How Bone Up Gets Made

Alrightyroo, so we've got a buncha recipes to share but first it probably makes sense for us to get into some of the nitty gritty details on how we do things around here - our approach to brewing is a lot more "one size fits all" than homebrewers tend to expect (as a general rule, homebrewers *love* making things more complicated than they need to be. I say this as a now-former homebrewer who used to be and/or still is the embodiment of that sentiment) but also that means we've got some very specific opinions on how things should be done.

Anyway, I'm gonna work off the baseline assumption that you have a handle on the basics of brewing, and I'm gonna try and use lots of section headers so you can pick and choose the parts you're interested in (I'm assuming nobody will try and exactly replicate the way we do things, but rather take our advice on some things and use their own system for the rest. There's no right way to brew, but there are lots of wrong ways to brew, et cetera), especially since this is gonna be equal parts instruction and manifesto. Also I'm just gonna start off with some bullet points because it turns out that me "being thorough" equates to "writing an entire novel."

THE TL;DR VERSION

In short, to emulate the Bone Up Style, you should:

- Use as thin a mash as possible (ours range from 2-3 pounds per gallon)
- Don't bother with water treatment (excepting a little citric to correct the pH)
- Mash as close to 149 as possible (single infusion)
- Pitch a 50/50 blend of Ardennes and London III (Wyeast numbers 3522 and 1318)
- Ferment at 75 degrees
- I think most of the other stuff we do is fairly standard (60min boil, 45min mash rest, probably some other stuff)
- But if you want Way Too Much Info on how we brew here, keep on keepin' on

THE LONG VERSION

A Note on Units

I don't know if this is common knowledge or not, but you'll frequently see me use the term "BBL" (pronounced "barrel"), which is the standard unit of measurement in the US brewing industry; 1 BBL is 31 gallons. I'm used to using it (and thus will frequently reference volumes in BBL), but also I hate it (because it's stupid). And since our brew length is only 3.5 BBL, we tend to use normal measurements a lot of the time anyway. The point here is that I'll be mixing and matching a bunch of different unit types, and it probably helps if you know what those units are.

Also of note, we use SG for gravity. Real breweries tend to use Plato and they sometimes get kind of sanctimonious about it, but I'd been homebrewing for over a decade when we started Bone Up and I didn't really see a compelling need to switch measurement systems (still don't, honestly). We tend to use shorthand notation for gravities in SG, so instead of 1.050 we'll just write 50, so if I use any two-digit gravity numbers here then hopefully you'll be able to figure out what I mean.

A Note on Recipe Volumes and Scaling and Substitutions

Our batch size targets a yield of 3.5BBL though our actual volume varies based on a wide variety of factors (kettle hops and dry hops decrease volume yield, adding fermentables during fermentation increases yield, sometimes our wort O2 is high and we have more yeast growth than usual and that nets us less beer, that sort of stuff).

Anyway, that's about 21x the size of your standard homebrew batch, which makes converting to homebrew scale kind of tricky, so I'm gonna be doing a lot of rounding and guessing when it comes to converting recipes - I intend to post our original recipes and then a suggested scaled-down version (and, if it's a beer that started as homebrew, I'll put up the original from Back In The Day).

Also of note, every brewery obviously has its own quirks and mash efficiency is one of the biggest things that's gonna change from one spot to another (our mash efficiency sucks, btw). I personally am of the mindset that for specialty grains it's more important to maintain the proportion of ingredients to batch size, rather than the proportion of ingredients to each other. So for example if it's a stout recipe with 20# of roasted barley in the original (making up 10% of the 200# in the recipe), a direct-ish scale to 5 gallons would be 9# of base malt and 1# of roasted barley. But if your mash efficiency somehow sucks worse than ours and you know you'll need 13# of base malt to hit that gravity, still stick with the 1# of roasted barley even though that throws off the base malt/roasted barley ratio. And we pretty much use the same base malt for everything, so even other “base” malts (Vienna, Munich, etc) tend to get treated like specialty grains when we do that calculation. Make sense? No? Cool, well that's the best I can do so if this seems like gibberish to you then just pretend I didn't say anything.

On a related note, our approach to kettle hops is not-dissimilar to how we approach specialty grains. Any late addition hops (i.e. whirlpool hops) will stay the same in terms of lbs/bbl regardless of changes in alpha acids from one crop year to the next; if the alphas on our late addition hops change, we compensate by changing the first wort addition to keep the IBUs similar (though honestly we’re at most moderately strict about keeping the same IBUs - we try to be within 10% of our target, but a change of 2-3 IBU in one direction or another isn’t really noticeable). I guess my broader point is that I’m operating on the assumption that you’ve got an at-least medium handle on how things behave in your own brewhouse (and/or that you have/use brewing software that will do the math for you). And if not, the likely worst-case scenario is that you’ll make something that’s close but not quite right, which really ain’t all that bad.

Speaking of consistency/accuracy, we’ve got a fair amount of leeway in terms of what our reported ABV is vs the actual ABV, specifically 0.3 percentage points in either direction (so if it says 6.0%, it’s actually somewhere between 5.7% and 6.3%). Maintaining consistency is really hard on a small manual brewhouse, but it turns out that hitting the exact same numbers every time is less important for perceived consistency than you might think, which is to say (again) that a lot of the stats for these recipes should be thought of as more of an “ish” kind of thing than a hard-and-fast requirement (not that we don’t do everything we can to get every number on-target. We do. But we can’t precisely control every variable that affects those numbers, so we do the best we can with what we have). Okay I think I’ve beaten this point to death by now. Let’s move on.

Oh wait, one more thing. When scaling recipes down I’ll be making a fair number of brand substitutions for specialty grains, since we use a lot of stuff that isn’t available at a homebrew level (for instance, we use Thomas Fawcett dark grains, but I don’t think homebrew shops tend to carry that so I’ll sub in Muntons). We’ve done a fair amount of experimenting with swapping out different brands of grains, so the substitutions won’t be made willy-nilly (for instance, Briess caramel malts are a pretty easy swap for Ireks crystal malts, but Proximity crystal is a completely different thing) and I’ll probably put in a note about the

substitution if I think it’s important. But in general, you’re gonna want to/have to tweak things to fit your system anyway, so brand swaps are just gonna be a part of that.

WATER

Water Chemistry

Water chemistry is *boring* so we decided to mostly not worry about it. All the water in the brewery goes through three filters (two sediment and one carbon, turns out there's lots of what appears to be sand in Everett's municipal water) before it reaches us, but it's very soft water with a pretty high pH so we'll just add a little citric to bring the pH down - usually 2.5 - 4.0 oz in the strike water (105 gal), and 1.5 - 3.0 oz in the sparge (65-75 gal).

Strike/Sparge Volumes

We use the same volume of strike water for every batch, mostly as an "it ain't broke, don't fix it" kind of thing (or, as we say by way of making excuses to ourselves for our own laziness, "it's farmhouse brewing"). Anyway we use 105 gallons for grain bill sizes that can range from under 200# to 350#, which to some extent means that we're getting more dextrins in our lower-ABV beers and less in the strong stuff (mash thickness doesn’t have a huge impact on dextrin profile, but it does more than nothing).

For reference, we have about 15 gallons (maybe a little more, but not more than 20) of space under our false bottom, and I don't think that's supposed to be factored in when figuring out grain:water ratios.

For sparge water, we use 65 gallons for 250# and under, and 75 for over 250#.

MASHING

Mash Temperature

This is a fun one for us. In general, we use the same strike water temperature for everything (159 degrees), with a theoretical target mash temperature of 149 (though in practice, weaker beers mash higher and stronger ones mash lower - we don't want session beers to be thin or strong beers to be chewy). We start to lose efficiency if we use that system for stronger beers, so we'll go up as high as like 161 or 162 if we're using 300# or more in a single mash, but in general it's one-size fits all.

Where this gets fun, though, is that we have a tall and skinny mash tun (which I very much do not recommend) which sits next to our direct-fire HLT, which has the heat on during most of the mash rest. Ultimately this means that we'll have a pretty broad temperature swing between the top of the mash tun and the bottom, sometimes it's as high as 10 degrees (especially once we've lost a few degrees at the top of the MT - we keep it lidded during the rest, but that only does so much). That all more or less evens out during vorlauf, so I call our system the "poor man's step mash" but having a big ol' temperature gradient in the mash tun is gonna have an impact that's hard to measure, so... good luck!

Mash pH

We're not real specific about our target mash pH - we theoretically aim for 5.2 - 5.4 but we tend to be on the higher end of that, and it's not uncommon for us to get a mash pH of like 5.55 and then need to add a bunch of citric to the sparge water to compensate.

A lot of the reasoning for this comes from the fact that our mash pH gets figured out by trial-and-error more than anything else. Our pH meter for some reason doesn't like taking the pH of very soft water (we've tried with distilled water too, to no avail), so it's hard to establish a baseline when figuring out citric additions (and there's no real point in adjusting the pH of the mash once you've mashed in, which is why we adjust the sparge - we really just want to keep it from getting too high). Combine that with the fact that many if not most of our beers are one-offs (and thus it's difficult to have a steady baseline for the grain's contribution to pH) and also the fact that our water pH changes seasonally (it also starts to change when we need to change the cartridges in our water filters) and the end result is that our real target is "don't let it get too high" and we just kinda roll with it.

It's also worth noting that we use the run-of-the-mill powdered citric that you can get at the grocery store. We need little enough adjustment that we never saw a point in getting the industrial stuff.

Mash Times
We rest for 45 minutes, vorlauf for 10 minutes, and then start transferring to the kettle (we used to go longer on both of those, but we did a bunch of A/B testing and reducing the times showed no difference). Sparging usually starts when the grain bed starts to peek out from under the liquid in the mash tun (usually about 10-15 minutes in) and ends when we run out of sparge water. If things are going smoothly, lautering (including sparging) takes about an hour (also we keep the kettle burner on this entire time, because it takes forever to bring a tall skinny kettle to a boil).

Sparging
Not a lot to be said about sparging, honestly. Our sparge water is usually in the low-to-mid 170's when we start, but we figure that getting run through pump/hoses and then sprayed over the grain bed should prevent it from getting things too hot. It was never a problem, so we never really messed around with it much.

BOILING, WHIRLPOOL, AND KNOCKOUT

Kettle Hop Additions
As a general rule, we use first wort hop additions instead of waiting until the boil starts to add our bittering hops. If you really get into the weeds with beer nerds, there's a lot to be said about whether one method or the other offers a "cleaner" bitterness and whether having the hops in for that little extra amount of time gives you IBUs (or if it causes polyphenols to bind to the isomerized alpha acids and precipitate out, giving you *less* bang for your buck IBU-wise) but the real reason we go with FWH is because we'll forget if we don't.

Other than FWH, the only time we add kettle hops is in the whirlpool. We dabbled with other kettle additions in the early days (we started in the days before NEIPA was a thing, so kettle hops still made some sense), but when we knock out we pump our wort through a strainer and if we use more than like 2# total in the kettle we tend to clog the bejesus out of the strainer so we pretty quickly started moving kettle additions into the dry hop.

Boiloff Rate
We boil off about 7 gallons during a 60min boil (down from about 120 gallons of wort to start off. We don't have a flowmeter so these numbers are very approximate), which equates to a roughly 6% increase in gravity during the boil (which, for about 85% of the beers we make, is about three points).

Whirlpool
There's a compelling argument to be made against whirlpooling with a pump when you're as small as we are, so we just hop up on a ladder and stir the crap out of the kettle until the pool is whirling enough. We go 15 minutes from flameout to starting knockout, regardless of how long we stir for.

Other Kettle Additions
Every brew gets antifoam when the kettle is full and again during whirlpool. If we're using carrageenan (we use Kick Micro-T, 22 per brew) that goes in 10 minutes before the end. If we're using servomyces as yeast nutrient that also goes in at 10, but if we're using zinc then that goes in at whirlpool (we've actually gotten better results from zinc than from servomyces, and zinc is much cheaper. But considering that we use 0.5 grams for a 3.5 BBL batch, it might be impractical at a homebrew scale).

Knockout, Dilution
Since our kettle is flat-bottomed we need to pump through a strainer on the way to the heat exchanger (so we don't clog the HX), then past an O2 stone and on to the fermentor (target O2 is theoretically 1 PPM per degree Plato, but our control over O2 flow rate is rudimentary so we kinda just aim for "more than ten"). We ferment pretty warm, so we're usually able to knock out within half hour/forty minutes.

Also of note, we realized that at the end of knockout the pump/hose/HX/strainer assembly would be full of theoretically-usable wort, so we rigged up a way to pump 180 degree water through that assembly to push all the wort out. We also realized that we could use that assembly to add a little volume to the fermentor, so we'll usually push an extra few gallons through (this essentially equates to a 4% dilution, so we've had to adjust some recipes).

FERMENTATION AND CELLAR SCHEDULE

Yeast
Our house yeast is actually a blend of two commercially-available strains: Ardennes and London III (technically we use BSI B-22 and A-18, but Wyeast 3522 and 1318 are the same strains as far as I know. Pretty sure the source breweries are Achouffe and Boddington's, respectively). And while we do harvest and repitch, we generally don't go more than 6-8 generations out, so commercially-available yeast should work pretty well.

Fermentation and Cellar Schedule
We ferment at 75 degrees across the board, with a soft crash to 65 once terminal gravity has been reached and the beer has been VDK- for 48 hours, then a couple days later we crash to 33 for about three days before transfer.

If we're doing a mid-fermentation dry hop (gotta get that biotransformation, bro) or adding fermentables during fermentation (usually fruit puree, sometimes varietal honey), we'll wait until the gravity is at least halfway to terminal which is pretty reliably 36-48 hours after pitching.

In general, our cellar schedule is as follows:

- Day 0: Brewday
- Day 1: Check to make sure the pH has dropped by at least 0.25 within 24 hours
- Day 2: Mid-fermentation dry hop or fruit addition, where applicable
- Day 4: Terminal gravity usually reached by this point
- Day 6: Drop yeast and trub from the fermentor cone

- Day 8: Turn down to 65
- Day 9: Drop yeast and trub from the cone, post-fermentation dry hop where applicable
- Day 10: Turn down to 33
- Day 12: Transfer to brite (not time sensitive, may happen after Day 12)

DRY HOPPING

Honestly we've tried a handful of dry hopping techniques and each one tends to be a mixed bag, so we'll bounce around a bit. We have top manway tanks, so the process itself is as simple as opening the FV lid and dropping hops in - while running CO2 through one of the ports in the top of the tank to keep O2 out - but that results in hell of O2 pickup kinda no matter what, so it isn't great. But all the dry-hopping equipment out there is expensive or convoluted to use or both, so we just kinda roll with the punches.

Also of note, we don't tend to dry hop at a rate of higher than 3 lbs/BBL. You start to see pretty strongly diminishing returns after that point, and we're not really chasing the hazebro crowd so we don't waste the hops.

Oh, and we learned the hard way that you absolutely should not spund hazy beers, that apparently causes them to clear right up. We had a whole damn summer of clear hazies until we stopped spunding.

Post-Fermentation Dry Hopping

Ideally we don't want hops getting mixed up with yeast or trub from the bottom of the tank, so once the beer is VDK negative we'll soft-crash it to 65 degrees for a day and then dump out all the solids that settle out. That's not gonna get everything out of the tank, of course, but it's been good enough to us. After that we tend to turn down within 24 hours because everything we've seen/read/heard suggests that longer contact times on hops (especially in a warm tank) are more likely to extract unfavorable flavors than favorable ones.

Rousing

Sometimes but not always if we want to get a little more bang for our buck from a post-fermentation dry hop, we'll rouse the tank during and/or after dry hopping by bubbling CO2 through the bottom of the fermentor (which is only advisable if you have a way to dump yeast/trub first - you don't wanna blow all that shit up into the tank). It'll usually get us a more pronounced hop character and it *usually* increases haze levels but also sometimes it makes the beer drop completely clear and there's no easy-to-discern cause for why that would be.

Mid-Fermentation Dry Hopping

Ideally we wait until the most active part of fermentation is over, but our threshold is once the beer has fermented at least halfway to terminal. In theory this provides the mixing action that we get from rousing, plus the active yeast will transform some compounds in the hops into other compounds. We don't actually have particularly strong feelings about this - it does get us some different flavors and the CO2 mixing up the hops is helpful, but also active fermentation strips out some aroma compounds - but really it's easier than rousing and not particularly better or worse so we'll do it fairly often (and if you want to see professional brewers get into a nerd fight, get them arguing about the merits of different dry hopping techniques).

MISCELLANEOUS

We've got some other stuff we do that might(?) not be commonplace, so here's some odds and ends. Most of this stuff stems from the fact that we *really* don't like adding things to the tank once knockout is done, because that's kinda asking for an accidental infection. We have our exceptions of course (dry hops being a big one), but overall the principle stands.

Ginger (and spruce tips, figs)
When we were coming up with a way to get ginger flavor into our beer, we reasoned that we should get the flavor infused into something that would then be more soluble in hot wort. In this case, we use honey: the relatively high water content of the ginger means that moisture is drawn out of the ginger and into the honey, and then we can dump the whole thing into the whirlpool and get pretty good ginger flavor into our wort (note: we haven't tried just putting ginger in the whirlpool, which may end up working just as well. But the honey sits for like three days, so we figure we're getting more flavor extraction. Plus the ginger also sits in the hot wort after all that). Anyway, we do this with spruce tips and figs sometimes too, and that also gives us good results, so maybe soaking things in honey is just a good way to extract flavors.

Vanilla
Working off of a similar principle to the ginger honey, whenever we use vanilla beans we first make vanilla sugar (made by slicing vanilla beans and cramming them into a bag of sugar for a few days). There are probably more effective ways to get vanilla flavor into beer, but we're not allowed to use alcohol-based tinctures and the sugar itself is very vanilla-forward so we said "it ain't broke, let's not fix it."

Fruit puree in the fermentor
This is pretty much the same thing as adding mid-fermentation dry hops, only it's giant bags of liquid glop instead of a small bucket of hops. Okay so maybe it's just the timing that's the same - add aseptic puree to fermentation once the beer is halfway to terminal and you'll preserve more fruit flavor than you would if you added the fruit to the kettle (and you won't have a bunch of fermentable sugar in your finished beer like all those breweries that keep accidentally making grenades).

Varietal Honey
On the few occasions we've used cool varietal honeys in our beers, we've added them around the same time we'd add fruit puree or mid-fermentation dry hops, only honey isn't exactly guaranteed to be aseptic so we came up with a really convoluted system to deal with it. Basically we dilute the honey and heat it to 180 degrees or so for an hour, then we put it into something sealed and let it cool overnight. The next day we'll pump it into the still-fermenting beer.

TBD

I've got a few other things I meant to add on here (which may or may not be of use to anyone), but this is long enough and it's taken me long enough to write that I'll try and tackle them at a later date (50/50 on whether that will actually happen). These topics include:

- Barrel aging
- Blending
- Brett beers
- Probably something else, let me know if I've missed anything