways to start the class is to ask the students, "What makes a great salesperson great?" The most popular answers are consistent across venues and audiences: "aggressive," "convincing," "great presenter," "money hungry."

I don't think anyone has ever given me the answer I am looking for. Great salespeople are naturally curious. They ask great questions, listen intently, and probe into points of interest.

Great salespeople ask questions of potential customers in a manner that does not feel interrogative. Instead, potential customers feel like great salespeople are genuinely interested. After all, if the salespeople are truly great, they genuinely take interest in the responses of their prospects.

Great salespeople educate potential customers through the questions they ask. Their questions are thought provoking and elicit introspection. "You know, nobody has ever asked me that before. Now that I think about it ..."

Great salespeople quickly build trust in order to earn the right to ask personal questions and to receive honest answers in return.

Great salespeople seek to understand customer goals, aspirations, fears, and struggles—all through tactical questioning.

Students often ask me, "Mark, how can I prepare myself to be a top performer in sales?" I offer the following advice: the next time you are at a wedding reception, a school networking event, or a party on a Friday night, approach a stranger and ask them questions. See how long you can question that individual without mentioning anything about yourself. If the individual walks away from the conversation feeling interrogated, you need more practice. If the individual walks away thinking, "Wow, that was a really smart and interesting guy," you are on your way to becoming a great salesperson.

So how do we test candidates for curiosity? There are many areas in the interview process, but I will highlight two especially important opportunities.

The first test of curiosity happens the moment I meet a candidate in the lobby. “Hello, Jess. My name is Mark Roberge. Thanks for coming in today.”

Does the candidate start with a question? Does the candidate ask me about my day? Did the candidate research my background and does she take the opportunity to reference an observation from her findings? Based on my responses, does the candidate follow up with smart, open-ended questions to learn more? If all of these things are happening, then this interview has started really well for me (and for her).

The second test of curiosity occurs during the role play. Let’s use the same role play that I set up in the previous section on coachability.

> Ring Ring.
>
> [Marketing Manager] “Hi, this is Mark.”
>
> [Candidate] “Hi, Mark. This is Jess at HubSpot. Did I catch you at a bad time?”
>
> [Marketing Manager] “I have one minute.”
>
> [Candidate] “Great. I am not sure how much you know about HubSpot. We have an all-in-one marketing platform that helps companies get found online and convert visitors into leads and customers using your website. We have worked with Company X and Y in your industry. I am calling to see if you would be open to a 10-minute call to assess your company’s online visibility.”

Disaster!

Let’s try that again. This time, Jess will be far more curious.

> Ring, ring.

[Marketing Manager] “Hi, this is Mark.”

[Candidate] “Hi, Mark. This is Jess at HubSpot. Did I catch you at a bad time?”

[Marketing Manager] “I have one minute.”

[Candidate] “Great. I noticed you downloaded our eBook about generating leads on Facebook. What specific questions did you have about Facebook marketing?”

Yes!

[Marketing Manager] “Oh, jeez. Um. I was just doing some research. I think I was looking for examples of B2B companies that had seen success using Facebook for business purposes.”

[Candidate] “Great. I am happy to share some of those success stories. Have you run any Facebook campaigns yet?”

Yes!

[Marketing Manager] “I have.”

[Candidate] “How did they go?”

Yes!

[Marketing Manager] “Okay.”

[Candidate] “Okay? What do you mean? ”

Yes!

> [Marketing Manager] “Well, we did generate a lot of new emails from the campaigns. Unfortunately, I am not sure the people we are attracting are really qualified for our service.”
>
> [Candidate] “Interesting. What types of people are qualified for your service? What types of people did the Facebook campaign attract? ”

Yes!

Does the candidate lead with great questions? Or does the candidate “show up and throw up,” as we say in the industry?

Does the candidate ask about the specific areas the prospect cares about? Or does the candidate bore the prospect with her company’s elevator pitch?

Curiosity is the ability to understand a potential customer’s context through effective questioning and listening.

Prior Success

Prior success is a history of top performance or remarkable achievement.

It is probably the easiest characteristic to evaluate, especially if the candidate is coming from a reasonably sized sales force. It is the most objectively measurable of the “big five” traits.

> [Hiring Manager] “I noticed you were an account executive at your last employer. How many account executives were there at the company?”
>
> [Candidate] “125.”
>
> [Hiring Manager] “Where did you rank?”

[Candidate] “Six.”

[Hiring Manager] “Wow! Impressive. What metric is that rank based on? Bookings? Attainment?”

[Candidate] “Bookings.”

[Hiring Manager] “And the rank is based on last quarter or all of last year?”

[Candidate] “All of last year.”

[Hiring Manager] “Very good. And your references will verify that performance?”

[Candidate] “Of course.”

I am looking for top 10 percent. If the candidate falls outside of that range, the candidate really needs to rank extremely high on the other key characteristics in order to earn an offer letter from HubSpot.

Evaluating prior success becomes more challenging when the candidate does not come from a reasonably sized sales organization or does not come from sales at all. In these cases, I evaluate prior success through other activities in the candidate’s life. How did the candidate perform academically in school? What was her class rank? What were the candidate’s standardized test scores? Was the candidate a standout performer on a varsity sports team? Was she the captain of the team? Perhaps she contributed to a major championship? Was the candidate a member of student government or a leader of an extracurricular organization? If the candidate is transitioning from a nonsales background, how did the candidate differentiate herself from her peers in her current or former role? What made her special?

On the HubSpot team, we had an Olympic gold medalist. We had a cello player from the Portland Symphony. We had a former comedian

who was featured on Comedy Central. These people pursued their life passions with exceptional vigor and performed at a top percentile level. These people are likely to bring that same passion and competitive drive to their role in sales.

Prior success is a history of top performance or remarkable achievement.

Intelligence

Intelligence is the ability to learn complex concepts quickly and communicate those concepts in an easy-to-understand manner.

Not every sales team needs intelligent salespeople. For example, in a commoditized buyer context, I would bet that work ethic, rather than intelligence, is the superior predictor of success. However, in the HubSpot buyer context, intelligence proved to be a strong predictor of sales success. In retrospect, I believe intelligence was a key trait because our industry was evolving so rapidly. For context, Twitter was a garage project when we first started selling HubSpot. Just seven years later, it is a $25 billion technology titan. That should give you a sense of how quickly the industry was transforming in the late 2000s. Our salespeople needed to keep pace as the industry around us took shape. They needed to understand new concepts and communicate to our target customers exactly how those concepts impacted optimal marketing strategies. Because most early-stage companies operate in rapidly evolving industries, I expect that intelligence would be a predictor of sales success in their buyer contexts as well.

I tested intelligence by effectively commencing HubSpot sales training during the interview process. I exposed candidates to new information early in the interview process and observed their ability to absorb the information and communicate it back to me at a later stage in the process. For example, at the end of my first phone screen with a candidate, I would send her training materials on the concepts of inbound marketing, SEO, blogging, and social media. I would ask her to learn the

material before our next interview. Then, I would be sure to reference the materials in our next role-playing session.

Here's an example of testing for intelligence and information retention:

> [Mark] "Jess, I noticed on your website that you offer SEO services. I always wanted to better understand how I could improve my business's ranking in Google searches. Could you explain how I might go about doing that?"

Her response would offer me a first impression of her performance on this characteristic. To reinforce my earlier point, I am trying to understand two things here. First, how well did she understand the concepts to which I had I exposed her? Second, how well did she communicate those concepts back to me in a simple manner? I would always ask follow-up questions until I eventually stumped the candidate. The deeper I was able get on a topic before her responses suffered, the better it meant she was performing.

Intelligence is the ability to learn complex concepts quickly and communicate those concepts in an easy-to-understand manner.

Work Ethic

Work ethic is proactively pursuing the company mission with a high degree of energy and daily activity.

Work ethic is probably one of the most difficult characteristics to evaluate. These are the three techniques I use to gain insight into each candidate's work ethic.

1. **Observations during the interview process:** A lot can be learned by simply observing a candidate's mannerisms and behaviors during the interview process. This is especially true for assessing work ethic. How quickly did she return phone calls? How quickly did she turn around deliverables (such as her resume, her assessments, or her feedback from the interview)? Did she push the pace of the interview

process or were we pushing her? All of these observations provide insights into the candidate's work ethic.

2. **Reference checks:** Conversations with former supervisors or peers represent opportunities to assess the candidate's work ethic. Don't ask, "Did the candidate work hard?" Instead, ask the following questions: "Here are four characteristics that may describe a candidate: coachability, curiosity, intelligence, and work ethic. Could you please rank those characteristics from strongest to weakest for this candidate? Why did you rank these characteristics in the order you did?"

3. **Behavioral questions:** I often use behavioral questioning to explore the level of rigor with which the candidate approached her responsibilities— for example, "Please tell me about your typical work day or work week. What are some of your must-do activities?"

Work ethic is proactively pursuing the company mission with a high degree of energy and daily activity.

It's worth pointing out one more time that these were the characteristics that made sales hires at HubSpot successful because they fit our sales context, not because we identified some law of sales hiring that applies to everyone. At a company of a different size, or one that sells pharmaceuticals, the criteria might vary widely. The point of all of this is to show the process and our methodology, and hopefully inspire a few sales teams to develop their own.

I was lucky enough to join a company like HubSpot that treated almost every aspect of business like a science experiment, and hiring was no exception. From the day I joined until the day we IPO'd, we followed the same process: we started with a hypothesis, we tested it, and then we iterated.

NOTES

Chapter 1

1. Kyle Poyar, "How Meetup Built Demand for an Enterprise Product," https://labs.openviewpartners.com/how-meetup-built-demand-for-an-enterprise-product
2. Tom Tunguz, "A Structural Pricing Competitive Advantage In SaaS - The Three Part Tariff," http://tomtunguz.com/three-part-tariffs/
3. Kyle Poyar, "What SaaS Companies Can Learn From Gym Membership Pricing," https://labs.openviewpartners.com/saas-membership-pricing/#.WNAgdBiZPOY

Chapter 2

4. Clearbit, "The Modern Guide to Lead Qualification," http://blog.clearbit.com/the-modern-guide-to-lead-qualification/

Chapter 4

5. Clearbit, "Automated Outbound Sales: How Segment Generates $10,000 in Revenue Per Day," https://www.youtube.com/watch?v=-Bo8RJANE2Y

Chapter 7

6. Mark Roberge, "The Sales Acceleration Formula: Using Data, Technology, and Inbound Selling to go from $0 to $100 Million," https://www.amazon.com/Sales-Acceleration-Formula-Technology-Inbound/dp/1511318201

Chapter 9

7. SmartRecruiters Call Scripts, https://docs.google.com/document/d/1GkxhxYkbSc917hrCuocjfS59a23KDwZVL4SgU08nHsI/edit

ABOUT CLEARBIT

Clearbit provides the data backbone for modern sales and marketing teams. For more information, visit clearbit.com or follow @Clearbit on Twitter.

ABOUT THE EDITORS



Matt Sornson

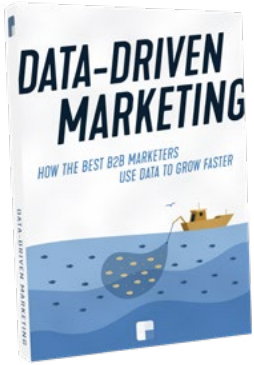
Head of Growth
at Clearbit



Michael Thomas

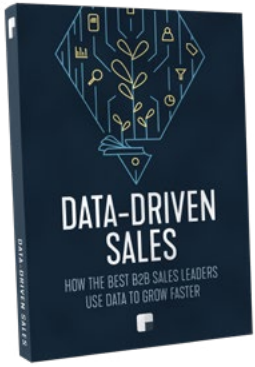
Michael Thomas is a writer and filmmaker. Previously he founded and sold SimpleData.

BOOKS BY CLEARBIT



Data-Driven Marketing
How the Best B2B Marketers Use Data to Grow Faster

Visit www.clearbit.com/books/data-driven-marketing to learn more.



Data-Driven Sales
How the Best B2B Sales Leaders Use Data to Grow Faster

Visit www.clearbit.com/books/data-driven-sales to learn more.

DATA-DRIVEN SALES is a blueprint for sales leaders and entrepreneurs that want to use the latest techniques to grow revenue and optimize their sales machine. Experts from companies like Hubspot, Digital Ocean, and Google cover topics from pricing to forecasting to lead scoring and more.

Including over 20 companies

yammer · logikcull · Segment · HubSpot · App Annie · DigitalOcean · Kahuna · GECKOBOARD

"Every chapter shows exactly how a company overcame their challenges and tactical ways you can do the same. Most sales advice in SaaS is generic. This isn't."

HITEN SHAH
Co-Founder at Product Habits

ISBN 978-0-9557163-0-0