FOUR MYTHS OF BUNDLING

NOTE: This version of the paper is now out of date. You can find the latest version at: https://coda.io/@shishir/four-myths-of-bundling/

Over the past few years, my opinions on paid subscriptions and bundling have changed 180 degrees. And I've come to realize that some of my previous misconceptions are fairly commonplace. So here's a discussion of 4 myths that I've heard regularly... and the 4 theses I would replace them with.

Some quick definitions I'll use throughout this doc:

- *Consumer = purchaser of the bundle (though in some cases, a consumer could be a business, we'll simplify terminology to consumer for now)*
- *Provider = provider of a good within a bundle*
- *Bundler = the organization that bundles goods together and sells them to consumers*
- *Bundle = something a consumer purchases that includes access to multiple goods generally (though not always) from multiple providers. For shorthand purposes, we will assume all bundles in this doc are billed on a monthly subscription basis.*
- *Myth-maker = our personified adversary for this discussion - an individual who will vehemently argue that all of these 4 myths are true*

Summary

Here's the high level summary of the 4 Myths followed by the 4 Theses (and corollaries) that I would replace them with.

Myth 1. Bundling is bad for consumers (as well as providers)	**Thesis 1: When done well, bundling produces value for both consumers and providers by providing access for (and revenue from) CasualFans.**
Myth 2. Revenue from bundles should be allocated based on usage	**Thesis 2: The most fair way to distribute revenue to providers in a bundle is by Marginal Churn Contribution, not Usage.** • Corollary 2.1: WholesalePrice (ProductX in BundleY) = RetailPrice (BundleY) * MarginalChurnContribution (ProductX in Population of BundleY) • Corollary 2.2: The best bundling outcome will have WholesalePrice (ProductX in BundleY) = RetailPrice (ProductX) * SuperFan% (ProductX in Population of BundleY)
Myth 3. Bundles will always feel like a rip-off to consumers since they represent a lack of choice	**Thesis 3: A consumer C will see a bundle as being "a good deal" if: Price(BundleY) < Sum(RetailPrice(X)) for all Products X in BundleY for which C is a SuperFan** • Corollary 3.1: For a Consumer to properly value a bundle, there must be a transparent (and reasonable) a-la-carte price for each of the products in the bundle • Corollary 3.2: A consumer C will see a bundle as being "a good deal" if: Price(BundleY) < IntegrationValue(BundleY) + Sum(CasualFanDiscountValue(X)*RetailPrice(X)) for all Products X in BundleY that the consumer sees as being part of the bundle
Myth 4. The best bundles are narrow and have very similar products so they make sense to consumers.	**Thesis 4: The best bundle is one that minimizes SuperFan overlap and maximizes CasualFan overlap**

Myth 1. Bundling is bad for consumers (as well as providers)

In the video space, this is generally phrased by our Myth-maker as something like: "Comcast is screwing their customers by forcing them to purchase channels Y and Z, when all they really want is X." Add moreover "The provider of Channel X would love to be a la carte, but Comcast won't let them."

Consider this scenario. Imagine there are ten goods (X1, … X10) each delivered as a monthly subscription. We have a choice to deliver them each a-la-carte, or to produce a bundle across all of them.

Now let's divide the population for each good into 3 parts. Imagine that for each good, each prospective customer is one of these 3:

1. SuperFan: This is someone who would pay the a-la-carte price for the channel. This means that they are fairly far along the price elasticity curve for the good (perhaps to the inelastic point), and also that they have the activation energy to go out and seek

out the good and purchase it.
2. CasualFan: Someone who would value the good if they had access to it, but either aren't willing to pay the a-la-carte price for the good, or don't have the activation energy to seek it out, or both.
3. NonFan: Someone who will ascribe zero (or perhaps negative) value to having access to the good.

Now if we offered these 10 goods a-la-carte, then:

- The providers would only provide service to (and collect revenue from) their SuperFans, and
- Consumers (regardless of fan type) would only have access to goods for which they are a SuperFan

This clearly doesn't maximize value, as consumers are getting access to fewer goods than they might be interested in, and providers are only addressing part of their potential market.

As an example, take the UFC. Ask a random person if they are a fan of the UFC, and you will likely get a binary answer - they will either say they are a huge SuperFan and pay to watch the fights every Friday, or they will ask you what the letters UFC stand for. This is because UFC has very few CasualFans - which I attribute to the fact that the UFC has very little bundled distribution. So if you are a fan of the UFC, you are likely a SuperFan. By contrast, take the NFL which has some SuperFans (e.g. subscribers to DirecTV's Sunday Ticket) but also has tens of millions of CasualFans.

So our first thesis is...

Thesis 1: When done well, bundling produces value for both consumers and providers by providing access for (and revenue from) CasualFans.

An obvious question is "ok fine, but what does 'when done well' mean?". Please hold this question till we get to Myth #3, we need a few more constructs first.

Myth 2. Revenue from bundles should be allocated based on usage

Our Myth-maker lives in the Bay Area tech scene where Usage is seen as the Great Equalizer, so when asked the question "how should revenue from a bundle be distributed among providers", our Myth-maker quickly and energetically exclaims: "By Usage of course!!"

Consider these facts. ESPN and the History Channel get almost the same amount of usage - i.e. if you looked at their rating points, viewership share, etc, you'll find them to be very comparable. And yet if you were to look at their carriage fees (the amount of money the cable company pays in fees to each provider out of your monthly cable bill), there is a ~20x difference. History Channel can get from $0.20-$0.40 per month per subscriber, while ESPN collects $4-6 per month per subscriber.

So is this unfair? Shouldn't they be paid the same, based on Usage?

Thesis 2: The most fair way to distribute revenue to providers in a bundle is by Marginal Churn Contribution, not Usage.

I propose that there's an alternate valuation scheme here. Colloquially, it is referred to as "Anchor Value". I would describe it more precisely as "Marginal Churn Contribution" (or alternatively, Marginal SAC contribution). In other words, I think that the best way to distribute money to provider X within a bundle is to ask "if I were to remove X from the bundle, how many people would churn?". Similarly, you could ask the same question from a subscriber acquisition perspective ("If I were to add X to the bundle, how many new subscribers would I earn?") but for simplicity, I'll pretend they are interchangeable for now and focus on the churn-based model.

If you believe this framework, then the reason ESPN would receive ~20x the carriage fee of the History Channel is because removing ESPN from the bundle would theoretically cause roughly 20x the number of subscribers to churn as would removing History Channel from the bundle.

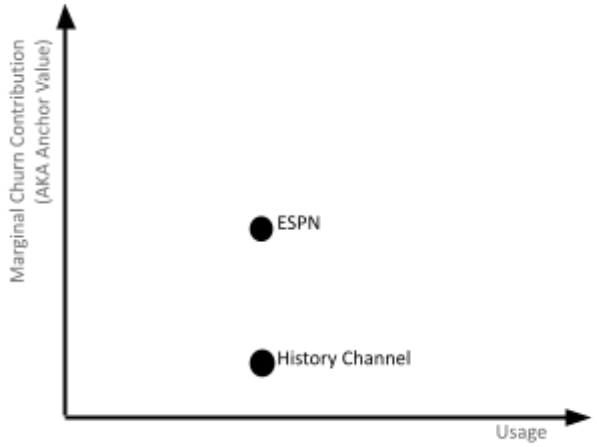

For example, let's imagine that if ESPN were removed from the cable bundle then 10% of subscribers would churn. If the average monthly cable bill is $50/month, and roughly 100M households subscribe to cable, this would mean 10M HHs canceling, or a loss of 10M * $600/yr = $6B in annual revenue at risk. So the cable companies agree to pay the $6B, but distribute it across all 100M HHs, resulting in a carriage fee of $60/yr or $5/month.

One way to calculate WholesalePrices

So not only do I believe that the correct way to distribute money in a bundle is *not* by Usage, I also believe that the value of a product in a bundle can be determined mathematically in a way that approximates the Marginal Churn Contribution. The general statement of this principle would be as follows:

> **Corollary 2.1:**
> **WholesalePrice (ProductX in BundleY) = RetailPrice (BundleY) * MarginalChurnContribution (ProductX in Population of BundleY)**

Here's the definition of some of those terms

- **WholesalePrice (ProductX in BundleY)** is how much revenue should be distributed from BundleY to the provider of ProductX per subscriber of BundleY
- **RetailPrice (BundleY)** is the price for the whole Bundle
- **MarginalChurnContribution (ProductX in Population of BundleY)** is the percentage of people from BundleY who would churn if ProductX were removed

WholesalePrice is an interesting concept. The idea is that every provider in the bundle gets paid a portion of the customer's bundle payment for _every_ customer. In some sense, we're distributing the SuperFan value of the bundle over the broader bundle population which includes both CasualFans and NonFans.

Putting this corollary in layman's terms, it says that if I'm going to distribute money from my bundle between products, I should run the exercise of removing that product from the bundle and seeing how many customers I lose, and establish a total amount of money I would lose. Then I should turn around and distribute that sum across the population of my bundle and pay the provider that amount per subscriber.

Another way to calculate WholesalePrices

So that's interesting and useful, but I think we can come up with an even more interesting and powerful corollary:

> **Corollary 2.2: The best bundling outcome will have**
> **WholesalePrice (ProductX in BundleY) = RetailPrice (ProductX) * SuperFan% (ProductX in Population of BundleY)**

Let's define the terms again

- **WholesalePrice (ProductX in BundleY)** is the same as above: how much revenue should be distributed from BundleY to the provider of ProductX per subscriber of BundleY
- **RetailPrice (ProductX)** is the a-la-carte price for ProductX. In this case, we'll further define this to mean the price at which SuperFans of ProductX will pay for ProductX if it were offered a-la-carte.
- **SuperFan% (ProductX in Population of BundleY)** is what % of the customers of BundleY are SuperFans of ProductX, when BundleY includes ProductX.

Now let's try an example. Imagine that I am trying to include a music service (ProductX) into my bundle (BundleY).

- Given that most music subscription services today are priced at ~$10/month, let's assume that the RetailPrice (MusicService) is $10.
- In round numbers, let's presume that ~10% of the population of BundleY subscribes to one of these music services. For now, we'll estimate that the SuperFan% of the MusicService within that population is 10% as well (this is important, we'll come back to it).
- Therefore, the formula says that the appropriate WholesalePrice (MusicService in BundleY) <= $10 * 10%. In other words, the WholesalePrice is $1.

This may seem counter-intuitive at first. How can I take a service that (successfully) charges $10/month and only give it $1/month per subscriber? Let's run through a hypothetical negotiation:

- First, the BundleY company will point out that although this is 1/10th the price that the MusicService is used to, the overall revenue collected will be the same since the new customer base is 10x as large.
- Then, the MusicService provider will retort that this is unfair because the BundleY company is providing the service to some CasualFans, who (by definition), ascribe some non-zero value to receiving the service, and they will argue that they should be compensated for that as well.
- And finally, In return, the BundleY company will argue that the MusicService no longer has to do any subscriber acquisition marketing, churn management, etc and therefore the price should actually be less than $1.
- And in some of these cases, they will agree and the negotiations will generally stabilize to the equation above: RetailPrice * SuperFan%.

In fact, I think this pattern of negotiation is roughly what each cable carriage dispute oscillates between and where it eventually settles.

Ok but I said that this holds true in the "best bundling outcome" - what does that mean?

Let's take two extremes:

1. **Fully Overlapped SuperFan base:** In this extreme, the existing population of BundleY has 100% overlap with the current ProductX SuperFan base
2. **Fully Distinct SuperFan base:** In the opposite extreme, there is zero overlap. In other words, there is no customer who is a SuperFan of ProductX who is an existing BundleY customer, and vice versa, there is no BundleY customer who is a SuperFan of ProductX

And just to make the math specific, let's assume that BundleY has 40M subscribers paying $10/month and ProductX has 10M subscribers paying $10/month.

Let's start with **Case #1: Fully overlapped SuperFan base**. So in this case, of the 40M BundleY customers, 10M of them are already SuperFans of ProductX and are paying $10/month for the service. There are two interesting perspectives:

- From the perspective of the provider of ProductX, this is a rough situation because if they bundle into BundleY they will likely lose 100% of their a-la-carte customers. So they expect to be fully compensated across the 40M customer base. So they will argue that they currently get $100M/month ($10/month * 10M subs) and so to retain that, they will want the same amount of money across the 40M BundleY subs, so they will want $2.50 / month. This is the same equation as Corollary 2.2 - the SuperFan% of ProductX in the BundleY population is going to be 25%, so they will want 25% * $10/month = $2.50/month. So far so good.
- But then we get to the perspective of the bundler. The bundler's thought process will be that after they add ProductX to the bundle, they are still going to have 40M subscribers, and so they'll be collecting the same amount of revenue. Where will this $100M/month payment come from? They could potentially raise prices - after all 25% of their customer base would still see a savings since they are currently paying $20/month for the two services, but what about the other 75%? This is tricky, and they will likely balk at this proposition.

Ok let's look at **Case #2: Fully Distinct SuperFan base**. Now in this case, none of the current 10M ProductX subs are part of the 40M BundleY subs. Let's look at the perspectives again, perhaps in reverse:

- The bundler will be very excited about this. They will reason that they can add ProductX to the BundleY and get 10M new customers. The existing BundleY customers will get ProductX "for free" and BundleY will now have 10M new customers. So BundleY will now net $100M more in revenue each month. Dividing that across the 50M subs, BundleY would be willing to pay $2/month in fees to ProductX. Another way to look at this is to look at Corollary 2.1 - the MarginalChurnContribution (or MarginalSAC contribution in this case) for ProductX in BundleY is presumably 20% - i.e. 20% of BundleY's overall customers (after ProductX is added) would churn if ProductX were removed. So the WholeSalePrice = the RetailPrice of the Bundle ($10) * the MarginalChurnContribution (20%) = $2/month.
- The provider's perspective is very similar to Case #1. They will likely lose all of their a-la-carte customers so they will expect to be compensated $100M/month. Given the new customer base will be 50M subscribers, they will be ok with $2/month in fees which happens to be the same as the RetailPrice of ProductX ($10) * the SuperFan% of ProductX in the new population (20%).

Another way to state this is that the bundle is "optimal" when Corollary 2.1 and Corollary 2.2 are equal: i.e.

RetailPrice (BundleY) * MarginalChurnContribution (ProductX in Population of BundleY)
= RetailPrice (ProductX) * SuperFan% (ProductX in Population of BundleY)

This is interesting because establishing SuperFan% is much easier than understanding MarginalChurnContribution. For the former, you can look at the a-la-carte subscribers as a proxy, whereas for the latter, you eventually have to remove the product from the bundle to get a proper estimate.

But perhaps even more interestingly, this leads naturally to a very counter-intuitive conclusion in Myth 4: that the best bundles are ones that minimize superfan overlap. But let's get through Myth 3 before going there.

Some things I'm leaving out for a separate discussion

For brevity, I'm leaving a few things out so I'll just list them briefly:

1. **How can one determine the correct RetailPrice and SuperFan% for a particular ProductX?** The theoretical best way is to launch the a-la-carte service, wait for it to hit a terminal # of subscribers, and then do the math. But that's often impractical. So thus far, I think this has mostly been done by negotiation, but I do think there are better ways to approximate these values today. However, I'll leave that to a separate discussion
2. **What about unstable SuperFan%?** The example above all assumes that the SuperFan% has stabilized. If the ProductX provider believes that they have 10M subs today but will have 20M next year, then the negotiation becomes trickier. For products like well-established cable networks, this seems reasonable, but for new products like a new music service, this is tougher. In the end this reduces to a similar question to #1 - how can one estimate SuperFan% with imperfect data. So separate discussion :)
3. **What if BundleY and ProductX have substantially different prices?** This is trickier as well since it's not as obvious that a-la-carte customers of ProductX will immediately purchase BundleY as a substitute. We'll explore this a bit in Myth 3, but there's some more framing to be done beyond that as well.

A closing example

Just to close this section, here's another real-world example: In January 2014, the WeatherChannel and DirecTV went through a contract dispute haggling over their carriage fee (AKA WholesalePrice in our constructs). The WeatherChannel put out a press release as a part of their dispute arguing that DirecTV would lose 1.6M subscribers without the WeatherChannel. The LA Times reported that the typical WeatherChannel carriage fee is $0.13 per subscriber per month, implying that the MarginalChurnContribution is ~0.1-0.2% (on a DirecTV ARPU of $50-100/mo). In other words, of the 20M DirecTV subs, all that is necessary to justify that carriage fee is for 20k-40k subscribers being WeatherChannel SuperFans and at risk of churn… much much less than the 1.6M than the WeatherChannel argued. This is an interesting case because the MCC is so small that it is very difficult to assess - this is no surprise since as bundles grow, the goal of the bundler is to add enough goods that any individual anchor would have very little impact on the bundle if it were to be removed. In doing so, it makes MCC very hard to measure, which leads to brinksmanship negotiations - i.e. sometimes ProductX actually has to be removed from BundleY to see the true MCC. One guess (just a guess) is that DirecTV took WeatherChannel dark, measured customer complaint volume, and then came back to the negotiating table with better data.

Myth 3. Bundles will always feel like a rip-off to consumers since they represent a lack of choice

Ok ok, so our Myth-maker agrees that MarginalChurnContribution makes more sense than Usage, and it makes some sense that bundling allows access for (and revenue from) CasualFans in addition to SuperFans. But our Myth-maker is resolute - these all sound like a bunch of frameworks for economist wonks and the average consumer won't see it this way. In particular, bundles represent a removal of choice - requiring someone to buy ProductX if they are purchasing BundleY - and removal of choice is always seen as a negative to customers, right?

Well there's a few clear counter-examples. For example, consider the McDonald's value meal. When a customer looks at Value Meal #1, their natural instinct is to do the quick addition: "how much would it cost to separately buy a Big Mac, fries, and a drink?" Then they think, "well I may not have wanted a drink, but I wanted a Big Mac and some fries, and I wouldn't really mind having a drink, so this looks like a good deal?"

A basic model for how consumers value bundles

That inner dialogue captures a very important concept for how consumers mentally process the value of a bundle, and leads to our next thesis.

Thesis 3: A consumer C will see a bundle as being "a good deal" if: Price(BundleY) < Sum(RetailPrice(X)) for all Products X in BundleY for which C is a SuperFan

Phrasing that differently, I believe that the way a consumer values a bundle is in the following steps:

- Examine the list of Products within BundleY
- Mark each Product for which I would anyways purchase that Product separately (i.e. Products for which I am a SuperFan)
- Look up the a-la-carte (aka RetailPrice) of each of those Products and add them up
- If that total is more than the cost of BundleY, then the bundle is a deal!

Consumers will often rephrase this sentiment as "I bought BundleY for ProductX1 and ProductX2, and I'm getting ProductX3 and ProductX4 for free!"

The importance of transparent RetailPrices

Once again, this leads to an important corollary:

> **Corollary 3.1: For a Consumer to properly value a bundle, there must be a transparent (and reasonable) a-la-carte price for each of the products in the bundle**

Taking our cable example, I think that the reason why consumers think that cable may be a bad deal is that they don't understand the true a-la-carte costs for each of the component products.

For example, take ESPN. Imagine that the SuperFan% of ESPN is 10% (it's difficult to get this research, but for a variety of reasons, it's hard to imagine it's significantly higher than that). This means that roughly 10% of cable customers would pay for ESPN at an inelastic price point. Applying Corollary 2.1, and assuming a carriage fee of $5/month for ESPN leads to a RetailPrice of ESPN of ~$50/month. Imagine if rather than thinking of your cable bundle as "buying 400 channels, one of which happens to be ESPN", these set of ESPN SuperFans could think of their $50/month cable bill as being for "buying ESPN for $50/month and then getting 399 channels for free". Repeat this exercise for the SuperFans of each product within the bundle, and you'll gradually recreate the current customer base of the cable subscription.

But what about CasualFan value?

But the close observer will notice that our formula is not yet complete. Here are a few other considerations:

- **CasualFan Value:** The consumer will ascribe *some* value to products in the bundle for which they are a CasualFan and not a SuperFan.
- **Negative Value goods:** In some cases, consumers could ascribe negative value to a particular product in a bundle. There is some good analysis of this phenomenon in Thinking, Fast and Slow where it was shown that in some cases, adding a good to a bundle actually reduced the value of the bundle. There are a few reasons why this can happen:
    - Goods that are undesirable to have access to - e.g. a kids content subscription that includes porn
    - Goods that are suboptimal relative to their alternatives - This is a more interesting case. For example, consider a subscription that includes a subscription to the WSJ when you prefer the NY Times. Even though you might think that a consumer should just value this portion of the bundle at zero, numerous studies have shown that consumers can sometimes ascribe negative value to these inclusions.
    - etc

We'll incorporate both of these concepts with a new multiplier in our formula:

- **CasualFanDiscountValue[ProductX]** = a constant from -1 to 1 that is multiplied by the RetailPrice of ProductX to determine the value that a specific CasualFan consumer puts on ProductX being in BundleY.

In order to keep our formula simpler, we can define:

- A Consumer is a SuperFan of ProductX if their CasualFanDiscountValue[ProductX] = 1

In other words, if you value a product's value in a bundle to be the same as it's a-la-carte RetailPrice, we'll count you as a SuperFan of that Product.

So our new formula is:

> **Corollary 3.2 (Initial): A consumer C will see a bundle as being "a good deal" if:**
> **Price(BundleY) < Sum(CasualFanDiscountValue(X)*RetailPrice(X)) for all Products X in BundleY**

Thinking about "Visible" and "Invisible" Products

Another concept to consider is the idea that a consumer is not an economist and they will not be able to do the above calculation exhaustively for every good in a bundle. This can lead to some interesting effects. For example, if a bundle contains a number of

products for which the consumer is a SuperFan, but the consumer first "sees" the products for which they are a CasualFan or NonFan, they may not see the bundle as valuable.

So we can adjust our formula to accommodate this concept as well.

> **Corollary 3.2 (Restated): A consumer C will see a bundle as being "a good deal" if:**
> **Price(BundleY) < Sum(CasualFanDiscountValue(X)*RetailPrice(X)) for all Products X in BundleY that the consumer sees as being part of the bundle**

While this may seem like a small adjustment, it is actually quite powerful. The reason is that in today's marketing world, one doesn't have to show the same offer to every potential customer - one can choose to highlight a subset of the products and introduce the others later. So as an extreme example, if you could properly determine the SuperFan products for a particular user, and just show them those, then I believe the consumer will value the bundle just on that basis.

This concept of how bundles are presented can be quite complex, as human beings don't generally behave as rational economic modelers. For example, the HBR article has some interesting examples where different models for presenting bundles caused overall perception of the bundle to be skewed.

Is there value in integration?

One last thing to include in our consumer value formula is the value of integration between services. After all, bundlers often provide value beyond just a purchase arrangement. For example, Netflix provides a browsing and recommendation experience across content from multiple providers, etc. Or as another example, Comcast provides you with a single piece of hardware for tuning to and recording content from multiple providers, etc. Similarly, Amazon Prime gives you a unified shipping and return experience across a wide set of providers.

So let's adjust our formula to capture this. The most accurate calculation for this would be combinatorial - we would take all products X1, … XN, limit them down the Visible set (see previous section), then compute every potential Combination of those products, and then calculate an IntegrationValue for every Combination.

For now, we'll summarize that calculation into:

- **IntegrationValue(BundleY)** = the value a particular consumer ascribes to the integration of the sets of products in BundleY that the consumer sees as being part of the bundle)

So now for our final, most complete version of the Consumer Value function

> **Corollary 3.2 (Final): A consumer C will see a bundle as being "a good deal" if:**
> **Price(BundleY) < IntegrationValue(BundleY) + Sum(CasualFanDiscountValue(X)*RetailPrice(X)) for all Products X in BundleY that the consumer sees as being part of the bundle**

Myth 4. The best bundles are narrow and have very similar products so they make sense to consumers.

Our Myth-maker is reeling a bit. Bundling can be good, money can be distributed fairly, AND consumers will sometimes see them as valuable. Fine fine, but intuitively, the best bundles must be narrow right? Bundles of similar products that "feel" like they go together?

Ah but the math says the opposite:

Thesis 4: The best bundle is one that minimizes SuperFan overlap and maximizes CasualFan overlap
This is especially counter-intuitive.

Let's start with an example.

- Let's ask our Myth-maker: "Please choose something to bundle with UFC."
- Our Myth-maker naturally says "how about a similar sport - perhaps boxing or football?"
- "Hmm, makes some sense, but why?"
- Myth-maker says, "well those are things that a UFC fan is likely to be interested in!".

But our thesis says the opposite! It suggests that the best bundle for the UFC is something like knitting, not boxing. Philosophically, this is because bundles are really about providing goods for CasualFans.

We can come at this a couple ways.

Let's start by returning to Myth #2. We saw that the "optimal" WholesalePrice was established when there was a completely distinct SuperFan base. This is when:

RetailPrice (BundleY) * MarginalChurnContribution (ProductX in Population of BundleY)
= RetailPrice (ProductX) * SuperFan% (ProductX in Population of BundleY)

In this case, the SuperFan base ProductX is being used to its maximum value by adding new customers to BundleY - in other words, MarginalChurnContribution is approaching SuperFan%.

Similarly, we can see this in Myth #3. We know that a consumer will see value in a bundle if:

Price(BundleY) < Sum(RetailPrice(X)) for all Products X in BundleY for which C is a SuperFan

Now if the Bundle is already valuable to C, then this trigger has already been achieved, so adding another product for which C is a SuperFan would be a waste of value. You are better off adding a product that someone else (ideally a current non-subscriber) would be a SuperFan of.

This diagram may help visualize. For example, it shows from the bundlers perspective that Product C is likely adding new customers to BundleY. From a consumer's perspective, if you are a SuperFan of A or B or C, the bundle is likely a good deal - e.g. if you're a SuperFan of A, then you "pay for A and get B and C for free".

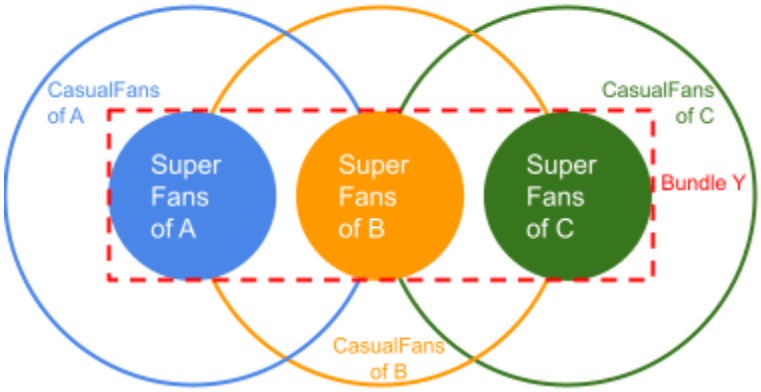

To take this to a natural extreme, I believe the absolute best world-wide bundle would be to design 7 billion unique perfect products for each of the 7 billion people on the planet, and then set it up in one big bundle. So everyone gets their perfect product, and also gets access to everyone else's perfect product for free. If there are no CasualFans (i.e. no one else's products are any good for me), then this is no worse than everyone just paying for their own product. But any amount of CasualFan overlap will justify the bundling exercise. Fun!

Extras

A few extra ideas that didn't make the cut:

- **Can Providers make independent decisions?** Another common (but more minor) Myth is that Providers can make independent decisions on being behind a paywall vs not. We experimented with this on YouTube - offering the ability for channels to offer a-la-carte subscriptions, and other platforms like Google Play and the Apple App Store have similar constructs. It initially seemed surprising that most creators were unwilling to put content behind an individual paywall, and yet were very excited to get deals that allowed them to distribute through a cable network like Nickelodeon or Disney. But using the theories above, this makes sense. The way they would phrase it is "if I put my channel behind a $5 paywall on YT, I'll at best keep 1% of my fans and it's hard for them to make up the value of the other 99%. But if I distribute on Disney, I don't give up any fans because everyone already subscribes to cable and has access to Disney." This leads to an important conclusion: **A Provider's decision to go "behind a paywall" is a function of (a) the composition of the distribution of that bundle and (b) the other participants in that bundle**
- **Other mechanisms for establishing MarginalChurnContribution? Other mechanisms for establishing SuperFan%?**
    - This got separated into it's own writeup: BUNDLING DRILLDOWN: ESTIMATING MCC and MSacC
- **What about bundles that depend on "Breakage" (customers not using the good) vs ones that don't?**
- **What about bundles where the cost of delivery is not zero - e.g. delivery services, etc?** Short version is that you have to add a "Minimum" to the WholesalePrice calculation which is the DeliveryCost of the goods. But it's also important to note that there

is often economies of scale in these delivery costs - e.g. shipping costs, warehousing costs, driver costs, etc often scale such that they become cheaper both as there are more subscribers (volume based) or when there are overlapping product needs (e.g. use just-in-time delivery services for urgent goods, and fill in empty spaces in delivery vehicles for goods that don't need to be delivered immediately). In any case the formula would be adjusted to something like:

WholesalePrice (ProductX in BundleY) = Minimum [RetailPrice (ProductX) * SuperFan% (ProductX in Population of BundleY), DeliveryCost (ProductX)]

Useful reading

- http://cdixon.org/2012/07/08/how-bundling-benefits-sellers-and-buyers/
- http://ebusiness.mit.edu/erik/Bundling%20Competition685305.pdf
- http://www.tvb.org/media/file/TV_Basics.pdf - Great stats on TV (HH share, etc)
- http://stratechery.com/2013/the-cord-cutting-fantasy/
- http://stratechery.com/2015/changing-unchanging-structure-tv/
- https://hbr.org/2014/06/how-to-succeed-in-business-by-bundling-and-unbundling/ - Jim Barksdale and Mark Andreessen on "there's only two ways to make money in business: One is to bundle; the other is unbundle"
- http://www.fiercecable.com/story/espn-will-cost-3630-sub-la-carte-world-priced-reach-analyst-says/2015-03-25
- https://app.box.com/s/c7fb19dzztqx7yzgrpbab58k7hxmczlf - Wharton Bundling Lecture - uses slightly different terminology (positive vs negative correlation) but fundamentally similar ideas

Table of Contents