CONFIANT

Ad Quality Report Q2 2018

NEW YORK - Sep 27, 2018 - Confiant, the cyber security company that protects the reputation, revenues and resources of publishers and platforms, releases ad quality and security stats.

Synopsis

Calling all Publishers: how much confidence do you have in the advertising demand that is coming from Real-Time Bidding (RTB) through your demand partners, SSPs and Exchanges? Malvertising redirects and video ads fraudulently positioned inside of display ad placements are a known phenomena in the industry, but how common are they really? What should you expect to see coming from the top tier Exchanges?

Confiant is a cyber security company that protects the reputation, revenues, and resources of publishers and platforms, with real time ad verification software. In Q2 of 2018, the company monitored over 60 billion programmatic advertising impressions. Leveraging its critical position as the first real time ad quality verification vendor in the industry, Confiant has captured unique insights into the quality of ads served by the programmatic marketplace. These insights benchmark the actual volume of malicious and misrepresented ads that attack publishers and their users every day. With this report, Confiant seeks to shed light on the state of demand quality in the market, and drive forward the discussion about how to clean up the ecosystem.

“We are used to seeing ad traffic quality reports supply valuable information to media buyers around what to expect from the DSPs and publishers they work with,” Louis-David Mangin, the CEO and Co-Founder of Confiant said. “We would like to equip publishers with information that could help them plan better, know what to expect from their demand partners and also possibly improve the status quo, pushing back on misrepresented inventory (e.g. IBV) and other issues.”

Methodology

Confiant analyzed more than 60 billion programmatic advertising impressions across over 2,200 sites, and over 1.5 trillion ad requests across multiple exchanges from April to June 2018, to compile the research contained in this report. Confiant utilizes its patent-pending real-time and offline verification products to measure ad quality while identifying creative fraud and malvertising across devices and channels.

Connecting malvertising, ad fraud & cyber crime

Cybercriminals leverage digital ad tech to make money in multiple ways. The simplest and most commonly known way is by using bots (software) to generate enormous quantities of fake traffic and ad impressions. There are many other forms of "cheating" that help criminals inflate the number of fake ad impressions they can exploit. For example, on websites they own and control, they can alter the code of the site to repeatedly load web pages, refresh ad slots every few seconds, stack dozens of ads on top of each other, or misrepresent low-cost display ad slots as high-CPM video ad slots, to name a few. They also use malicious code to trick fraud detection and viewability measurement technologies into labeling non-viewable and fraudulent impressions as "valid." Basically the bad guys are using tech to make "rotten apples" sellable as "fresh apples."

With this report, Confiant aims to shed light on a previously unquantified and unreported set of practices. The core of our data comes from having observed the criminals compromising real impressions from reputable sites and also attacking users directly, both for their own profit. The same attack methods are used, except that the criminals now take on the role of buyer and middleman. Instead of controlling the whole site, the criminals seek to control the ad slot to trigger forced redirect attacks to users, to stack misrepresented banner slots with fraudulently arbitraged video ads, or exploit the latest serious threat to the industry: cryptojacking. Increasingly, Confiant has observed the malicious actors layering multiple attacks into the same attack, with cryptojacking being the fallback if the forced redirect or the misrepresented IBV stuffing doesn’t deliver.

The industry has experienced a huge surge in malicious attacks since late 2017. The main reason behind this surge is quite simple: it is highly profitable. New attack methods have been developed, as well as new attack vectors, to capitalize on ever increasing attack surface presented by the open programmatic marketplace. Malvertising -- the portmanteau of malware and advertising -- is ubiquitous: every publisher, large or small, receiving programmatic ads is dealing with these issues.

In the study that follows, we show examples of middleman misbehavior -- e.g. misrepresenting in-banner video (IBV) -- and demand-side malfeasance -- e.g. allowing ad creatives with malicious code (malvertising) to slip by. We detected these forms of fraud with a the combination of our patent-pending real-time impression-level scanning technology and our proprietary offline scanning analysis, despite the fact that all the SSPs and DSPs studied had other fraud detection technologies already in place.

Expanding the ad fraud lexicon

Every ad fraud related discussion typically gets the air sucked away by our industry's hyper focus on bots and the buy side. Rarely is the fraud that victimizes the sell side discussed. To most, Ad fraud is about real ads served against fake sites or non human impressions. Unbeknownst to most advertisers, many criminals use the same methods to attack publishers, perpetrating fraud not as a fake seller, but rather by acting as a fake buyer. One example of this which we focus on in this report is when the criminals misrepresent their intentions to serve a banner ad into a banner slot and instead repurpose the inventory, unknowingly to the publisher, to sell a video ad instead.

Fraudulently Misrepresented In Banner Video (a.k.a. IBV) are a scourge on publishers both because of their network load & latency, including the fact that they are an outright fraud committed against publishers- and, we believe, often on the advertiser too. IBV are unlike pre-roll ads that load right before a video plays within a media player: IBV ads play outside of a player with no video content to follow. The favorite inventory of the fraudsters to misrepresent are smaller placements in the top banner or sidebar of a page, and the IBV is almost always set to autoplay. The heavy network load of these misrepresented video creatives is compounded by the standard mechanism by which IBV is usually served, which entails multiple concurrent VAST auctions running through a display ad slot. At best these misrepresented IBV slow down the site and are bandwidth hogging intrusions to the user. At their worst, they can severely impact the user experience by causing sites to stall or cause crowding in pages, making it difficult for users to consume their content.

To be clear, Confiant is not saying all IBV is fraudulent. IBV, when done properly with both the advertiser and publisher knowingly executing on a banner ad that runs featured video creative asset, can be part of valuable media buy strategy to scale up inventory. In these cases the video will often not play automatically, but will require a user interaction to start playing and the video content will be adapted to the smaller placement and shorter attention span. The Misrepresented IBV flagged by Confiant is not this. Misrepresented IBV auctions are attached to a hijacked banner creative and the winning video is often never even seen. Unscrupulous bad actors will often run this resource intensive process in a continuous loop in order to maximize their take.

Assessing the programmatic industry's ad quality

Quantifying the scale of the issues affecting publishers and their users has always been a challenge for our industry. Never before has any benchmark been done on how many actual creatives are bad, how many users are affected, nor where the issues come from. This report will be the first of many, with every new analysis delivering new insights upon which we, as an industry, can act collectively to win the arms race that is at the core of beating the criminals behind these attacks. Fixing the systemic quality issues that pervade the OpenRTB programmatic market will require the industry to unify on higher quality standards and technology - across the board. Though these issues are pervasive, Confiant can now report the first clear assessment of how many attacks actually occur every day.

For Q2 2018, our data showed the following of open programmatic impressions:

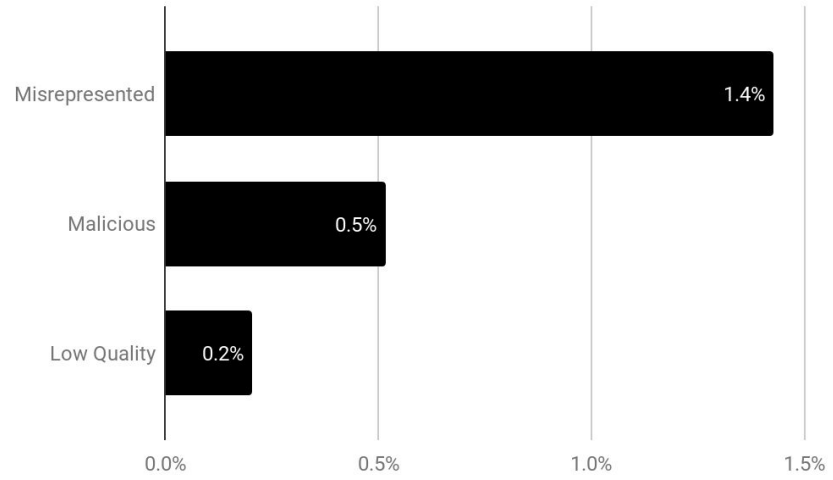

Chart 1 - % of Q2'18 programmatic impressions flagged by Confiant

"In Q2, we saw that over 0.5 percent, or 1 out of every 200, of programmatic impression were malicious. This was exceeded by the 1.4 percent of impressions that were misrepresented IBV, defrauding the publisher by serving unwanted in-banner video into display ad slots." added Mangin. "Our data shows the issue was pervasive across over a dozen ad exchanges, though with significant variances from the worst to best performer per tier likely reflecting the varied average CPM fill rates of the different exchanges."

To put these percentages in context, if the open marketplace is serving 1 trillion programmatic impressions per month, then 5 Billion of those are malicious, 15 billion are misrepresented IBV, and 2 billion are so low quality as to be blocked outright. Each of those impressions is a ruined user experience - that is 1 out of 50 sessions for unprotected publishers. **It is critical to note that that these numbers are additive to all the other ad fraud the industry deals with.**

Mapping the supply side spread of the infection

Beyond answering the question “how much?”, Confiant also analyzed its data seeking to answer the question “where from?”. Diving more deeply into data on 12 of the top SSPs/Exchanges, we observed a significant variance in their effectiveness at protecting their publishers from these issues. Grouping these results into two groups yielded the following:

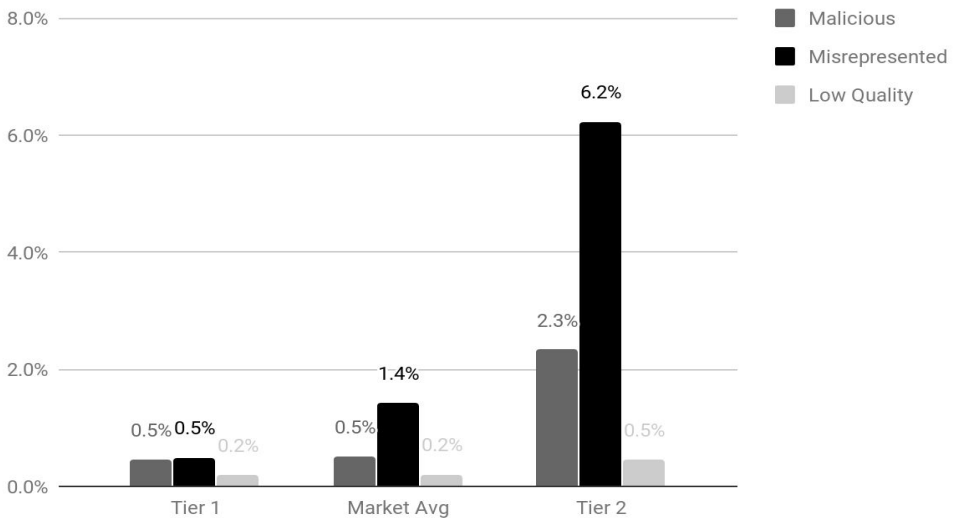

Chart 2 - Malvertising and misrepresented IBV rates of leading SSPs

Our data clearly shows two tiers of SSPs: Tier 1’s, of which 8 of the 12 were categorized as, show rates of malicious and low quality creatives that are equivalent to the market averages and almost 3x better for misrepresented IBV. Tier 2 SSPs (4 of the 12) show rates of malicious over 4x higher than the market and misrepresentation rates over 10x higher than Tier 1 players. The reality is that not all exchanges serve the same constituents, neither publisher-wise nor the same advertising buyers either. Even though they strongly outperformed Tier 2 players, the eight exchanges we classified as Tier 1 failed at being significantly better than the market average for both Malicious and Low Quality ads. Given the market averages encompass all the impressions verified by Confiant in Q2, including those from walled garden players like Facebook Audience Network (who constrain the use of third party javascript and therefore are safe from the bad actors attacks), this calls out the inadequacy of the industry’s current methods given no Tier 1 SSP came away looking clean.

All of the monitored ad exchanges heavily invest in traditional offline tag scanning technologies, and several are even Confiant clients, protected using Confiant’s 1st generation tag scanner. The data proves what many in the industry have known for a long time: the criminals currently have the upper hand on offline scanning based security solutions. Shifting the industry to a new paradigm is required, one which delivers real time verification to all levels of the industry. Confiant’s 2nd generation real time verification solution begins that process.

Where the buy side fits into this issue

Although publishers hold supply side platforms most responsible for the issues being discussed in this report, all players are intimately aware of the fact that the SSPs are not the actual source. It is the buy side and the DSPs who are the doorman to the criminals and bad actors. The pie charts below summarize the composition of issues detected by Confiant coming from the Top 5 DSPs as ranked by leading research firms. As many attacks can not be properly attributed to individual DSPs, due to varied methods used detect the malfeasance itself, we have summarized the composition of issue detected instead of their volumes (as we did for SSPs, for whom we have stronger indicators of identity). In no specific order, the top 5 DSPs' violations detected by Confiant had the following distributions:

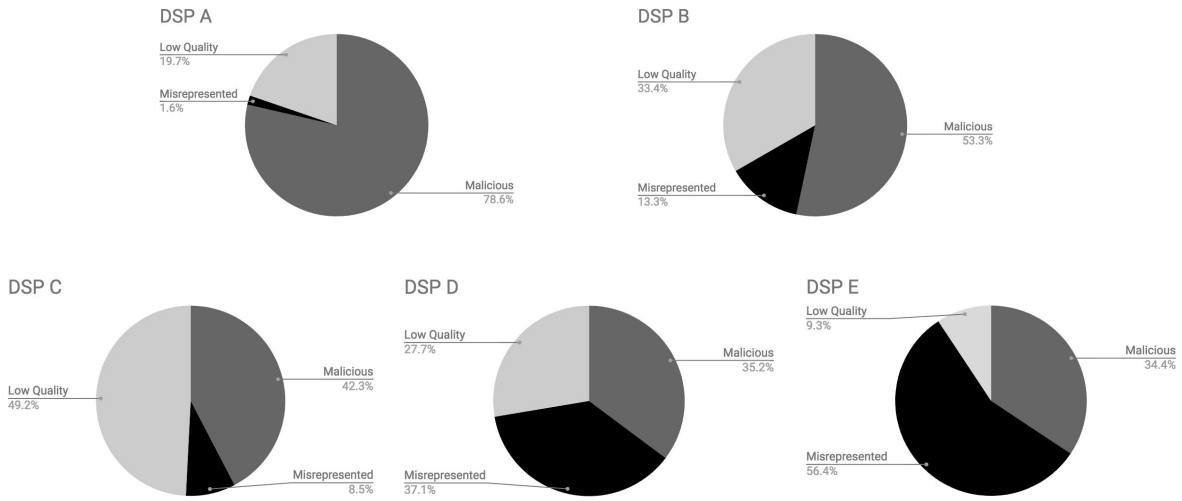

Chart 3 - Breakdown of flagged impressions for the Top 5 DSPs

As the data shows, there appears to be a negative correlation between a DSP serving a Malicious ad vs. a Misrepresented ad. DSP E had the lowest Malicious rate but the highest Misrepresentation rate, whereas DSP A served the highest Malicious rate and lowest Misrepresentation rate. This further illustrates the limits of the industry's current approach to these problems. Like the SSPs, all of these DSPs most definitely also scan their tags regularly. Malicious ads, fraudulently misrepresented ads, and low quality ads are often lumped into one large 'creative quality' issue, which DSPs erroneously believe can be solved just by scanning. Unfortunately, doing so obfuscates the real breadth of issues. One size does not fit all when it comes to these problems. Each issue is a distinct attack vector that requires its own set of calibrated tests to catch, and an uncompromisable data set to confirm that the bad actors have not evaded the tag scanner by detecting its non-human signature.

That it is so challenging even for Confiant to reliably identify the DSP enabling the bad actors is a further testament to the work that remains to be done in creating a properly transparent ad ecosystem. OpenRTB 3.0 makes serious strides in this direction, but will continue to have the inherent weakness of relying on self declared information.

Diving deeper into bad ads

Understanding the scale and vector of these bad ads as they wind through the ecosystem is a major step. Understanding the mechanisms and motivations of the criminals is just as, if not more, critical. The industry is replete with myths as to the criminals' methods, and Confiant has analyzed its data with a view to establishing facts and backing up (or disproving!) some of the more pervasive ones.

Industry Belief #1: Bad ad attacks get launched on Friday nights

A common perception is that weekends tend to be more active in terms of higher malvertising rates. While this could be true, the mass of data collected during this Q2 was not able to support that claim. Our data shows that bad impression volumes tend to spike and drop during different week days.

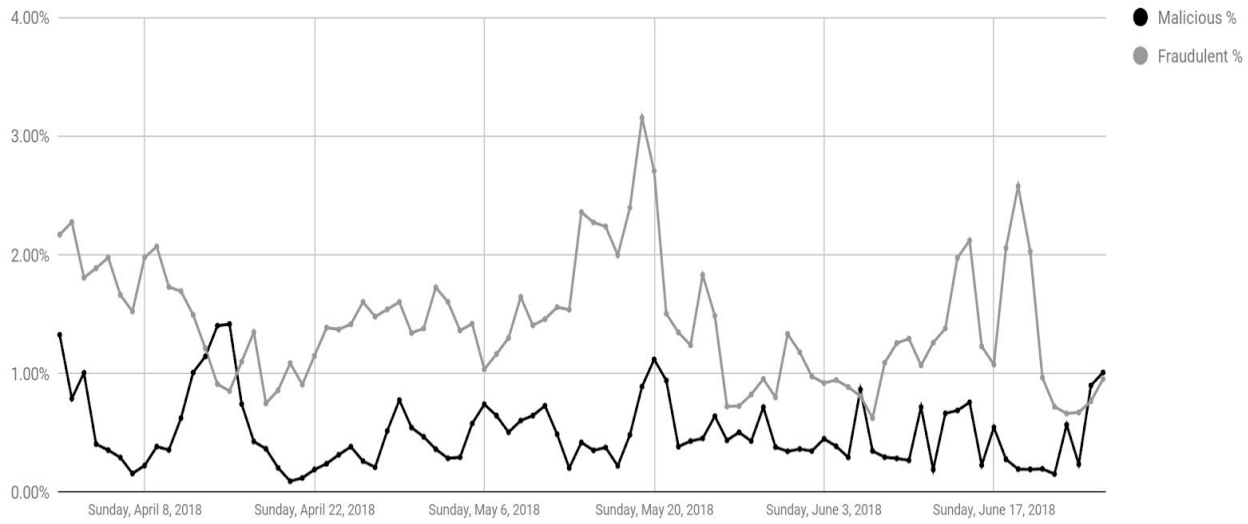

Chart 4 - Weekly trends of malicious and misrepresented IBV ads.

The key takeaway that our data shows is that it's not straightforward to predict when malvertisers are more likely to buy media. The biggest one day spike occurred from Sunday June 24th to Monday June 25th. If we rank by size of spike, it isn't until the fifth position on this list that we see a Friday to Saturday appear. The reality is that the data shows different sites and different users will attract attention at varied times. The social and technical engineering required by the bad actors to get access to the ecosystem are consistent and constant.

Industry Belief #2: Screening for infected landing pages is important

With 1 out of 200 programmatic impressions being malicious, Confiant dug deeper to break down what specific type of attack was being executed by the criminals. Our data shows, as expected, that the most common form of malvertising today are mobile redirects. Forced mobile redirect attacks attempt to solicit users to donate their personal data via fake gift card forms that engage the user with a non-existing reward. Every other attack method is a fraction in comparison, including desktop redirects (most typically to tech support scams), cryptojacking (hijacking the browser of unsuspecting users for the purposes of mining cryptocurrencies), ad stacking (showing hidden iframes with extra ads), cookie stuffing (forcing 100s of cookies onto the user) and other even more esoteric attacks. Infected landing pages came in second to last, comprising 3% of the total security issues flagged by Confiant, yet it is often held as just as important to screen for as to the other more pervasive issues.

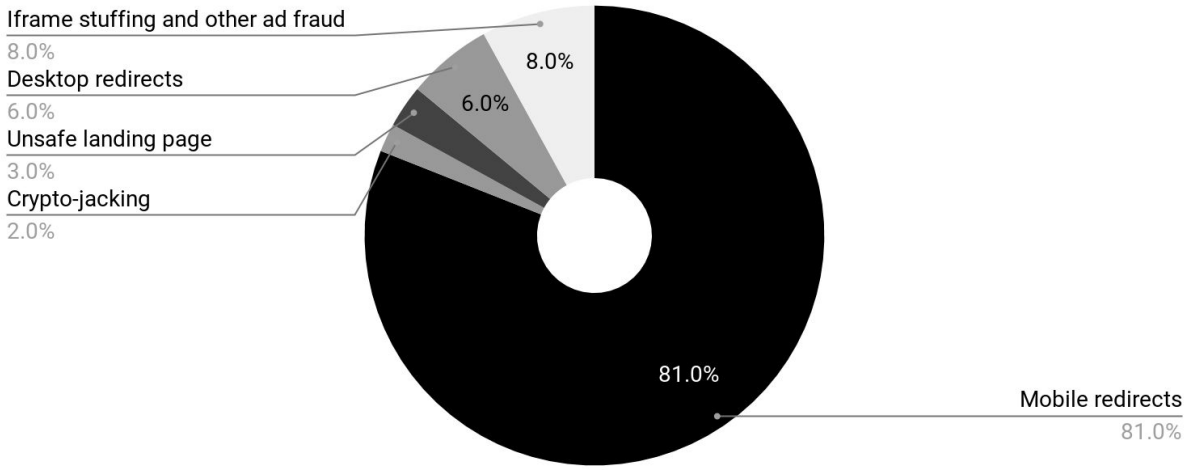

Chart 5 - breakdown of malicious ad types.

Browser exploit kits, which were the leading attack method for years, are now virtually extinct in malvertising attacks due to the overall improvement in browser security.

Industry Belief #3: In banner video (IBV) ads carry high CPMs

One would imagine that being the victim of misrepresented high-CPM video ads would at least yield higher than average returns for publishers. Not even. Our data shows that the bad acting intermediaries who deliver these fraudulently misrepresented IBV ads do not pass any of these lucrative margins down to the publishers. IBV fraudsters simply buy at the cheapest CPMs and pocket the arbitrage-driven margins, severely impacting the user experience in the process.

- Fraudulent Misrepresented IBV CPMs are 54% lower than Market Average.
- Malicious CPMs are 57% lower than the Market Average.

Industry Belief #4: High floors can protect a publisher

Divining the optimal floor price for programmatic impressions is not known as the favorite pastime of any ad operations team. Inherent to that question is the balance between blocking low CPM creatives along with the malicious creatives. Our data shows that malicious and fraudulent IBV impressions tend to trade in the price range that aligns with the distribution of the majority of total programmatic impressions, and this implies that the common practice of raising floors with the attempt to avoid these bad ads is not effective most of the time.

Average threshold values can be misleading though, since averages are sensitive to outliers and don't take into account the distribution. In this case, small subsets of impressions with extremely low/high CPM prices can throw the average off of where the majority of impression prices are. To correct for this, those outliers were discarded by Confiant for this analysis. This notion is better visualized in the chart below, showing CPM distribution across all traffic clustered into quartiles.

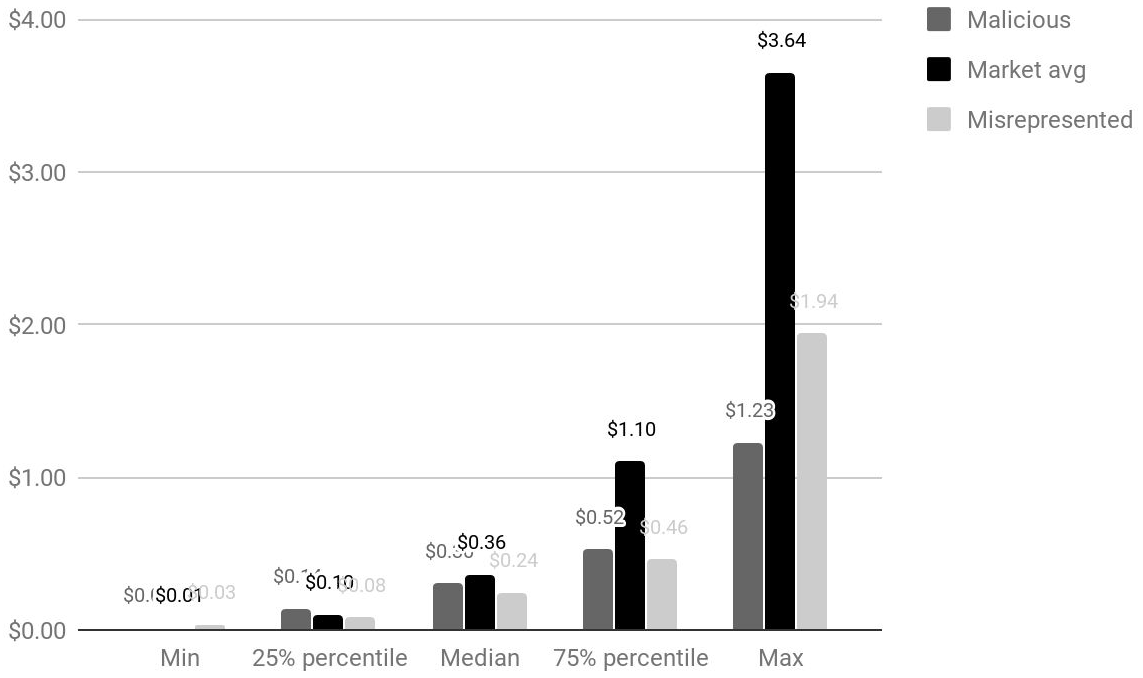

Chart 6 - Impressions distribution by CPM

Looking at the graph it's easy to see that "bad" (malicious and fraudulent) CPMs paid to the publishers are close to market average in the first and second quartiles. For the first quartile (the 25% percentile) Malicious ad impressions actually pay 36% higher CPMs. If a publisher were to want to set a high price floor, for instance $0.50 as an example, this would indeed block nearly 75% of the bad ads, but it will also block over 50% good ads, a revenue impact that can not be easily recovered.

Conclusion

Fraud in our industry is not a new topic of discussion. Clarity on the crimes of the bad actors pretending to be good advertisers is. Forced mobile redirects and fraudulently misrepresented IBV are just two of the latest attacks that the criminals utilize. These criminals are constantly adjusting their activities to align with whatever business model and strategy is most lucrative, and aren't afraid to branch out to hurt all parties. Trying to predict when and where bad ads are going to appear is not straightforward nor easy to rationalize - that's why a technically sophisticated approach, as Confiant has built, delivers granular control at the individual impression level post auction. Though the industry has started on the path to effective protection - Confiant today is protecting over 2,200 sites, our report clearly shows there are serious strides still to be made. Having revolutionized ad quality verification with our real time methods, Confiant has since been building back up from the micro (impression level) towards the macro (industry level), so as to deliver a holistic approach that delivers on our vision of maximum protection. Scaling up across more publishers and platforms, improving the detection methods to shorten the feedback loop, and many other upgrades are required before ad ops professionals will sleep soundly on Friday nights.

About Confiant

Confiant is a cyber security company that came out of a recognition that the world's most sophisticated advertisers aren't Verizon or P&G, but criminals using the industry for their own, selfish ends. These criminals are hijacking programmatic advertising and giving publishers a bad name.

Confiant protects the reputation, revenues and resources of publishers and platforms with always-on anti-malware software that verifies desktop, mobile, and video ads. Our sole focus is on helping advertising platforms and publishers rid the world of malware. This focus enables us to evolve quickly and meet our clients' needs for defeating the bad actors trying to undermine the industry.

We were the first to come to market with a technology that does not just detect the malicious activity, but actively blocks it. We believe in the intelligent application of technology to fight back and make digital media safe for everyone.