Data-Driven Beer Planning

By Adrian Febre
Owner, Head Analyst
Bits to Barrels

The Evolution of Brew Planning

- Early on, brew planning for many of us is simple: if the number of kegs of a core beer seems low, we brew more
- As time goes on, we may upgrade that system, and figure that if we sell two kegs of beer in a given week, and it takes four weeks to brew and package, we should brew it when there are at least eight kegs left
- Perhaps we even enter this logic into a spreadsheet at some point, which applies that basic arithmetic to all of our core SKUs
- But the next step up after that is usually brewing software that is great at handling the brewing process but not as great at forecasting “deadline” brew dates
- We found out that we needed a better solution the hard way

The Problem

How We Used to Plan Brew Days

- The beauty of having a finite number of tanks and a growing number of core beers is that you barely have to plan at all, since you always brew what you think will run out first and hope nothing bad happens
- That worked well for us when we had a number of core beers roughly on par with the number of fermenters, and which were all ales (which is to say, quickly fermenting beers)
- Only brewing ales also meant that all beers took the same approximate amount of time, so your smallest pile was what you brewed first
    - Even then - is it your smallest pile of slow-moving beer, or the second-smallest pile of slightly faster moving beer?

When That Stopped Working

- That method stopped working once we introduced a popular lager
- The first problem was that we needed to plan brew days far in advance, but how far, exactly?
- The second problem was that we also had to plan brew days far enough apart from each other so as not to cause inventory issues (imagine brewing that lager every week - you'll get an 18 car pile-up in a few months), and it's this competition between **inventory efficiency** and **low risk** that really drives the need for thorough, math-driven beer planning

Cont'd

- For the sake of argument, here's a list of factors that may signal a general brewery's need for more rigorous planning:
- Brewing lagers (though not one-offs, like an annual Festbier)
- Brewing a large number of core beers
  - By "large" I mean more than your number of fermenters, roughly
- Brewing fast-moving beers, like Hazies
- Brewing beers whose popularity is highly variable

What a Picture's Worth

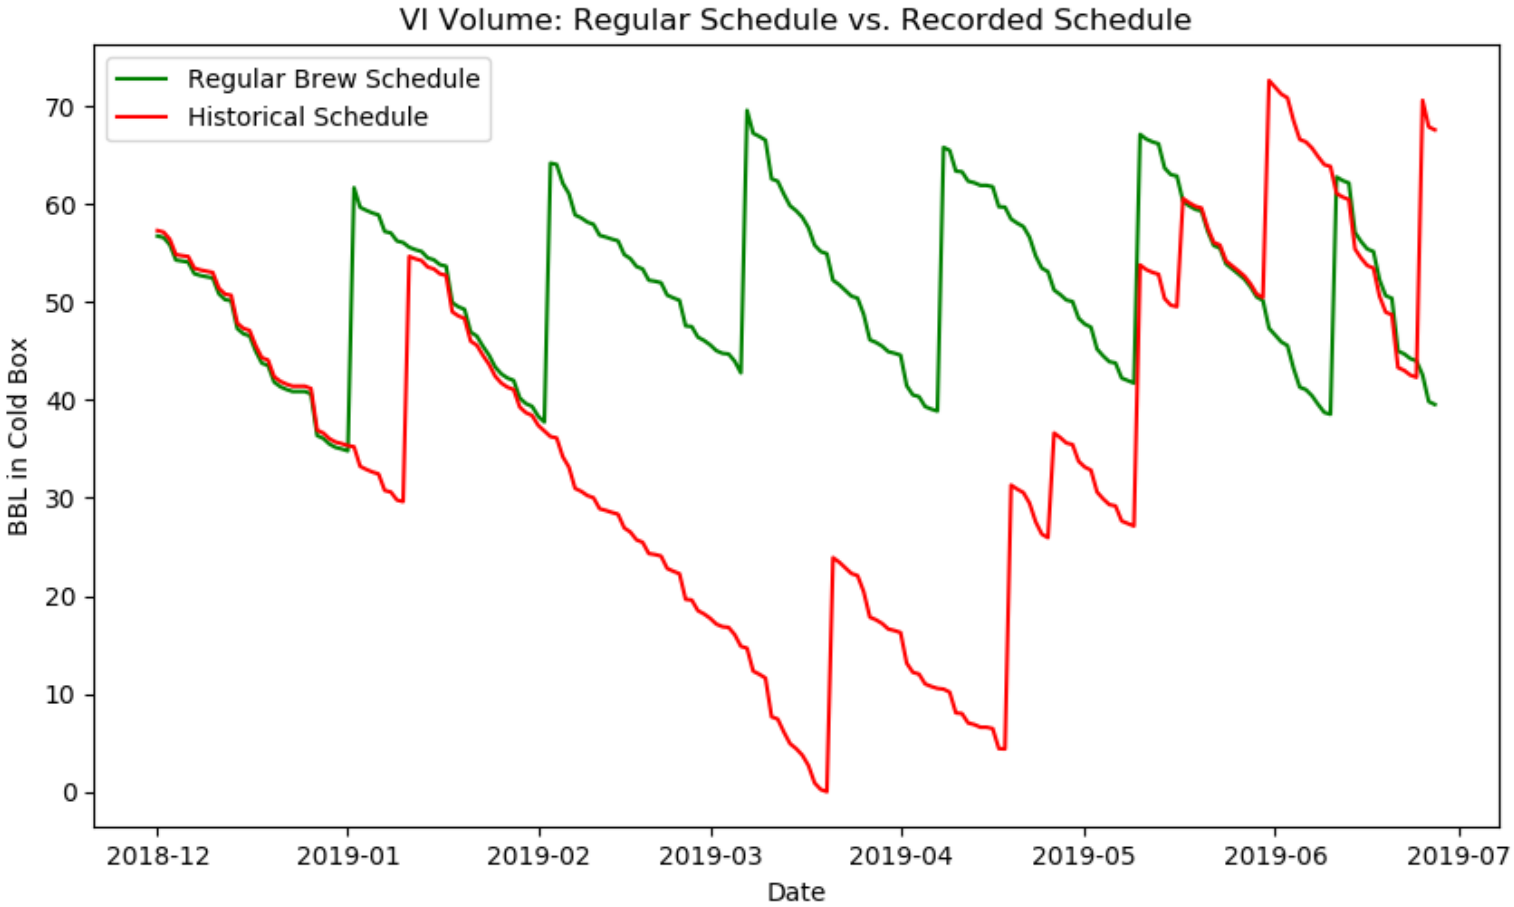

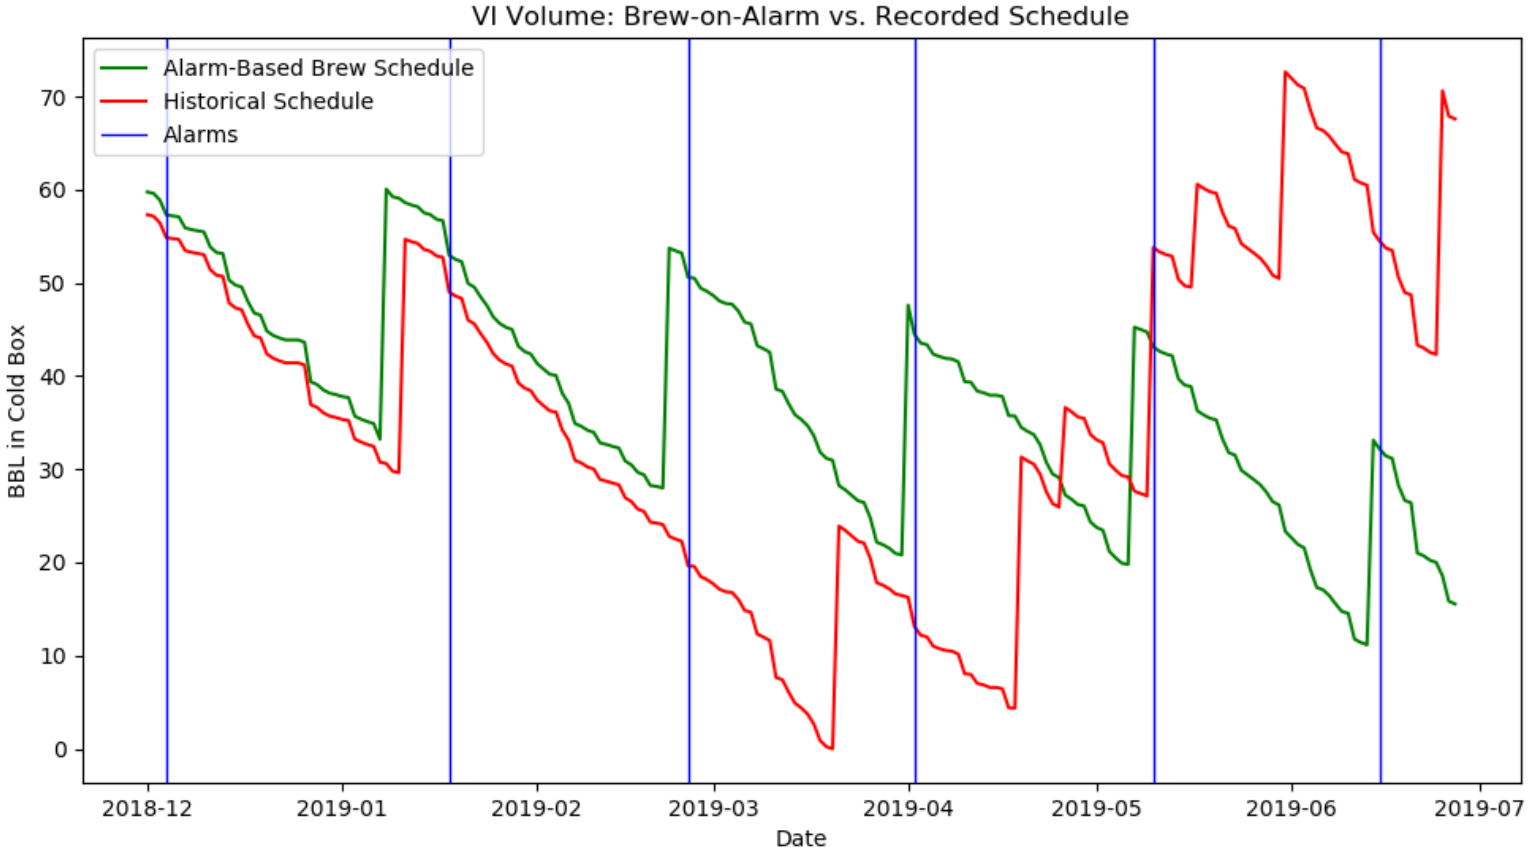

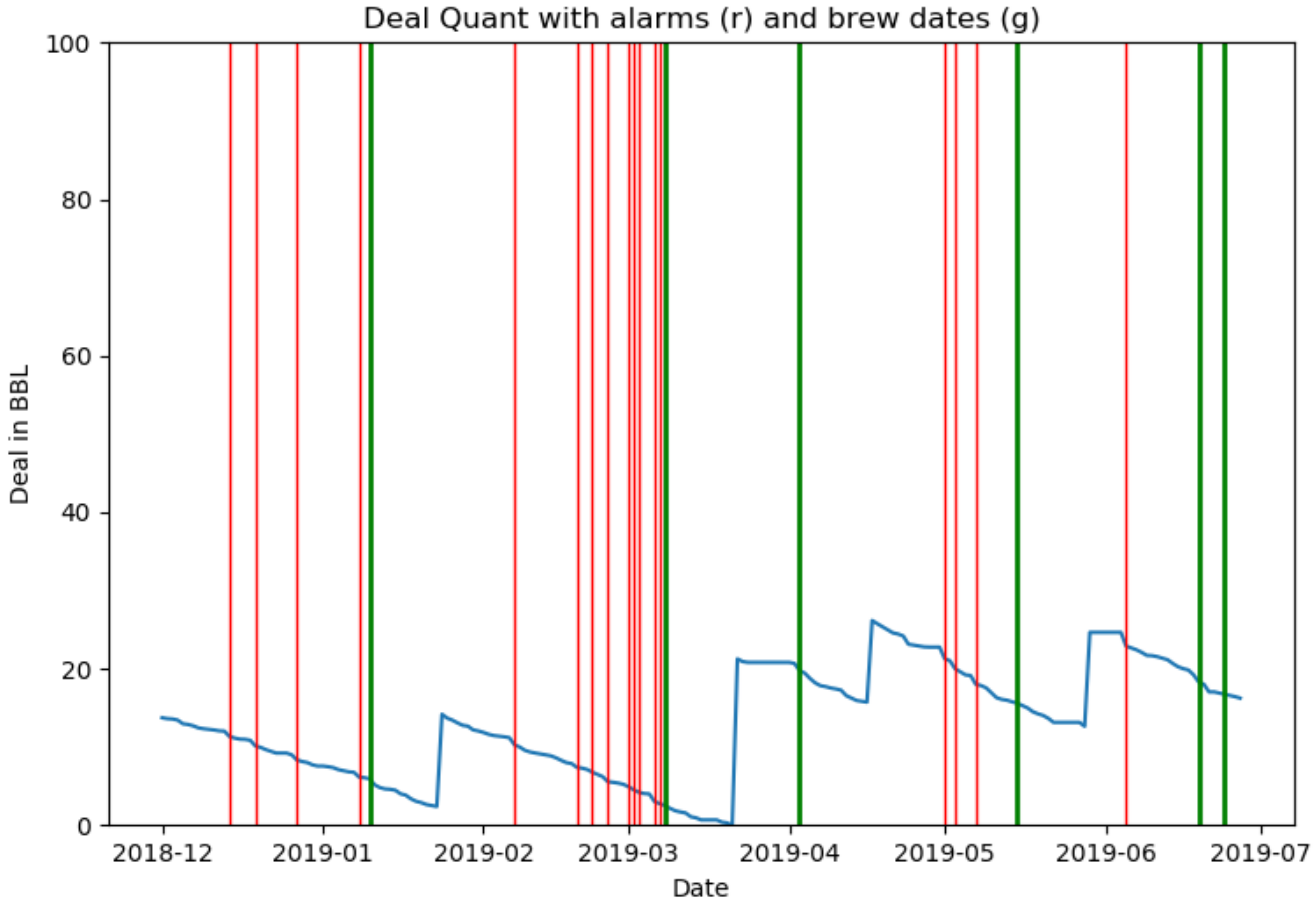

Why Me Worry?

So What if We Run Out of a Beer or Two?

- I can't imagine the owners at MacLeod were alone in thinking, hey, what's the big deal with tossing a few kegs of local, craft Light Lager on if ours runs out? Or letting people drink the Session or Double IPA if the flagship is out for a week?
- Here are a few thought experiments:
- Situation one: our flagship Hazy IPA, which was our most popular beer at, say, .350 BBL/day, drops, leaving our flagship IPA at .250 BBL/day to pick up the slack (and the sum, .600 BBL/day, was what we were selling).
- What are the odds that our customers now buy .600 BBL/day of *just* the standard IPA?
- Or heck, what happens if the standard IPA drops, leaving the Hazy?

Cont'd

- Situation two: looking at some numbers, let's say that the revenue per typical BBL is about $2,000 (that's basically true for $8, 14.5 oz "16 oz" pints at ~90% efficiency).
- Now let's say our crazy popular .300 BBL/day Flagship American Light Lager runs out and gets subbed for a .200 BBL/day Light Lager (which would be quite high for us).
- That's a drop of $200/day, or about $6,000 per month, in revenue from swapping out just one beer for another!
- Situation three: what if *both* things happen, and we replace our sole IPA with a local craft IPA? It'll cost more and, in our experience, sell less, so we'll experience a drop in revenue *and* an increase in costs.
- And this is actually better than situation one!

Two More Graphics

- Let’s look at one particular year’s revenue vs the number of can options available, in case the last few years hasn’t sold any of you on the idea that more cans = more revenue
- And let’s look at one possible use case for this beer planning structure (though something only possible with code)

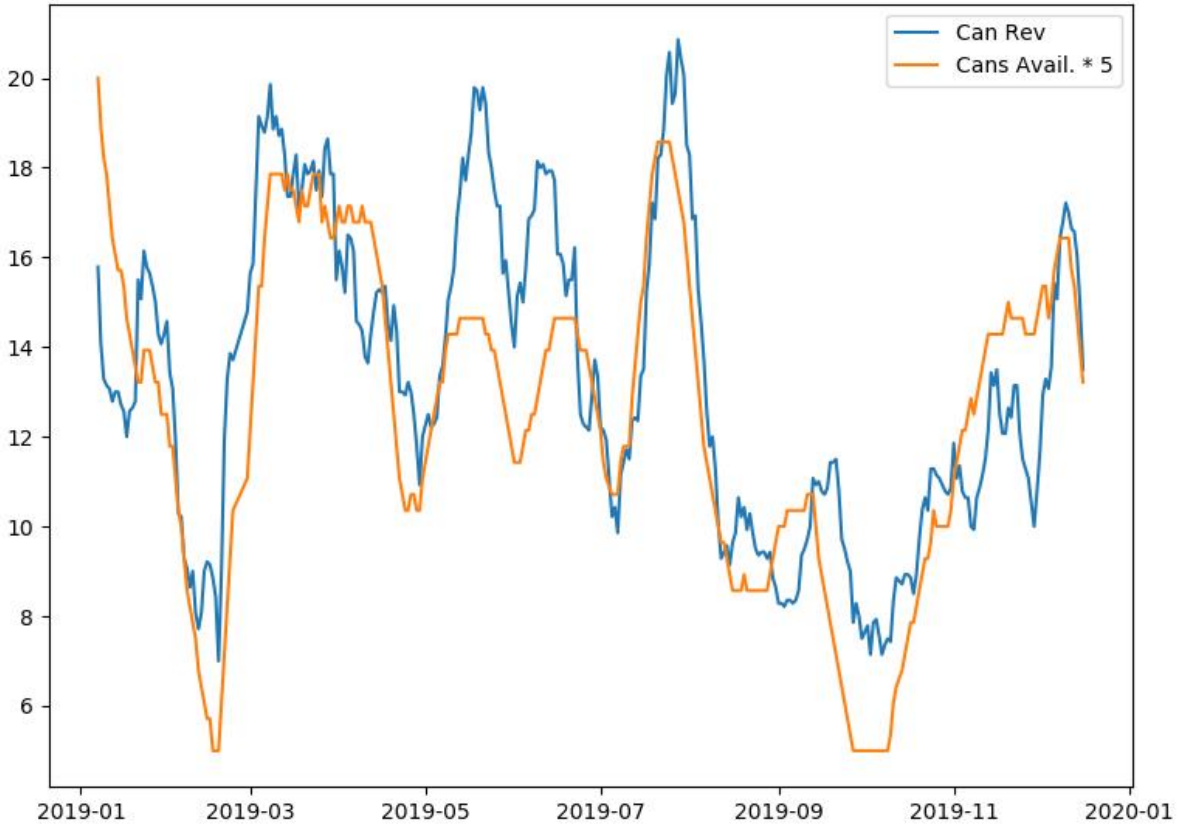

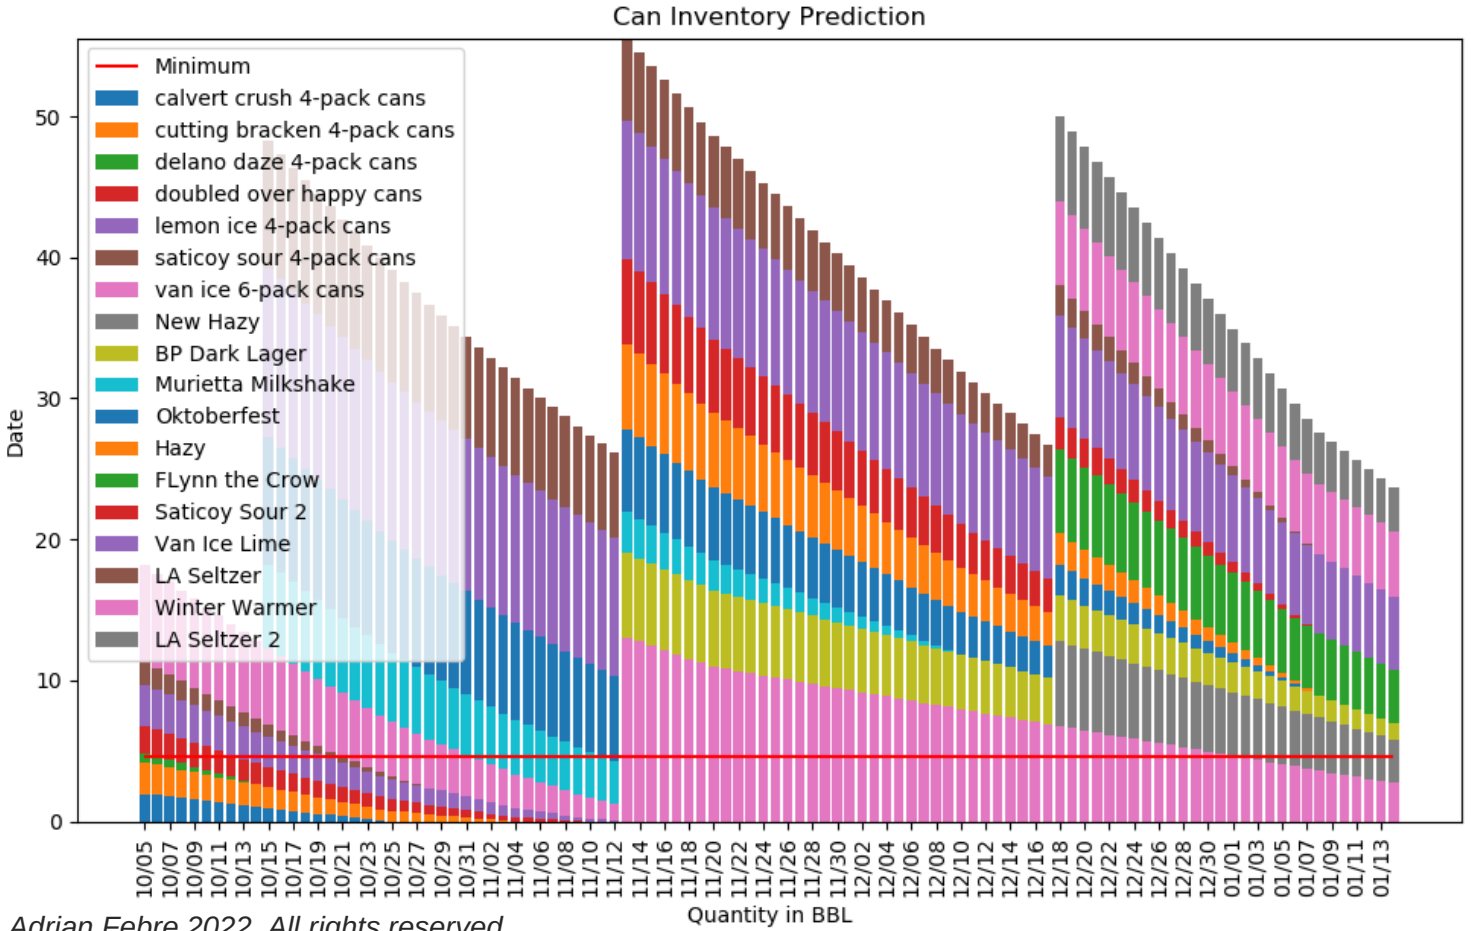

The Broad Solution

What's the Mathematical Issue With the Gut Check?

- You can think of inventory over time as a hill, with inventory as the height a given moment, and **demand as the slope**
- Getting the height is easy, it's getting a useful estimate of the height *and* the slope that's key to brew planning
- The "gut check" method that we once used relies on completely legitimate inventory data, which is good, but implicitly relies on a guess or an imperfect estimate of *demand* (in kegs per week, say), which is half the battle
- In order to beat the gut check approach, we need to get better estimates of beer demand (in BBL/day), and maybe even a sense of how much those estimates vary over time
- There are two sources for this number, which we'll now dive into: Inventory Data and Sales Data

The Data

Inventory Data

- Inventory data has a few big advantages: it's very clean, generally, and builds in things like draft loss (i.e. foam, dumped beer)
- But, for small breweries, the unit (a ½ BBL keg) is so big that we lose information
  - One keg per week is .071 BBL/day, but two kegs is obviously twice that - .143 BBL/day
  - If the "true" average rate is .080 BBL/day, the first number's close but the second one is way off
- Which is to say, inventory data is useful for estimating the demand of popular beers, but isn't as useful for less popular beers
- Though a note: it is a fantastic way to track cans, since cases are pretty small units
  - Because of this, for instance, comparing this data to can sale data could help you track "lost" cases

Sales Data

- On the other extreme, what if we take the unit size to basically zero, from BBLs to Pints? We get really beautiful, smooth data for demand
- While sales data *does* accomplish this, it's riddled with small errors
  - Categorical errors, like not putting a new can SKU in the "Cans" category
  - And Bartending errors, like comped beer and pouring technique/efficiency issues
- Still, it's a very powerful data source for slower beers, and for us it has turned out to be more than accurate enough, with some tweaking, to be used as our primary source of demand (BBL/day) data

So, How Do We Actually Use This?

- We'll talk about plugging these in below, but:
- Inventory data
  - We can update this information each week, and manually copy and paste, say, the number of kegs or cases of each SKU into a spreadsheet
  - Few breweries have so many SKUs that this is a massive headache, but automation is possible with the use of code
  - This is business as usual
- Sales data
  - This is a tougher nut, though there is a manual method that makes this granular data brute force-able

Brute-Forcing Sales Data

Pints to BBLs Calculator			
	One Unit (oz)	Number sold	Total volume
Pints	14.5	97	1406.5
Half-Pints	9	40	360
Crowlers	32	11	352
Growlers	64	3	192
Sixtels	661	1	661
Tasters	2.5	28	70
		Net Beer Volume, oz	3041.5
		Net Beer Volume, BBL	0.767
		Days in Period	14
		BBL/day	**0.055**

1) Pull up your sales data
2) Assign each unit (“Pint,” “Sixtel”) a volume in fluid oz, say, or mL
    a) Even grams, a proxy for mL
    b) Pop quiz - if 1 mL of water weighs 1 gram, what’s the mass of 1 mL of flat beer?
3) Multiply the volume of each item by the number sold, and sum all of those totals to get the number of oz, say, of each SKU sold in each period
4) Voila, you have BBL/day for any beer or SKU for any time period

The Metrics

Useful Metrics

- What's a metric?
  - A number that summarizes useful information about the data
  - For us, mostly BBLs, Dates, and Percentages
- Quantity on Hand
  - How many days will what I have last?
- The Point of No Return (PONR) Brew Date
  - When do I need to brew in order to *just* package a new batch when the last keg kicks?
- The Buffered Brew Date
  - When do I need to brew in order to package with some number of BBLs of beer left?
- Average Turnaround
  - How fast, in days, am I going through each beer?
- % of Capacity
  - If we call 100% the case when brewing a batch takes as long as selling it, what % am I at for each beer right now?

The Flow

Setup

1. Take inventory
2. Gather information on beer in tanks and their brew dates
3. Download or pull up sales data
    a. Either run code/upload the data to an app to automatically generate BBL/day, or use the Brute Force method
4. Enter the inventory data, volumes in tanks, brew dates, and demand (BBL/day) data into some sort of “setup” sheet

Punchline

1. Go over your quantity on hand fields, noting beers that you have a ton of, as well as beer that you're running low on, in terms of both BBLs and days
    a. It can help to look at last week's metrics, since a spike in demand (say, for a Irish Dry Stout on St Patrick's Day) can skew your numbers
2. Take a look at turnaround and capacity - are the turnarounds for any popular beers approaching their limits, i.e. the amount of days it takes to brew those beers? % of Capacity is the other side of this coin
3. Finally, read through the list of beers and their recommended brew dates. Buffered brew dates represent the "ideal" for most breweries, with PONR representing the do-or-die brew day
    a. In our experience, that PONR brew day is quite accurate, though if you don't count racked kegs in inventory, you have a built-in buffer - we usually have one and a half kegs "hidden" from inventory

Final Thoughts

What this did for us

- Implementing this beer planning system resulted in an **88% drop in beer outages**, calculated as the sum of days for which each core beer was out in six month before-and-after periods
  - I.e. being out of one core beer for ten days is the same as being out of two core beers for five days each - the sum of these in B was 12% of the sum in A
- We've been able to expand our offerings from about 10 taps and 0-3 can varieties to about 20 taps and 6-10 can varieties, with only **11 fermenters**

Inventory Efficiency vs Minimal Risk

- Given the PONR and Buffered brew dates, what happens when we have a choice between the two? Which is to say, if we have some breathing room and can brew early, which list's **order** should we follow?
- That's a great question, and it comes down to a decision between Inventory Efficiency and Minimal Risk
- By using the **PONR brew date**, we're brewing to replace beers right as they run out, which is inventory efficient but obviously risky
- The minimization of risk of the **Buffered brew date** strategy is less obvious, name aside, but it's easily seen with a slow beer. If we go through one keg in, say, two whole weeks for a certain beer, we'll have brewed a staggering two months early if we use a 2 BBL buffer target, which is safe (imagine the distro team suddenly selling three of your four remaining kegs), but inefficient - you're storing a full batch in your cold box for two extra months

What This Framework Doesn't Do

- This sheet doesn't address every single planning issue under the sun, so let's talk about a few
- My singular focus when starting as Analyst was making sure that our core beers dropped as seldom as possible, so any one-offs, or non-core SKUs in general, weren't incorporated into the sheet, directly
    - Granted, this is addressed by the inclusion of Categorical pseudo-SKUs like "Lager" and "Hazy," so this may not be a huge issue
- However, if you want to fill your taps with in-house beer, you need to ask a different question, which is how many taps you'll have over time (though the data used to work this out, demand and inventory, are the same), so for that you'll need a separate sheet
- Further, this sheet doesn't handle "complicated" beers well, like that flagship lager of ours. Truly popular lagers may require filling multiple tanks simultaneously, on a staggered schedule, and that requires its own (very similar) solution, which lives as a tab on my planning sheet
- This also doesn't address cost, which is yet another sheet

Cheers!

Adrian Febre
adrian@bits2bbls.com