Redacted

Redacted

Ex. No.
UPX0509
1:20-cv-03010-APM

Redacted

REDACTED FOR PUBLIC FILING & ABRIDGED

On Tue, Jan 1, 2019 at 7:43 PM Prabhakar Raghavan Redacted@google.com> wrote:

Thomas I read through this with some interest and am have a naive question that you might advise me to discuss in person: all of the discussion on advertisers' reactions to our pricing changes seem to presume that this is a 2-person game between the advertiser and google and - at some point on the pricing curve - advertisers could abandon us even if behooved them to stick with us. But isn't it really 3 players - the advertisers, us, and our competitors? I.e., the discussion seems insensitive to where else the advertiser could obtain traffic of similar quality and price. Why is it ok to ignore this - isn't the real point where an advertiser switches us out determined by the surround market?

On Sun, Dec 30, 2018 at 7:54 PM Thomas Iljic (via Google Docs) Redacted @google.com> wrote:

Thomas Iljic has invited you to **comment on** the following document:

How should AQ think about Pricing?

2018 Holistic Pricing Context Doc

Redacted

# How should AQ think about Pricing?

Jan 2018 - Redacted @

**Note:** This was reviewed and brought to resolution (preferred path) in AQER on 2/1 and 2/8. The outcome of these investigations, 6 months later, were then summarized in August here.

## I - Strategic considerations

**Pricing today**

If we think of *pricing* as the action of defining the cost our advertisers' pay per unit of value (e.g. cost per click or cost per conversion), then pricing on google.com happens today in small touches all throughout the year.

Most AQ launches, in addition to trying to improve the value for users and advertisers (e.g quality improvements driving better conversions, new formats driving higher CTRs, etc) have some CPC or CPA component. We'll call this *"incidental"* pricing. We also directly affect pricing through tunings of our auction mechanisms, in general through the three levers that are format pricing, squashing or reserves. We'll call this *"intentional"* pricing.

As a result pricing on google.com is happening in a semi-controlled, semi-organic way that might not be ideal: We worry about how these changes stack up over time for individual advertisers. We worry that this might eventually - and unwillingly - lead us to a state where advertisers are unhappy with their performance. And we worry this might lead them at some point to modify their spend on google.com, possibly reallocating budget to other channels on timeframes we can't track (e.g. quarterly, annually). How can we assess the cumulative effects of our launches?

A side effect of these worries is a recent discomfort among leadership with the idea of intentional pricing launches. Historically we've aimed at maximizing long term value for both advertisers and google through intentional exploration. Redacted

Redacted

**Types of risks at hand**

When it comes to pricing there are two broad categories of risk we should consider:

(Type A) - Advertiser lowering their bids or modifying other settings (e.g. use of match types, campaign pausing, …) to get back to a given ROI, leading to less revenue for Google than the initial impact hinted to.

(Type B) - Advertiser sentiment and perception of the product's performance degrading significantly enough that it leads to spend shifts away from Google (clear loss), and arguably in a more permanent way.

The first is a risk we can likely manage: advertisers actions are a way to tell us that we are moving away from what makes sense for them, and the lost revenue is the cost of our exploration. Incremental launches and monitoring should help us manage that risk.

The second is more problematic: the responses could be much more drastic and hard to recover from. They are also likely to happen on time frames (quarterly, yearly) on which we don't have much visibility. We need better tools to think through Type (B) risks in the context of launches stacking up over time.

**Advertiser ceilings**

It's reasonable to imagine that at any given point in time there is some price or ROI ceiling above which our product doesn't make sense anymore for an advertiser. Arguably this ceiling changes over time based with the macro-economic context, inflation, etc.

Whether or not the advertiser has an exact sense of where that ceiling stands at a given point in time (= CPC or CPA value above which marketing objectives can't be met), if we bust past this ceiling or even just give the impressions that we might, we can expect adverse advertiser reactions.

If we want to limit Type (B) risks we need to have pricing principles across *all* our major launches. In particular, special attention should be payed to cumulative CPC changes that are the largest source of escalations and articles complaining about performance.

**Risks vs Types of launches**

It's useful to breakdown pricing launches into the changes they have on the click / cost curves. A launch can be thought of as a combination of two changes:

1. Changes that move the advertisers along the click cost curve: these can be thought of as ROI changes where we change the value given while prices are kept *relatively* stable (e.g. the price of a hundred clicks remains a dollar).

2. Changes that modify the click cost curve (e.g. move it up or down, stretch it): the prices for the same number of clicks changes.

Redacted

But it's important to realize this doesn't fully alleviate Type B risk. It's not because advertisers could do a certain action that they will realize they can, or should, do so. The overall sentiment of the product might still be affected by the change, especially as it accumulates. In other words: regardless of nature of the pricing - incidental or intentional - we are still exposed to both type A and type B risks. Even if we decided to not to do pricing exploration, we need a framework to think about cumulative effects.

We think of this similarly to how we think about blindness on the user side: Redacted Redacted allow us to capture the long term effects of changing prices, similarly to how we measure changes in user sightedness Redacted In the case of user blindness, a macro model allows us today to vet long term effects of small changes and helps us decide whether we want to tune individual launches. In the case of advertiser sensitivity to prices we don't have a macro model yet, in other words we can't really evaluate long term implications of small price changes. By controlling the pricing environment and running principled exploration we pave the way there while making sure we're not getting into a bad state.

## II - Paths forward

So what are the options for Ads Quality? Redacted Redacted

Redacted

## Main document changes and comments

| Page 1: Comment [1] | Thomas Iljic | 1/29/2018 7:27:00 PM |
|---|---|---|

Omkar input: having RPM is generally viewed as something that's desired for launches. The drift in prices is also somewhat reinforced by our current approval process.

| Page 1: Comment [2] | Thomas Iljic | 1/29/2018 7:28:00 PM |
|---|---|---|

Omkar input: and also other things changing (advertisers coming in, auction pressure). Even if we did nothing prices would change.

| Page 1: Comment [3] | Lars Hirsch | 9/6/2019 6:45:00 PM |
|---|---|---|

How do we track budgets? Is there a dashboard?

| Page 1: Comment [4] | Thomas Iljic | 9/9/2019 4:46:00 PM |
|---|---|---|

Redacted

| Page 1: Comment [5] | Lars Hirsch | 9/6/2019 6:44:00 PM |
|---|---|---|

How do we know how much we "can" capture without advertisers diverting budgets? What is a "fair" price?

| Page 1: Comment [6] | Thomas Iljic | 9/9/2019 4:50:00 PM |
|---|---|---|

Redacted

(b): simplifying a lot, we don't. The HP effort is focused on stabilizing the current CCC under the observation that it's working OK for our advertisers, but it doesn't make statements about what/where the ideal curve (the ideal prices) should be. There's another line of work started this year, Value Based Pricing which goes after leveraging better understanding of value to set better reserves (that goes more in this latter direction).

| Page 1: Comment [7] | Thomas Iljic | 9/9/2019 5:01:00 PM |
|---|---|---|

Btw, "fairness" in pricing is a terminology I've been pushing away from because it's not concrete enough / there's a lot of subjectivity to the notion. Setting appropriate reserves however (e.g. that make sure we don't let ads go for pennies when we know there are worth at least $X...) is I think an interesting angle we're currently pursuing with Value Based pricing.

| Page 2: Comment [8] | Lars Hirsch | 9/6/2019 6:50:00 PM |
|---|---|---|

Or changing their mix of keywords/products/channels to optimize ROI given observed change in performance?

| Page 2: Comment [9] | Thomas Iljic | 9/9/2019 5:04:00 PM |
|---|---|---|

Absolutely. And there's a question of time frame here that you'll see come up in the following docs I sent you.

Redacted

All of that will surface in the Holistic Takeaways document :-)

| Page 2: Comment [10] | Lars Hirsch | 9/6/2019 6:46:00 PM |
|---|---|---|

How do we track advertiser actions / advertiser responses?

**Page 2: Comment [11]** **Thomas Iljic** **9/9/2019 5:05:00 PM**

Redacted

**Page 2: Comment [12]** **Jeremy Kubica** **1/31/2018 4:08:00 PM**

For type B risk, I'm not convinced there is a predefined ceiling for all advertisers. I would expect that some percentage of advertisers have anchored on their current metrics (ROI, click volume, etc.) and will consider significant deviations from these loses even if the new ROI point makes sense from a rational point of view.

Or, put another way, the future ceiling for an advertiser to reduce spend on the network might depend on the path of pricing changes we took to get there.

**Page 2: Comment [13]** **Jeremy Kubica** **1/31/2018 3:52:00 PM**

I think this makes sense given a perfectly rational advertiser, but breaks down when considering advertiser's perception of transactional utility. Even if they are below their economically optimal ceiling, they may see changes in overall pricing as a loss.

**Page 2: Comment [14]** **Jeremy Kubica** **1/31/2018 3:58:00 PM**

In other words, I think the type B risk is pretty significant.

**Page 2: Comment [15]** **Adam Juda** **1/30/2018 3:35:00 AM**

I wonder how much value plays into this in reality. For example, I think there is a belief that when auctions become thicker and increase prices, that might be safer than raising the prices in a thin auction, because in the former we are likely closing in on the "actual" value of the click, whereas in the latter prices might exceed the true value.

**Page 2: Comment [16]** **Chris Monkman** **1/30/2018 4:51:00 AM**

why does value matter for this mechanical division of how pricing changes impact advertisers?
competition (thick auctions) is an instance of case 2.
isn't value more relevant to Type A vs B risk quantification?

**Page 2: Comment [17]** **Adam Juda** **1/30/2018 4:01:00 PM**

In my mind, if the CPC after a change is still _way_ lower than my value per click, do I really care if the launch corresponded to a movement along the existing click-cost curve vs. the curve itself moving? I'd still be highly profitable after the change, so why does this taxonomy matter in such an instance?

Also, if my prices exceed my value, I might bail on advertising on that thing vs. lowering my bids, which feels like B risk. Otherwise, if only A risk, I question how everything is not always A (as usually there is some lower bid that would make ROI attractive, just that volume could be near absent).

**Page 2: Comment [18]** **Omkar Muralidharan** **1/30/2018 4:50:00 PM**

Is another way to put this, that different mechanisms for price increase can have very different risks for exceeding value? If we see a price increase from more competition, it feels safer than a big price increase from reserves we train on our end, since in the first case we know that someone in the world is willing to pay that much.

**Page 2: Comment [19]** **Chris Monkman** **1/30/2018 5:25:00 PM**

For folks who are currently paying way less than their value, I'd expect them to be volume constrained (be at the highest point on the click-cost curve). So, launches that move them along the curve (by extending the curve with more possible clicks) feel very safe since I'll sell strictly more or the same clicks as today. If I move the curve up, there may now be a volume that is unprofitable, so I could sell fewer clicks than today, and could even be revenue negative post response. I think the distinction in terms of curve changes does help to think about the kinds of risks we are creating.

I basically believe that advertisers are primarily going to focus on the increased CPC, and the secondary the lost volume. Class (2) does create more risk for volume, but they both cause CPC which is the primary issue. So, while the "TANK principle" may seem a bit safer, I think we shouldn't believe it is that safe, and instead spend the CPC increase in the most efficient way possible.

**Page 2: Comment [20]** **Adam Juda** **1/30/2018 7:49:00 PM**

> If we see a price increase from more competition, it feels safer than a big price increase from reserves we train on our end, since in the first case we know that someone in the world is willing to pay that much.

At the moment the above feels so to me.

> while the "TANK principle" may seem a bit safer, I think we shouldn't believe it is that safe, and instead spend the CPC increase in the most efficient way possible.

I agree that it's not safety guaranteeing.

**Page 2: Comment [21]** **Adam Juda** **1/30/2018 3:33:00 AM**

Given that the true click-cost curve is somewhat mythical, should I just be thinking about this as any change where ROI and conversion volume move in opposite directions? By contrast, there could be other launches where both don't move in opposite direction?

**Page 2: Comment [22]** **Chris Monkman** **1/30/2018 4:55:00 AM**

while saying "launches where CPD is down but Conversions are up" is a fair summary of this, I think the click cost curve is a useful way to think about these things.

the click cost curve exists. it's not exactly smooth or convex, and certainly a large number of advertisers are not at optimal points on the curve, but I think it is something advertisers observe as they change their bids.

Redacted

**Page 2: Comment [23]** **Thomas Iljic** **1/30/2018 5:16:00 PM**

I think bid landscapes, albeit smoothed out versions of click cost curves, gives most of our advertisers a sense for what those are. But I fully agree with Chris' point that it's likely not enough for them to really reason through what operating point they could *get back to*, rather they use it to see where they could go volume wise. That's why Type B risk is I think still high for such changes.

**Page 2: Comment [24]** **Adam Juda** **1/30/2018 5:25:00 PM**

Ok. Then who should be thinking about launches this way (given that reasoning through a curve is tricky for advertisers)? Googlers? Advertisers?

**Page 2: Comment [25]** **Thomas Iljic** **1/30/2018 5:27:00 PM**

I think it helps because if they do not realize it and escalate we, and AMs, can help them get back there without changing our systems. There are options open to reach that operating point they like. From that perspective there is a milder risk.

**Page 2: Comment [26]** **Adam Juda** **1/30/2018 8:35:00 PM**

I think this plays nicely into my request elsewhere around developing comms for Sales et al. to know how to respond. I think most would not know when and how to respond to situations where a simply bid change could get someone back to where they once were.

**Page 3: Comment [27]** **Adam Juda** **1/30/2018 3:34:00 AM**

Any thoughts on when we know we have constructed the "right" click-cost curve to want to preserve it?

**Page 3: Comment [28]** **Chris Monkman** **1/30/2018 4:58:00 AM**

I personally see no way for us to articulate what the best end state is. We can only say whether a change in the curve is good or bad based on stickage and risk.

**Page 3: Comment [29]** **Omkar Muralidharan** **1/30/2018 3:15:00 PM**

I think the right curve, in aggregate, is locally revenue maximizing (similar to how the right coverage means that the marginal impression is about long term revenue neutral)

**Page 3: Comment [30]** **Adam Juda** **1/30/2018 9:10:00 PM**

So I buy from a change-management perspective that being able to communicate to an advertiser that they can return back to where they were has value. I guess for me the open question then is, "When might we recommend this action vs. not do so?"

**Page 3: Comment [31]** **Lars Hirsch** **9/6/2019 7:40:00 PM**

Are there impacts to incentive compatibility and value-maximization properties of the auction that we should care about, from this tuning?

**Page 3: Comment [32]** **Adam Juda** **9/6/2019 8:29:00 PM**

I'll let Thomas speak to the value-maximization question, as that part seems harder! On IC, the approach we've taken lately with launches Redacted

Redacted

**Page 3: Comment [33]** **Omkar Muralidharan** **1/30/2018 4:51:00 PM**

as discussed offline - I think it helps a lot. it's hard to imagine an advertiser just storming out without at least trying some bid changes, and if they do in (1) they will find their old happy point. but agree that it doesn't eliminate the risk.

**Page 3: Comment [34]** **Jeremy Kubica** **1/31/2018 4:00:00 PM**

Do we have metrics on the distribution of how often non-autobidding advertisers adjust their bids? I feel like your point above should be the case (that advertisers will try different points), but it would be nice to see if this happened in practice.

**Page 3: Comment [35]** **Omkar Muralidharan** **1/31/2018 6:23:00 PM**

I don't think we have sensible metrics yet, but are trying to get them (not sure if we will succeed, change metrics are hard).

Redacted

Redacted Redacted @google.com who had a lot more conversations with them than I did

**Page 3: Comment [36]** **Lars Hirsch** **9/6/2019 7:37:00 PM**

Are we investing in building long-term models? If not, should we?

**Page 3: Comment [37]** **Thomas Iljic** **9/9/2019 4:44:00 PM**

Redacted

**Page 3: Comment [38]** **Adam Juda** **1/30/2018 3:37:00 AM**

What are the mechanisms by which you envision the ceilings being past in this scenario?

**Page 3: Comment [39]** **Thomas Iljic** **1/30/2018 6:19:00 PM**

Redacted

**Page 3: Comment [40]** **Adam Juda** **1/30/2018 9:12:00 PM**

As discussed just now, I'm all for monitoring this better, as I see no reason for us being in the dark. What's less clear to me is how to define a ceiling, and whether any action surely need be taken when such a ceiling is reached.

**Page 4: Comment [41]** **Thomas Iljic** **1/29/2018 7:41:00 PM**

Omkar input: we could explicitely make the analogy to blindness.
Also point out we're not aiming at getting one big bang launch but centralizing decision making.

**Page 4: Comment [42]** **Thomas Iljic** **1/30/2018 12:46:00 AM**

Sunita input: What leads to this belief? --> E.g. hinted to by queries in the same space having drastically different prices because of differences in competition, or responses to Expensive/Free formats AE, etc.

**Page 4: Comment [43]** **Chris Monkman** **1/30/2018 4:59:00 AM**

Redacted

**Page 4: Comment [44]** **Jean-François Crespo** **2/1/2018 1:35:00 AM**

Redacted

**Page 4: Comment [45]** **Ankur Jain** **2/1/2018 1:46:00 AM**

+1. We may be blinded to the cumulative pricing advertiser response in short-term but eventually should notice it on our holistic pricing dashboards, perception surveys and pricing explorations.

| **Page 4: Comment [46]** | **Omkar Muralidharan** | **2/1/2018 1:59:00 AM** |
|---|---|---|

I think there are few important differences

Redacted

-- just for power to track changes we need bigger things

| **Page 5: Comment [47]** | **Jean-François Crespo** | **2/1/2018 1:37:00 AM** |
|---|---|---|

How do we know what these margins are?

| **Page 5: Comment [48]** | **Thomas Iljic** | **2/2/2018 2:12:00 AM** |
|---|---|---|

For example through financial reports, consulting studies about different industries, etc.

| **Page 5: Comment [49]** | **Jeremy Kubica** | **1/31/2018 4:12:00 PM** |
|---|---|---|

I think this is an important point. One thing we have consistently said about advertising on Google is that it is a good value compared to other forms of advertising. I expect that keeping this differential in advertiser ROI is a very good thing to the long term health of our product.

| **Page 5: Comment [50]** | **Jeremy Kubica** | **1/31/2018 4:12:00 PM** |
|---|---|---|

+1 to this

| **Page 5: Comment [51]** | **Lars Hirsch** | **9/9/2019 6:47:00 PM** |
|---|---|---|

Is there also a programmatic way of detecting this, rather than relying only on interviews?

| **Page 5: Comment [52]** | **Thomas Iljic** | **9/9/2019 7:35:00 PM** |
|---|---|---|

Redacted

**Header and footer changes**

**Text Box changes**

**Header and footer text box changes**

**Footnote changes**

**Endnote changes**