Marketing Fraud Benchmarking Report

By Adam Sell, HUMAN Director of Digital Marketing

RENEGADE

HUMAN

2021 Marketing Fraud Benchmarking Survey and Report

The greatest advantage that a marketing team can give to their organization

is to fully understand the buyer's journey. Buyers aren't responding to lists of features, data sheets, and as much as marketers wish they were, they're not making buying decisions based on individual pieces of top-of-the-funnel content. The path from awareness to consideration to decision can be a meandering one, with numerous entry points and channels along the way. Every buyer is different.

To that end, new marketing ideas have to demonstrate quickly that they're having an impact. They need to demonstrate a greater awareness of the brand, capture names for nurturing, or increase leads' velocity through the funnel. If these campaigns aren't doing one of those things, they go away. It's the "fail fast" principle, applied to marketing's efforts to find and retain new customers.

It's increasingly less and less about how a marketing department spends money and more and more about how they test tactics and messages to find what works best, and then to continue to optimize both for the audiences they're trying to reach.

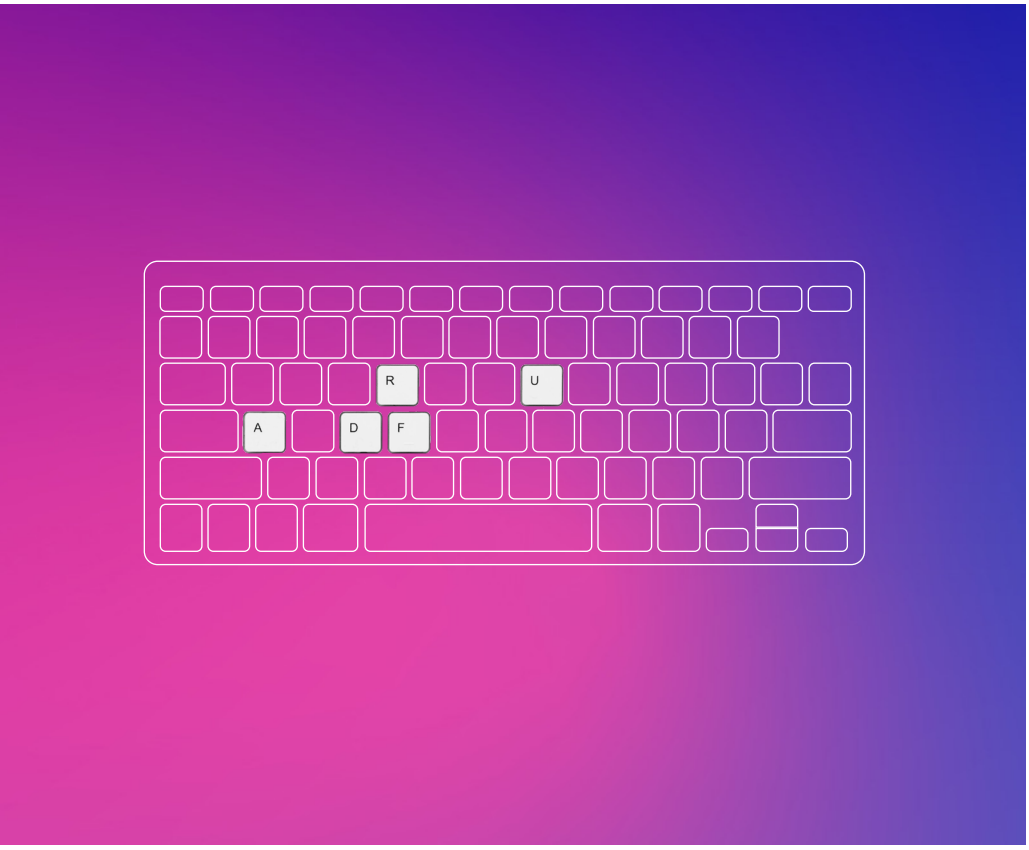

The digital marketing industry spent a third of a trillion dollars in 2019[1], and the ongoing COVID-19 crisis will likely push that number ever higher as consumers remain in the safety of their work-from-home circumstances. The unemployment spike that followed the onset of the pandemic didn't slow growth in holiday shopping in 2020[2], which underlines the role that digital marketing plays in consumers' daily lives, even as circumstances shift.

New marketplaces, testing new tactics, and commensurate increases in spending are frequently accompanied by a rise in fraud: there's simply too much money changing hands across the ecosystem for cybercriminals to ignore it as they always follow the money. The pie continues to grow and the criminals eye it with a thought about how big of a slice they can take while remaining unnoticed.

That phenomenon is marketing fraud. It's fraudsters messing with the tools of modern digital marketing, creating significant downstream effects. Not only can marketing fraud siphon off a significant portion of performance marketing budgets, but it can spoil the data that marketers use to make further decisions. It's the worst type of double whammy: first the marketing team wastes good money after bad because fraudsters have infiltrated the martech stack, and then later decisions are based on data that's no longer indicative of actual human interactions with the campaign.

Marketing fraud comes in numerous flavors, some of which are eminently relatable to any internet user in the 21st century. But some fraud models aren't as easily understood, and continue to fly under the collective radar of even the most senior performance marketing leaders.

HUMAN and Renegade recently conducted a survey of 129 digital marketing leaders with performance marketing responsibilities to better understand how familiar marketers are with marketing fraud models, the scale of marketing fraud threats, and their response to the challenge.

[1] Global Digital Ad Spending 2019, eMarketer, March 28, 2019

[2] Reviewing 2020's Holiday Shopping Season, Adobe Analytics, January 12, 2021

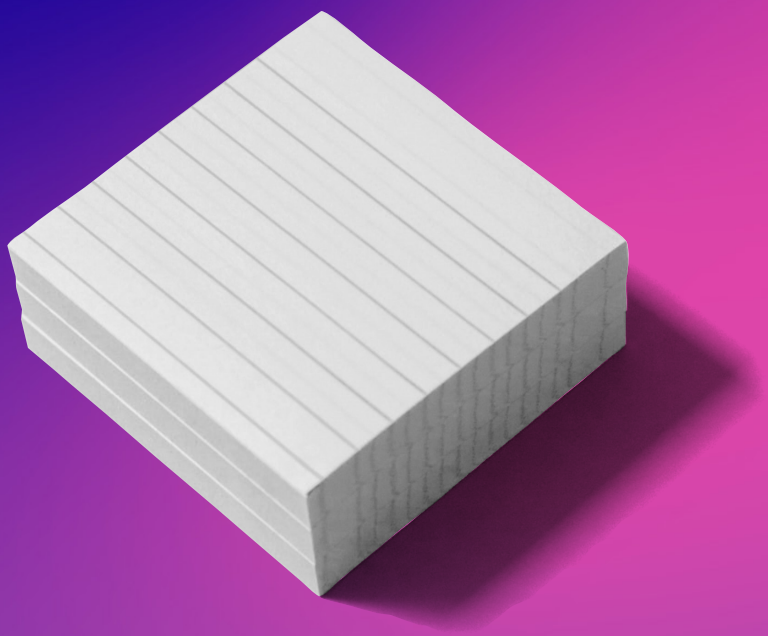

1

Executive Summary

Marketing fraud is a bigger problem than marketers realize, and it impacts companies of all sizes and maturities.

Today's marketers are asked to be both creative and analytical in proportions never before required. There's an expectation that marketers will produce brilliant, off-the-wall campaigns that outperform earlier work in every meaningful way.

That pressure creates a difficult, razor's edge situation for marketers. To meet ever-escalating demands, they work with partners to drive traffic, lead-generation targets, and other key performance indicators. But new tactics and a dramatic rise in spending are frequently accompanied by a rise in fraud. Marketing fraud is the phenomenon of bad actors (sometimes including corporate competitors) messing with the tools in the modern martech stack, wasting downstream spend and the data on which further business decisions are made.

Marketing fraud is a bigger problem than marketers realize, and it impacts companies of all sizes and maturities to the tune of millions of dollars a year for many organizations. A recent survey of 129 performance marketing leaders—sponsored by HUMAN and conducted by Renegade—uncovered what marketers are seeing and doing today, as well as where the gaps in protection lie. The positive is that the challenge is surmountable, promising significantly higher conversion rates and better business results if the right level of protections are put in place.

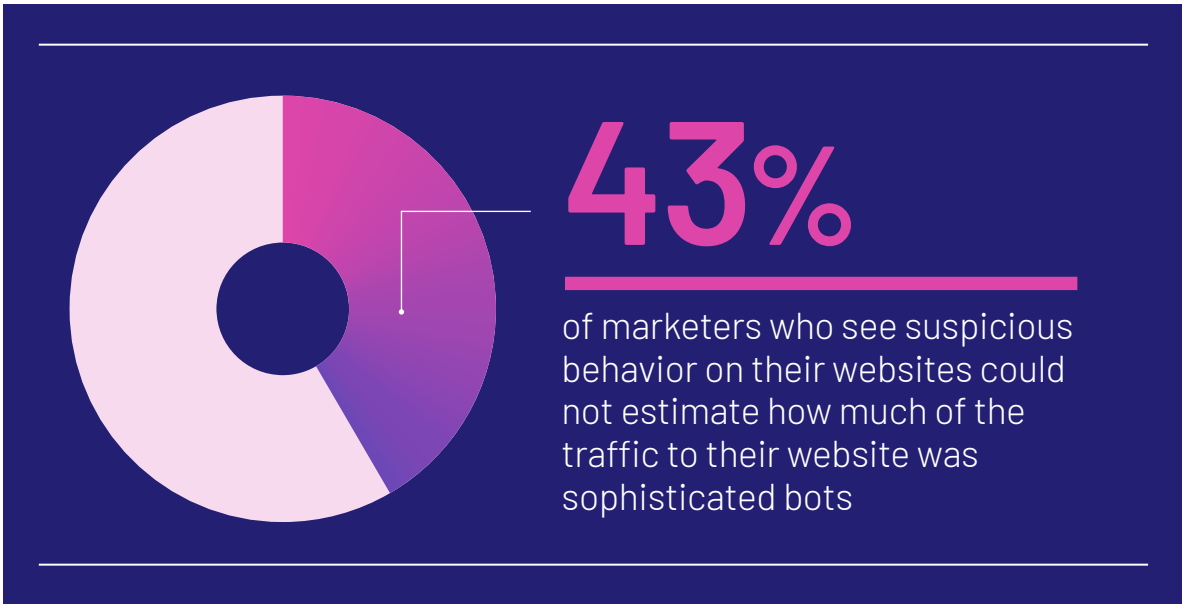

Two-thirds of marketers experienced some kind of marketing fraud in the last year.

The Marketer's Challenge

Digital marketers work tirelessly to make the best possible decisions using the data they have in hand based on the decisions they made before. It's a feedback loop that relies on accurate data and accurate insights into the customer's journey to inform future spending decisions.

Dirty data, therefore, has a pretty dramatic domino effect. If the data on which decisions are made is inaccurate, the decisions themselves may be faulty. And the performance of those campaigns impacts future campaigns, and so on. Marketers love to test new partners and tactics, and without the confidence that the KPIs that result are true, it's impossible to optimize for future campaigns.

What's more, it's generally up to small teams to manage those processes. There are only so many hours in the day to create these campaigns and measure their effectiveness before the next cycle begins; marketers don't have the luxury of doubting their data before making the next decision.

Where Does Fraud Come Into It?

All marketers are currently experiencing the impacts of marketing fraud on their websites: dramatic traffic spikes unconnected to new content or events, steep increases in traffic associated with marketing campaigns, time-on-site metrics that vary widely based on traffic source, lower than expected conversion rates, and complaints from the sales team about inbound or web-based leads. Based on the survey results published in this report, as many as **40% of marketers are seeing suspicious behavior on their websites and campaigns**. A solid 43% could not estimate how much of the traffic to their website was sophisticated bots, as opposed to the humans they intended to reach.

More than a fifth of marketers surveyed—22%—believe their first-party database is 25% (or more) populated by fake and fraudulent contacts. All of those fake or fraudulent contacts living in the database can have significant downstream impacts. Money is wasted trying to retarget and convert contacts that either don't actually exist or never truly engaged with the brand in the first place.

And what's more, there's a regulatory compliance element that comes into play: the easiest source of "working" email addresses for a fraudster is a data breach. This may result in a slew of real, working email addresses added to the first-party database, but without having opted into the subsequent nurture campaigns that attempted to push these contacts down the funnel.

A third of marketers surveyed experienced media buy fraud in the past 12 months, and the percentage grows even higher when the marketing budget exceeds $5 million a year. But a full two-thirds of marketers experienced some kind of marketing fraud in the last year. Media buy fraud, lead-gen fraud, retargeting fraud, and incentive program abuse all ranked among the top challenges for today's performance marketing leaders.

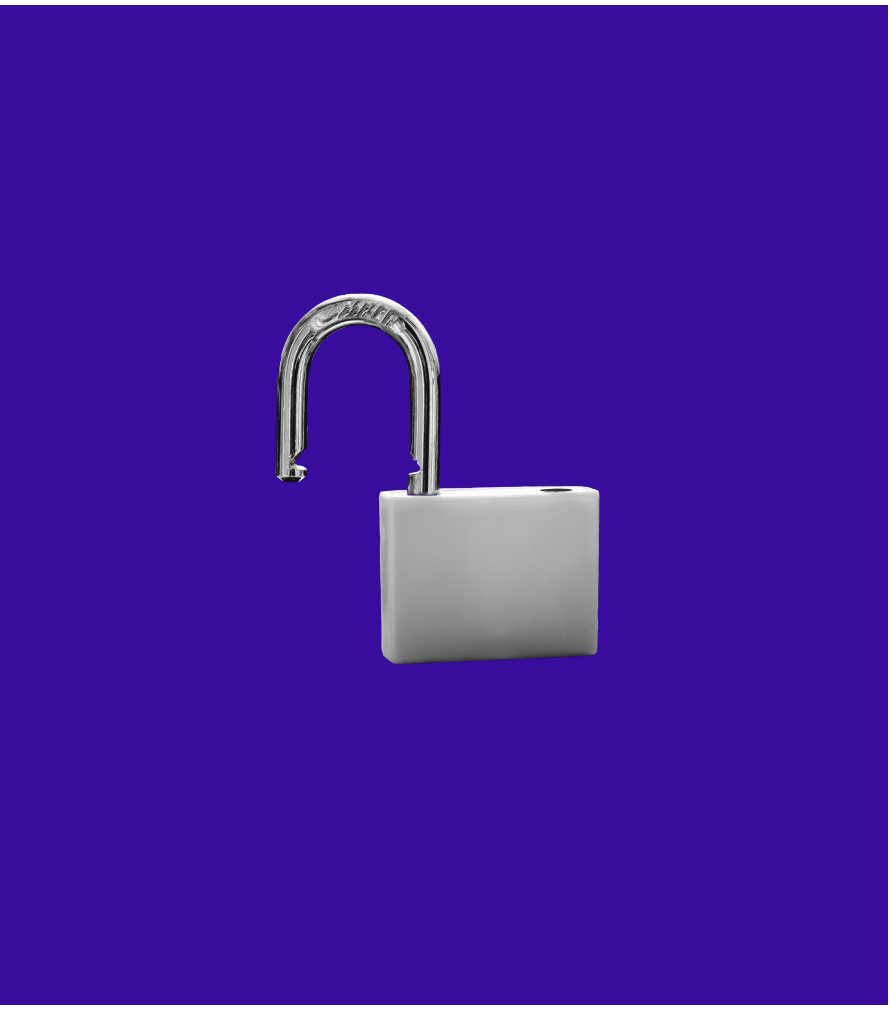

Additionally, the problem compounds over time if left unchecked. Less than half of respondents regularly scrub their databases for junk or fraudulent contacts, all the while adding more and more names to their lists through lookalike audiences, retargeting, and other tactics.

Perhaps the greatest cause for concern, though, is the lack of ownership of the problem: less than half of marketers talk with their cybersecurity teams when campaigns are producing unexpected or problematic results, and there was no consensus over which department should own the response. With no consistent ownership of how to ameliorate marketing fraud, fraudsters' efforts (which are growing ever more sophisticated and difficult to identify by the day) will only get more successful.

Solving the Challenge: Drive Significantly Higher Conversion Rates

→ It's not a doom-and-gloom situation, though. That lack of consensus presents an opportunity for marketing leaders: When you take control and address the problem head-on, you'll have a dramatic impact on all of your marketing campaigns and metrics. Dirty data results in poorly-informed decisions, so taking concrete, preventative steps to ensure the data reflected in your performance metrics is clean will enhance conversion rates, prevent data privacy compliance issues, result in better-informed campaign/budget allocations, and perhaps most notably, drive better business results.

2. Research Methodology

HUMAN and Renegade teamed up to survey 129 digital marketing leaders in the fall of 2020, each of whom expressed a responsibility for performance marketing within their organization. The goal of the survey was to benchmark the understanding of marketing fraud among digital marketers and to assess what tactics marketers were using to combat fraud models.

The survey was conducted via SurveyMonkey, and questions were reviewed by SurveyMonkey's research team before the survey was fielded to ensure they complied with the latest research protocols. Invitations were sent to a highly-qualified list of marketing leaders, and only those who expressed specific responsibility for performance marketing initiatives were permitted to complete the full survey.

While HUMAN and Renegade supplied lists of names to be invited for participation, the survey results were anonymous.

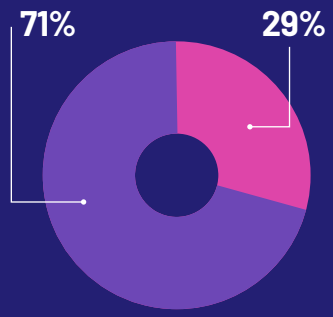

- Self-identified as director or above
- Self-identified as managers

More than half of respondents spend more than

$5Million

on digital advertising. More than a quarter spend more than $20 million.

129 survey participants

60% work at organizations with more than **1,000** employees

7% work at companies with fewer than 100 employees

More than half of respondents' websites see a monthly average of more than a

quarter-million visitors

Numerous industries were represented among survey respondents, with Retail/CPG composing the largest single percentage.

- Banking & Financial Services
- Retail & Consumer Products
- Software, Information and Communication Technology
- Consulting & Strategy
- Insurance
- Advertising, Arts & Media
- Engineering
- Administration & Office Support
- Trades & Services
- Education & Training
- Healthcare & Medical
- Manufacturing, Transport & Logistics
- Marketing & Communications
- Hospitality & Tourism
- Real Estate & Property
- Sales
- Science & Technology
- Accounting
- Sport & Recreation

3

What is Marketing Fraud?

Before we can talk in detail about marketers' awareness and perceptions of marketing fraud, we need to define the phenomenon and explore a number of its various fraud models to fully understand the challenge.

At its heart, marketing fraud is the misuse of any of a marketing professional's tools with the aim of diverting money or goods from the marketing technology ecosystem into the hands of a non-participant (fraudster) in that ecosystem.

That misuse is most efficiently carried out by sophisticated bots, which fraudsters deploy to scale up fraudulent activities. These bots often live on consumer devices, sharing web browsing histories, purchase histories, and other characteristics that make them difficult to identify amid the noise of general web traffic. They're often carrying out fraud without the device's owner's knowledge and with few if any signals that would clue a device owner into the scheme.

Malware is the most common vehicle for sophisticated bots to arrive on consumer devices, making it ever more crucial that users pay close attention to where their apps come from. More frustratingly, however, these bots often have complex persistence mechanisms, making them difficult to remove. After all, if you were a fraudster, you'd want to keep your tools out in the field as long as you could.

These bots can conduct a broad variety of fraud attacks, including many complex forms of advertising fraud and collection of sensitive user information. Their "work" in the spectrum of marketing fraud is of particular interest, and was the focus of the survey.

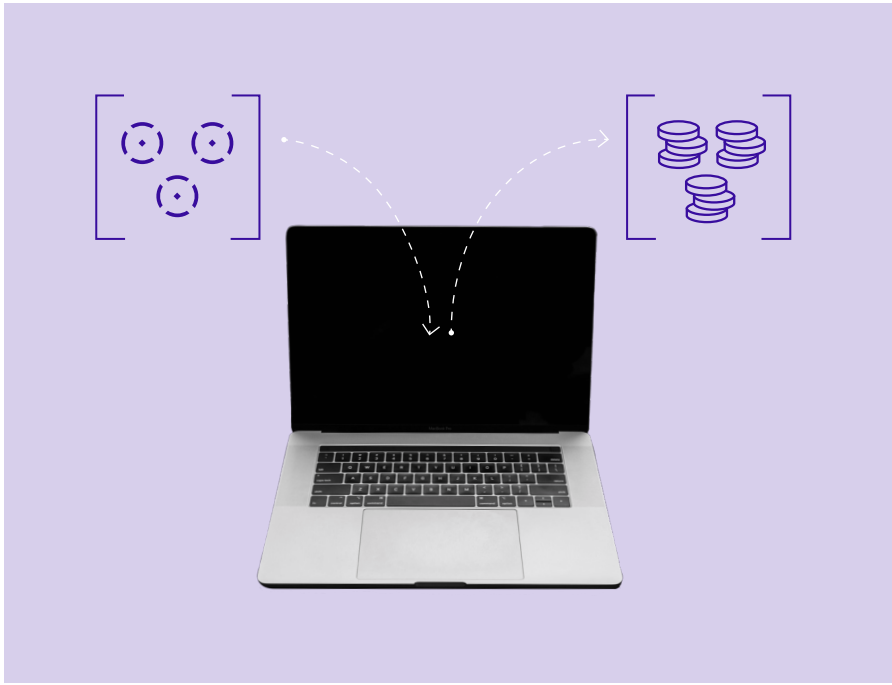

Marketing fraud is the misuse of any of a marketing professional's tools with the aim of diverting money or goods from the marketing technology ecosystem into the hands of a non-participant (fraudster) in that ecosystem.

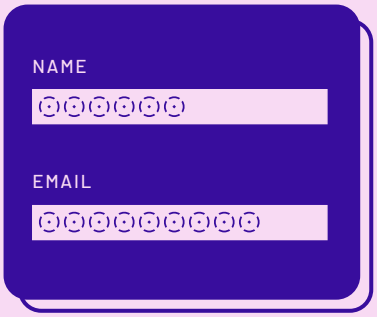

Common Marketing Fraud Models

Lead-Generation Fraud

Marketers at organizations of all sizes are familiar with the struggle: some of the names and email addresses in the first-party database are just flat-out garbage. Best-case scenario, it's just a case of someone wanting a piece of content but to avoid the subsequent email marketing outreach, and at a small scale, gibberish emails are easy enough to manage.

But worst-case, those fake emails are not only annoying, they're genuinely troublesome. And at scale, they can cause downstream effects that marketers want to avoid at all costs.

Consider: many organizations with sophisticated digital marketing programs employ retargeting software to ensure that their brand stays top-of-mind for potential customers. Those platforms rely on the cookies and first-party data that's captured on an initial visit to the organization's website, and "follow" the customer all over the web.

If those cookies or that captured first-party data isn't accurate—or even real—the retargeting efforts are going to fail. That's money spent on trying to reach and follow potential customers that may not even exist.

The challenge is compounded when you consider how those bots get the names they use to fill out those forms.

New data breaches are uncovered every week, with thousands of names, credentials, and PII exposed. There's a lucrative market for these stolen names and addresses, and marketing fraud is a big part of why.

What's a little scary for modern marketers is that those stolen names and addresses correspond to real people. These are people who don't necessarily know your brand or offerings at all, but who suddenly start seeing your logo in their inbox because a bot filled out a form.

With GDPR and CCPA concerns weighing heavily in the minds of many digital marketing leaders, this concept of stolen contact information finding its way into your first-party database is worth considering.

Retargeting is just one way that dirty first-party data in the DMP can cause problems, though. Many organizations are keen on audience extension opportunities - using a platform's captive audience to find and market to potential customers who share a lot of the characteristics of the folks in the database.

Fake or fraudulent contacts in the database upon which a lookalike audience is built can result in an organization targeting contacts who not only don't recognize the brand, but may not even work in the industry or role that the marketer is trying to target. It's another instance of dirty data throwing good money after bad.

Inventory Fraud

One of the most frustrating experiences, and one with which so many of us are familiar, is the disappointment of seeing a "sold out" error immediately after something goes on sale. Concert and movie tickets, video game and console releases, new designer shoes...all of these and seem to sell out within mere seconds of going on sale.

That's not an accident. That's bots.

Inventory fraud has a number of flavors, but the most relatable one is the use of bots to swoop in and buy up limited-release items before humans could possibly complete the process. And fraudsters don't often stop at just one bot to get one item: if a fraudster can do something once, they can do it several (hundred, thousand) times over.

It can be incredibly profitable for a fraudster, too - they effectively corner the market on limited-edition items and legally resell them on after-market websites. It's scalping for the modern era.

What makes it more frustrating than "traditional" scalping, however, is that consumers' frustration is, as often as not, directed at the brand rather than the fraudsters. In the 21st century, it's presumed—rightly or wrongly—that the brand is responsible for implementing sufficient measures to ensure humans get the items they search for digitally.

Fake/Automated Account Creation

"New customer discount".

Those words are like music to the ears of many fraudsters, especially when the site or brand that's offering the discount has high-value or limited-release items available for the sale.

In a similar fashion to how sophisticated bots can swoop in right at the moment of a new release, they can also create accounts in large batches to take advantage of special discounts available only as incentives for new customers. And in much the same way those products end up on third-party resellers at a steep markup, the discounts offered can give the fraudster a financial incentive to purchase anything they can with that markdown before pivoting to the third-party marketplace.

Brands can suffer, too, from the potential for sophisticated bots to carry out a negative review campaign en masse. A brand's rating on any major review site is intended to be a badge of accuracy and trust between the brand and its customers. But an attack of fake accounts, each submitting a negative review of a given item or of the brand itself, can quickly tank that trust.

At that stage, it becomes the marketer's role to work with the organization that manages and stages the reviews to root out which are fake and which are legitimate in order to restore accuracy to the system. That's time the marketer isn't spending on other initiatives.

Anything worth doing for a fraudster is worth doing many, many times over.

Marketing Fraud's Scale

One May Be Too Much

Marketing fraud is perhaps most accurately and reliably measured as a percentage of overall website traffic. The advertising industry coined the term invalid traffic (IVT) to refer to any traffic that's not driven by humans. IVT can include bot traffic that's a net benefit to most marketers, though: search engine web crawlers are bots by definition, and fall under that bucket of IVT.

1-40%

Average SIVT rates vary widely by industry and website, but HUMAN has observed a range of 1-40% fraudulent traffic on key campaigns.

So the advertising industry coined a separate term to refer specifically to the malicious bots themselves: sophisticated invalid traffic, or SIVT. The distinction is that bots that fall into the first bucket but not the second are very easy to spot, and for most marketers, easy to block if the situation warrants.

SIVT percentages can vary widely from one website to the next, depending in large part on the tactics that the marketing team uses to drive traffic.

HUMAN' internal research uncovered an SIVT rate of 37% for a specific campaign for one luxury automaker. The tactics the automaker used to push people back to their website were resulting in a staggering number of sophisticated bots filling out forms, clogging their databases and frustrating their sales teams.

Now, 37% may seem like an extreme, but even 1% SIVT has significant impacts, especially at the scale at which digital marketers work. That 1% has the downstream effects outlined above, including taking up space within a DMP, retargeting platform costs, brand reputation costs, potential compliance concerns, and wasted time and effort on the part of sales professionals working to follow up with fake or fraudulent information. That 1% is only a leading indicator for the impacts of marketing fraud throughout the go-to-market organization.

Marketing fraud doesn't have to be the "cost of doing business," it can be mitigated entirely.

The Receipts

Translating from SIVT percentages to actual dollars on the line can be a complex proposition, including taking an inventory of what marketing services are in use, exploring how an organization is performing on social media and search engines, and examining what campaigns are in progress and what audiences they're targeting.

On the whole, though, it can scale up quickly. Using a conservative SIVT estimate, HUMAN concluded that leading retail organizations lose as much as $7.3 million each year to sophisticated bot-based marketing fraud, and another $7-8 million on wastage within the marketing tech stack.

In short - the greater the toolkit, the greater the risk. The more tools a marketing department uses, and especially which depend on clean data within a DMP, the greater the potential losses as a result of marketing fraud.

4

What are Marketers Seeing Today?

4.

What are Marketers Seeing Today?

Perhaps the most fascinating thing that our research evealed was that marketers are, for the most part, aware of marketing fraud as a phenomenon, but believe the impact of it is merely noise. They've encountered or experienced one form of fraud or another, but they haven't necessarily rated marketing fraud as a challenge in dire need of a solution. But the actual scale of marketing fraud often catches marketing leaders by surprise.

We asked marketers about specific observations that correlate to instances of potential marketing fraud, and we found that marketers see fraud both inside and outside their organizations. There are numerous signals hiding in plain sight that can clue a savvy marketer into potential fraud. And, as referenced earlier, there are numerous forms of marketing fraud that the average consumer might encounter on any given day.

What's Going On With My Site?

Web traffic. It doesn't rank among the top metrics that digital marketers live and die by, but it does serve as a leading indicator for those critical KPIs.

The challenge is that web traffic isn't a "what you see is what you get" situation. There are a number of phenomena that can hide amid seemingly-typical web traffic metrics that are actually indicators of fraudulent activity.

For instance, dramatic spikes in web traffic that don't appear to correlate to a specific new piece of content, new event, or campaign may be indicative of fraudulent activity. As much as marketers (the author included) would love to believe that the adoring public might spontaneously discover the value of what an organization has to offer and begin arriving at the website en masse, it's not a realistic expectation absent a driving force.

Symptoms of marketing fraud

Dramatic traffic spikes unconnected to new content or events

•••

Steep increases in traffic connected to marketing campaigns

•••

Time on site metrics that are drastically different depending on traffic source

•••

Your Sales team is complaining more about the quality of website captured leads

•••

Lower than expected conversion rates

An influx of new traffic may be hiding the very fraudsters that marketers want to prevent.

Unexpected and unexplainable spikes in web traffic are a common characteristic of fraud: a wave of bots may arrive on the site as the result of a campaign partner's tactics, or advertising efforts may have become the unintended victim of a fraudster's campaign. One in five marketers surveyed noted that they'd experienced traffic spikes like these.

Fraud amid web traffic isn't limited to attacks that appear out of the clear blue sky, though - shiny new marketing campaigns, content, and product releases provide fantastic air cover for fraudsters.

In a vacuum, it seems like it would be a great moment: the new campaign surrounding a high-profile product release is performing great. The website is hitting metrics never before seen by the organization and the product is flying off the shelves. But as we learned earlier, inventory fraud is a model of marketing fraud that many people have experienced, and this is precisely how those bots fly under the radar.

It's important to remember that the people who run these attacks are smart. If they weren't, they'd get caught in a hurry. An influx of new traffic may be hiding the very fraudsters that marketers want to prevent.

How the web traffic arrived, too, can be a clue, as well as how long the sessions last. It's common for an organization's website to have a broad variety of traffic sources, and it's common for those sources to spend differing amounts of time exploring. After all, a visitor who arrives based on searching for what you offer may spend a great deal longer on your site than someone who clicked through from news coverage.

But there's a limit to that discrepancy. Time-on-site metrics can be a very obvious fraud signal: if all of the traffic from a particular referral site has a time-on-site average of zero, that's a clue that something isn't quite human about what's going on there.

In the b-to-b world especially, there's an old aphorism that the sales and marketing teams will

butt heads based on the volume and quality of the leads that are passing from one side of the go-to-market operation to the other. Sales may argue that the leads aren't good enough or that they don't fit the ideal customer profile, while marketing may argue that sales isn't doing enough with the leads they've been given.

That conversation can be a healthy one: marketing should always be iterating on what levers they pull in response to sales' feedback about lead quality and volume. And sales should be giving each lead a genuine opportunity to convert before chalking it up to a bad program from marketing.

When sales' feedback is, time and again, that the leads coming from web-based marketing initiatives aren't responding, though, that may be a sign of trouble. When many organizations have lead scoring algorithms in place to notify sales when a lead is theoretically warm enough for outreach, a fake lead slipping into the system can result in a fake MQL...which in turn leads to wasted time on the part of the sales team trying to follow up.

And finally, conversion rates. These are the numbers that campaign managers and inbound marketers obsess the most over, and they're a key leading indicator when things are going well. But numbers aren't always as they seem - fake leads that come in can impact those conversion rates in unpredictable ways.

Much like the case with unexpected and unexplainable traffic spikes, marketers (again, the author included) want to believe that they can be surprised by a gangbusters performance of a new piece of content. On the flip side, content that has a

40%

Of marketers could not estimate how much of the traffic to their website was fake or fraudulent.

much lower than expected conversion rate is also a potential sign for concern.

Two in five marketers surveyed reported experiencing lower than expected conversion rates on new content - while nobody can perfectly predict how content will perform, most marketers have a ballpark in mind. Missing the mark completely—in either direction—is a red flag.

There's a classic planning matrix about knowns and unknowns - you know what you know, you know what you don't know, and there are things that you don't know you don't know. Smart planning comes down to checking off the three boxes that include knowns, and accommodating as best you can for the unknown unknowns.

Pick the top 3 KPIs (key performance indicator) for your role

Revenue	66%
CPA (cost per acquisition)	54%
LTV (lifetime value of a customer)	42%
CPC (cost per click)	32%
MQL (marketing qualified leads)	25%
CPM (cost per thousand impressions)	24%
Pipeline	19%
SQL (sales qualified leads)	18%
Net new logos	6%

What is your annual spend on digital media related to performance marketing?

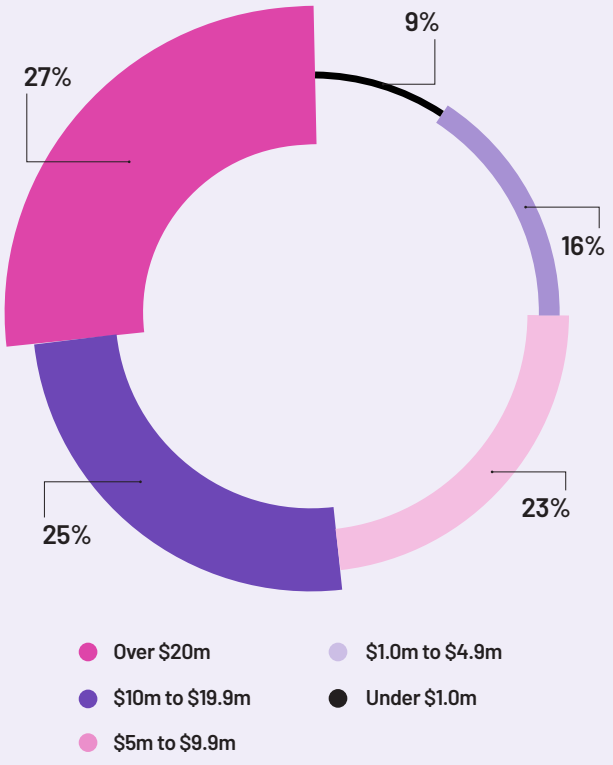

Marketing fraud is one of those unknown unknowns.

More than 40% of marketers surveyed couldn't estimate how much of the traffic to their website was fake, and an additional 17% of marketers ballparked the figure between 11% and 40% of traffic. And similarly, 37% of marketers don't have a sense of how legitimate the contacts in their database are, with an additional 22% expressing that they think their marketing database is 25% or more fake.

That's largely a challenge of tooling and prioritization. Marketers are asked to do so many things across such a broad variety of tactics that it's no surprise that there just aren't enough hours in the day to think about whether all the visitors to the site and all the names in the database are actually human. Marketers have to spend their precious hours setting up the next campaign, assessing new partners' capabilities, managing the tools in the toolkit, creating the next new piece of content...the list goes on.

Check off all types of bot fraud you're aware of (select all that apply)

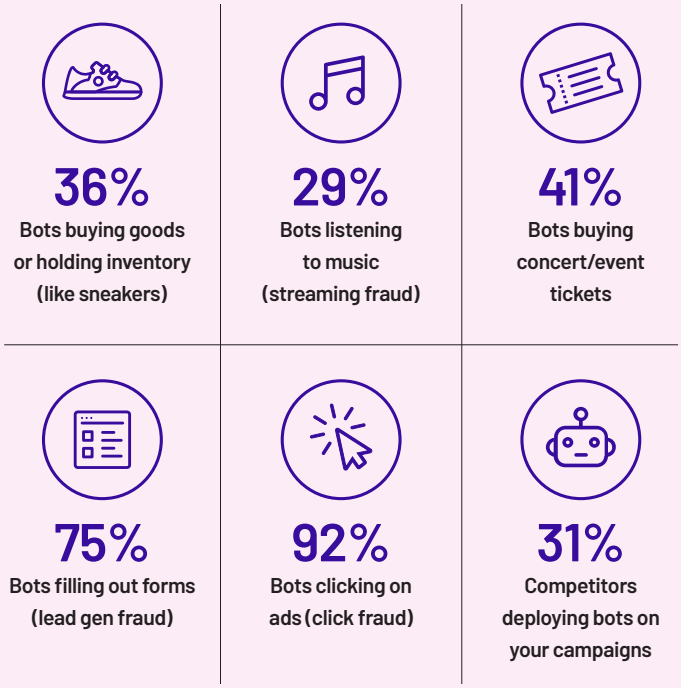

Not to mention, there are precious few tools available that can do the job of bot-or-not decision-making. Assessing whether or not a given lead is real or fake is the very definition of something modern marketers don't have time for.

What's frustrating is that despite that resource shortage, the impact of those fake contacts and fake web traffic could be pretty dramatic.

Performance marketing is one of the fastest-growing corners of the marketing ecosystem, and the respondents to the survey are spending significant amounts of money on this particular tactic. More than a quarter spend more than $20 million on performance marketing every year.

Think about where that money might go, though - above, we saw that 17% of marketers suspect that between 11% and 40% of web visitors aren't human. How much of that $20 million is being spent on driving non-human traffic? More than one in five marketers believed their database was 25% (or more) full of bots. How much of that $20 million is being spent generating leads that aren't real?

Cost-per-acquisition (CPA) is one of the top KPIs for marketing leaders. If as much as 25% of the database is gibberish or fake, how skewed are the CPA metrics for any given campaign?

Hey Joe, What Do You Know?

→ When it comes to specific awareness of various models of marketing fraud, marketers know what they've seen and they know they've experienced it. We asked marketers about their

Have you experienced any of these types of marketing fraud in the last 12 months? (select all that apply)

Media Buy Fraud *bots clicking on ads/paid search results*	36%
Lead Generation Fraud *bots completing form-fills with fake or stolen contact info*	23%
Retargeting Fraud *paid retargeting of bots instead of real humans*	20%
Inventory Fraud *inventory holding or inventory exhaustion for resale*	7%
Incentive/Referral/Loyalty Program Abouse *bots gaming incentive programs*	12%
Competitive Assaults *marketing fraud including any of the above organized by competitors*	5%
None of the above	37%

awareness of a variety of fraud models that are perpetrated by sophisticated bots, and while marketers are familiar with some models, others are a little less recognizable.

The vast majority of marketers surveyed knew about click fraud as a sophisticated bot-based tactic, and three-quarters of marketers know about lead-generation fraud.

The flip side? No more than two-in-five marketers were familiar with any of the other types of marketing fraud that we asked about. That included sneaker fraud, streaming fraud, inventory fraud, or targeted attacks by competitors.

When it comes to what types of marketing fraud performance marketers have experienced, the numbers suggest that most survey respondents know they're being hit.

Nearly two out of every three respondents said they were the victims of some kind of marketing fraud in the last twelve months, with media buy fraud, lead-generation fraud, retargeting fraud, and incentive program abuse topping the list.

The bigger spenders among the survey respondents, perhaps obviously, were significantly more likely to experience media buy fraud and incentive program abuse.

One respondent mentioned that they had been the victim of a type of carding attack. Carding is a form of fraud in which a fraudster steals credit card numbers and uses them to buy something to resell later for cash. This particular respondent was a victim in the middle step of that process: this fraudster was using their site to test those stolen cards and ensure they worked before carrying out the remainder of their scheme.

5

What are Marketers Doing Today?

5.

What are Marketers Doing Today?

To date, there have been a limited set of tools designed specifically to tackle the challenge of marketing fraud. Marketers have a general awareness of fraud, though their responses to the situation are piecemeal, scattered across a broad variety of point solutions aimed at particular symptoms of fraud.

The Toolkit

Many respondents to the survey identified ad verification partners as a part of their marketing fraud toolkit. And while the services these partners offer can go a long way toward preventing the siphoning of marketing budgets (particularly those dedicated to display, native, and in-app advertising), their capabilities are limited at preventing fraud that's unassociated with advertising and appears to arrive on a marketer's website organically.

Other respondents identified CAPTCHA-style tools and hidden honeypot fields as their primary mechanisms for spotting marketing fraud. But CAPTCHA/cognitive challenges are defeatable (either by sophisticated platforms or outsourced to human puzzle-solving farms), and honeypot fields only work if fraudsters happen to have their bots configured with certain characteristics (like JavaScript, for example).

Regardless of the tactics and tools identified, however, the bulk of marketers who participated in the survey were merely responding to fraud, not preventing it outright. Remediation is important, of course, but response time, resources, and toolkits limit how much marketing fraud-based damage can be undone.

That's not to say that marketers should stop working to make themselves whole following an advertising campaign that's fraught with bots. But when the downstream impacts of marketing fraud are managed manually and platform-by-platform, the fraudsters will "win".

Winning, to a fraudster, isn't about outright conquest. At its core, fraud is an economic problem: it's easier and more lucrative for some people to

Changing the economics of cybercrime:

The way to stop fraud is not through simply throwing more tools at the problem, it's through changing the equation in the first place. Consider: when a fraudster's tactic is discovered and blocked, they have to start over again, either with a new victim or with a new tactic. Every fraud scheme is, therefore, a race against time. How long can they operate before they get shut down? And how much money can they make in that time?

Reduce those last figures—the time to operate and the amount of money available—and fraud becomes a less appealing option. There will simply be less reason to bother with the effort of finding new mechanisms and approaches.

The way that we make that shift is by shortening (or even eliminating) the amount of time in which a fraudster's operation can succeed. It's in developing a technological solution to the economic problem of fraud.

develop and deploy the tools to steal money from complex systems than it is to work within the legal bounds of those systems. And the likelihood of—and potential penalties for—getting caught are slim enough not to serve as effective deterrents.

One of the most worrying statistics from the survey centered on how marketers managed the fake data that had made it into their first-party database. Fake contacts sticking around within first-party databases are the cause of the bulk of the downstream impacts of marketing fraud, so it's incumbent on—and imperative that—marketers regularly scrub their first-party database to ensure that these contacts are removed.

But despite the downstream threats that marketing fraud poses, fewer than half of marketers surveyed scrubbed their database with any sort of regular cadence.

These fake names are piling up, risking compliance issues, siphoning marketing dollars away, and crushing the metrics on which future campaigns are built. And what's worse, it cost money to get those names in the first place. If clawbacks are the only mechanism a marketer has for managing fake contacts, it will be an even less effective tool unless marketers dig into their databases regularly.

How Contacts Are Used and Found

→ That first-party database, the marketer's crown jewels, has numerous applications for the modern, multifaceted marketer. Virtually every tactic a marketer wants to run relies on the development, maintenance, and hygiene of this crucial database.

For example, a full ninety percent of marketers surveyed indicated that they ran email marketing campaigns using this first-party data. And that's a logical use case for that information. But as alluded to above, if that database isn't kept clean of junk and

fake emails, there may be notable problems later. (Read on, MacDuff - our next chapter will address some of those problems.)

Marketers will recall the old "rule of seven": prospects need to encounter a brand at least seven times before the messaging will stick and a true purchasing opportunity arises. And it's perhaps with that classic axiom in mind that so many modern marketers take advantage of retargeting/remarketing tools. Of those surveyed, more than 80% said that retargeting was a part of their toolkit. For fake contacts, though, it may be a case of throwing good money after bad.

What do you do with first-party data?

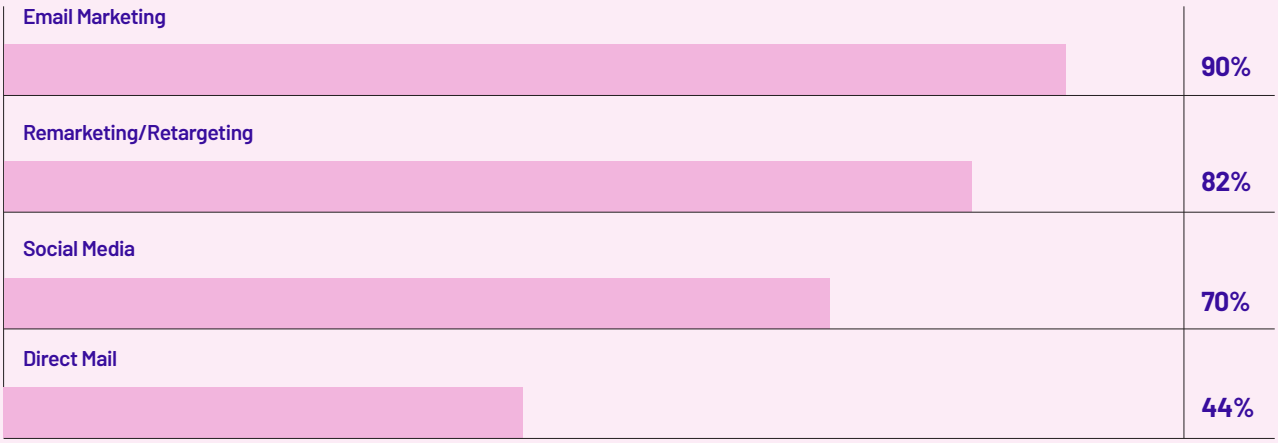

And in order to create and grow the database on which these campaigns are built, most marketers turn to partners who have first-party databases of their own. Three out of every four respondents to the survey took advantage of similar audience or lookalike campaigns to help grow their own contact lists.

That particular tactic, in light of what we know about marketing fraud, shows a fantastic faith, both in the partner's database and in their own. If the list of contacts on which the lookalike campaign is built includes fakes, how effective will the lookalike campaign be at hitting the targets? And if the partner doesn't have robust marketing fraud protections built into their own system, who's to say that the contacts being provided aren't themselves fake?

Retargeting/remarketing tools are invaluable to marketers, helping to ensure brand awareness and affinity, and encouraging return visits and conversions.

But sophisticated bots can result in wasted spend on these key tools, as they pursue contacts all over the internet. The cookies that power these tools are only as good as they are accurate.

Marketers need to be sure it's human activity—not bot activity—that planted the cookie.

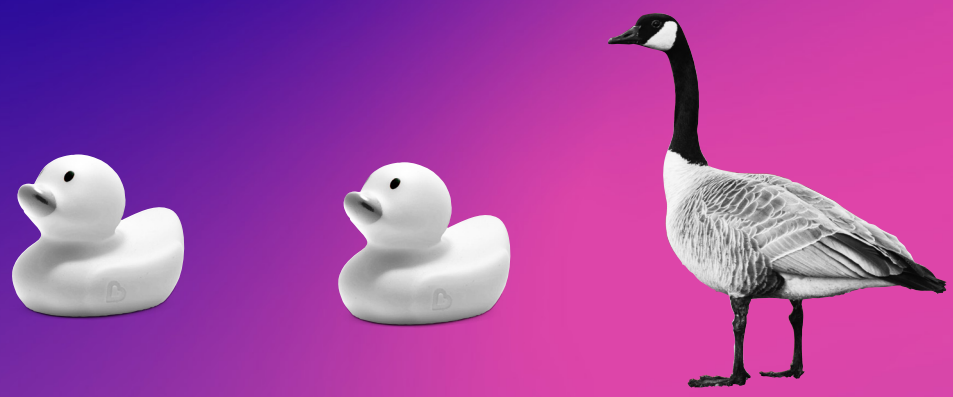

6

Causes for Concern

6.

Causes for Concern

→ The survey revealed that there are both causes for concern and causes for hope. While it's great that there's a general awareness of marketing fraud and about several specific models of fraud, some marketers have plenty of opportunity to be doing more to prevent fraud's impacts on their organizations.

Let's start with some of the causes for concern.

Possession is Nine-Tenths of the Problem

→ The short version: there is no consensus among marketers as to who owns the management, prevention, and eradication of marketing fraud within their organizations.

The longer version: respondents were fairly evenly split among three different departments as being the final decision-makers. A little more than a quarter of respondents identified the Chief Digital Officer (or the Digital department as a whole) as holding the ultimate responsibility for marketing fraud. One in three respondents suggested that it was the CISO (or the Security/Information Security department as a whole), and two in five suggested it was the CMO (or the Marketing department as a whole).

Perhaps more concerning, though, is that 12% of respondents did not know whom in their organization was responsible for marketing fraud management.

(As you may have inferred, respondents could choose more than one department for this question.)

Marketing fraud is, admittedly, not a very sexy thing to manage. Data hygiene and campaign optimization can feel more like table stakes than like major game-changers.

But getting marketing fraud under control can be the "boring" thing that facilitates a slew of "fun" things for a marketer. Ensuring that the campaigns are reaching and converting humans saves money,

Respondents identified the following departments as having authority to manage marketing fraud:

40%
Marketing Department

33%
Security Department

27%
Digital Department

12%

did not know who held final responsibility to manage marketing fraud.

freeing up funds for longer/broader campaigns, or for new, creative, eye-grabbing ideas.

The corollary to the above question about who is responsible for marketing fraud is to ask whether those departments are at least speaking with each other when metrics look awry. Less than half of respondents, though, consult with security teams about unexpected or unpredictable campaign results.

It's an *opportunity*, not a threat. Marketers are familiar with security teams' insistence on being involved and up to speed on any activities that might prove to be an entry point into the complex ecosystem they protect. Give security teams the opportunity to weigh in on why a marketing campaign is performing unpredictably.

There's also a harm avoidance instinct that should come into play with respect to marketing fraud. Yes, it protects budgets. Yes, it reinforces compliance initiatives. But it's also a brand safety measure. Consider the potential impacts of an influencer becoming the unwitting accomplice to a scheme that targets your organization: you may find yourself with a different type of mess to clean up.

Nearly half of the respondents to the survey weren't sure if marketing fraud had caused any impact to their customers' experiences or to their own reputation. It's an understandable gap in knowledge – if marketing fraud hasn't been a top priority for marketing leadership, less quantifiable impacts of marketing fraud will be difficult, maybe even impossible to measure. Those who did suspect they knew the answer to the question, though, overwhelmingly believed marketing fraud had not caused an impact to CX or their brand reputation.

Marketing fraud is a ticking time bomb for brand safety. When it goes off, the fallout will be significant.

According to PwC, ***44% of CEOs rank data privacy*** *among their* ***top three policies*** *most impactful to their businesses [1].*

[1] Top Policy Trends 2020: Data privacy, PwC, November 2019

60%

of marketers believe themselves to be average or worse at preventing marketing fraud.

What We Don't Know Can Hurt Us

We've written throughout about the downstream impacts of marketing fraud and how damaging they can be to marketing budgets, revenue operations, and brand sentiment. But how well do marketers believe they're doing at managing the problem and those downstream impacts?

When asked about how concerned they were about compliance with new privacy regulations like GDPR, CCPA, and others, the vast majority of survey respondents (75%) expressed concern about compliance. And they're right to be wary: the regulations carry dramatic penalties and public relations nightmares.

But as we saw earlier, too few respondents regularly scrub their databases of fake/fraudulent contacts. It's an obvious cause for worry when you consider how those contacts may have reached the database in the first place. Only half of those surveyed knew about the potential compliance implications of fake contacts.

And it's the practitioners who are most worried about compliance issues. The folks who live and breathe with that first-party database, the folks who are carrying out the email marketing and retargeting and social media lead-generation programs: they're far more worried about compliance than senior marketing leadership.

Generals: please listen to your troops on this one. Compliance isn't a checkmark, it's an ongoing conversation and an ongoing effort.

When it comes to the self-assessment, marketers aren't too confident of their fraud avoidance abilities. A solid 60% described themselves as average or worse at preventing marketing fraud, and as we've seen, even those numbers may reflect a limited awareness of the breadth of marketing fraud tactics. Nearly one in ten respondents couldn't rate themselves at all in marketing fraud prevention.

7

Causes for Hope

7.

Causes for Hope

While there are, yes, several reasons that marketing leaders should be concerned about the results of this survey, there are also several reasons to be optimistic about the future. There are a few things that digital marketers are doing today, in the course of their normal everyday work, that make marketing fraud easier to identify and trace, which in turns makes it easier to cut off the sources of that fraud.

Similarly, as explored earlier, there are a number of remediating actions that marketers are actively pursuing today, each of which can contribute to the overall elimination and prevention of marketing fraud.

Marketers should read this report as an opportunity rather than as a litany of woe. Implementing ***new anti-fraud measures can have enormous impacts*** *on the marketing technology stack and on the bottom line.*

Championship-winning teams in any sport have to excel at both offense and defense in order to keep their opponents at bay. Many of the remediating actions that marketers are taking today could best be classified as defense: they help claw back the budget that's been stolen or retroactively clean up the mess that fraudsters leave behind. But going on offense and proactively taking steps to limit fraudsters' effectiveness - that's a whole new ballgame.

Let's talk first about the defensive tactics that are working and how they can play into a winning strategy. The two causes for hope uncovered in the survey break down into tracking and awareness.

Tracking

→ One of the hardest things to answer in many marketing fraud circumstances is "where did the fraudulent traffic and contacts come from and how did they arrive on the site in the first place?"

67%

Most marketers are using UTM parameters to manage campaign attribution. This same mechanism can help fight marketing fraud.

As it turns out, many digital marketers already have a piece of the answer in place.

Nearly two in three marketers are using UTM parameters on a regular or near-constant basis to measure the performance of the campaigns they're running across the internet. And those parameters can be configured to arm marketers with an immense amount of information about the source of the traffic captured in those campaigns.

UTM parameters themselves aren't enough to spot and stop fraud - there's no mechanism contained within them or within the analytics tools that process them that can alert on marketing fraud or block or challenge suspected traffic. But they can be instrumental in figuring out which campaign drove the fraudulent traffic when paired with a tool that can spot it in the first place.

Earlier, we wrote about the need for marketers to commit to scrubbing their first-party databases of fake contacts on a regular basis to ensure that any fraudulent contacts that find their way in are wiped clean before compliance issues become a major concern. And while there weren't enough marketers taking that action already, it is a tactic that many marketers have in their toolkit.

Similarly, the broad variety of solutions implemented in the piecemeal approach that most respondents identified play a part in the fight against marketing fraud. Not enough marketers are using them today, and none of the solutions are a panacea. But a slingshot stopped a Goliath in his tracks, and every small effort made toward marketing fraud prevention and remediation is a positive one.

Awareness

→ A series of cartoon PSAs from the 1980s used to tell kids, "now you know, and knowing is half the battle". If that's true, digital marketing leaders might actually be in decent shape. On the whole, marketers showed a reasonable awareness of the challenge, especially to the forms of marketing fraud that may be most pernicious.

And many folks have experienced marketing fraud personally, but may not have been aware that there was a name for the phenomenon. The respondents who failed to secure concert tickets, who sought the latest video game console and found an empty shopping cart, who were frustrated at meager conversion rates... they've all seen marketing fraud firsthand.

Not to mention actively seeking out information like the results of this survey. If knowing is half the battle, now it's time to saddle up and get the other half done.

8

Stopping Marketing Fraud and Driving Better Business Results

8.

Stopping Marketing Fraud and Driving Better Business Results

Throughout this report, we've described many of the various tactics that marketers are using today to combat the specific models of marketing fraud they encounter. And we've described the challenge in that piecemeal approach, particularly of the solutions that are remediating—rather than preventative—in their response to marketing fraud.

What's needed for modern digital marketers to truly get a handle on fraud and protect their budgets, reputation, security practices, and compliance is a robust preventative solution that spans the breadth of all of the digital marketing tactics in the quiver.

HUMAN Marketing Integrity is that solution.

HUMAN Marketing Integrity helps marketers uncover and eliminate bots from within their digital marketing efforts and landing pages. Using our multilayered detection methodology, HUMAN detects today's sophisticated bots that mimic human behavior and skew metrics. HUMAN Marketing Integrity helps marketers maximize engagement with humans, maintain cleaner data platforms, enhance lead generation, and deliver more efficient results.

9

Conclusion

9.

Conclusion

It's important that the takeaway from this report not be one of blame or failure on the part of digital and performance marketers or their colleagues. This report outlines an opportunity for marketers to educate themselves on the problem and step up and become their go-to-market organization's heroes. Protecting marketing budgets from unnecessary wastage is a noble aim in and of itself, but when enhanced with compliance initiatives, data hygiene best practices, and user experience safeguards, it becomes an absolute must-have.

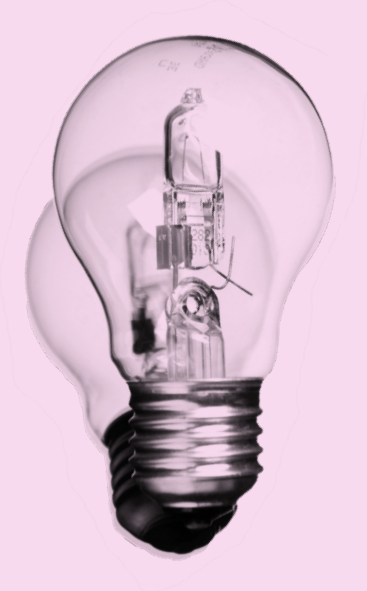

Ask yourself:

- **Do my numbers not make sense?**
- **Do I see odd or inexplicable traffic spikes on my website?**
- **Are my conversion rates lower than expected?**
- **Are my retargeting campaigns falling short?**
- **Is my data filled with leads that are fake?**

If the answer to any of these questions is "yes", you may have a sophisticated bot problem.

HUMAN Marketing Integrity is the solution to the challenge of marketing fraud. Marketing Integrity can identify even the most sophisticated bots pretending to be humans before they infiltrate your first-party database and wreak havoc throughout your marketing technology stack.

Contact us today for more information or to learn about a free marketing fraud risk assessment.

About Renegade LLC

Quite possibly the savviest B2B marketing agency in NYC, Renegade has been helping CMOs cut through since 1996. Wielding the wisdom gleaned from over 400 CMO interviews, Renegade is unique in its ability to simplify brand stories and translate these into goal-busting campaigns. To learn more, visit **renegade.com**.

About HUMAN

HUMAN ensures digital engagement with real humans. Our Human Verification Platform collectively protects enterprises from sophisticated bot attacks and fraud - faster and more accurately than any solution - because of our superior detection technology, global visibility, and threat intelligence. We verify the humanity of more than 10 trillion interactions per week for some of the largest companies and internet platforms. HUMAN recently secured a strategic growth investment from Goldman Sachs Merchant Banking Division, in partnership with ClearSky Security and NightDragon. Join our mission to protect digital engagement and your digital business better with HUMAN. To learn more, visit **www.humansecurity.com.**