REPORT

Beware the Artificial Impostor

A McAfee Cybersecurity Artificial Intelligence Report

May 2023

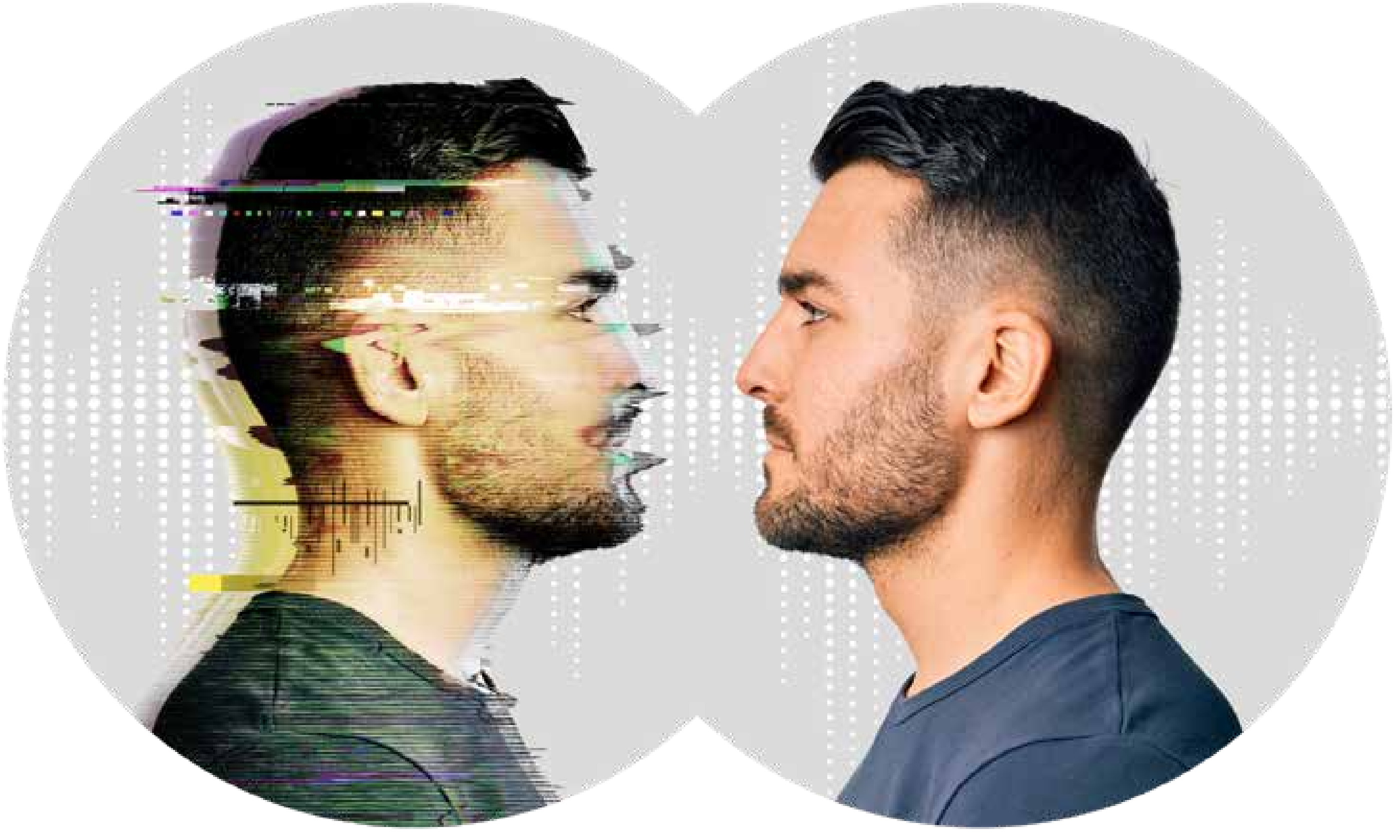

Table of Contents

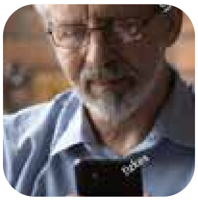

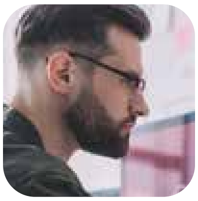

Summary

In the first months of 2023, McAfee observed a new threat emerging, with attacker groups using artificial intelligence (AI) voice-cloning technology to impersonate people and convince their family members or loved ones to send money.

This scam is an updated version of an old one. The 'Hi Mom' or 'Grandparent' scam, in which a cybercriminal impersonates a family member in distress and asks their target for financial support. This type of scam has gained pace in recent years and new technology is making it harder to detect. In 2022, imposter scams resulted in losses of $2.6 billion[1] in the U.S, with thousands of reports of people being swindled by those pretending to be friends or family. This trend has been seen around the world—in the U.K. a version of the scam carried out on WhatsApp alone resulted in £1.5 million[2] in reported losses over a four-month period last year, while in Australia the "Hi Mom" texting scam resulted in an average loss of $5,742[3] per victim in 2022.

The increased availability and sophistication of artificial intelligence tools is changing the game for cybercriminals. Today, these imposter scams, which were previously carried out by text, voice, or email, now feature a cloned AI voice of the target's loved one. With inexpensive and easy-to-use AI voice-cloning tools, fraudsters are delivering bespoke messages through calls or voicemails and fraudulently asking for help in the form of financial assistance.

In the past, those wishing to create these assets needed to have both the time and technical ability to replicate voices, but applications and tools that produce near instantaneous, highly believable results are now just a few clicks away.

In April 2023, McAfee commissioned research with more than 7,000 adults worldwide to better understand the level of awareness and first-hand experience of AI voice scams. McAfee conducted research in the U.S., U.K., France, Germany, India, Australia, and Japan. In addition, security researchers from the McAfee Labs team conducted an in-depth analysis and review of AI-cloning tools to evaluate the pervasiveness of the technology used in these scams and to discover how they could help consumers better protect themselves.

Report contributors: This reported was created with help in the form of AI voice cloning technology experimentation with analysis and insights from the following members of the McAfee Labs team: Vallabh Chole, Christy Crimmins, Oliver DeVane, and Abhishek Karnik.

Topic One

How AI voice cloning works

Topic One: How AI voice cloning works

With the rise of messaging platforms such as WhatsApp and video-based social media platforms including TikTok, it's now more common than ever for people to share recordings of their voices on the internet.

In fact, McAfee's survey found that 53% of all adults share their voice online at least once a week, with 49% doing so up to 10 times in the same period. The practice is most common in India, with 86% of people making their voices available online at least once a week, followed by the U.K. at 56%, and then the U.S. at 52%.

Do you share your voice online?

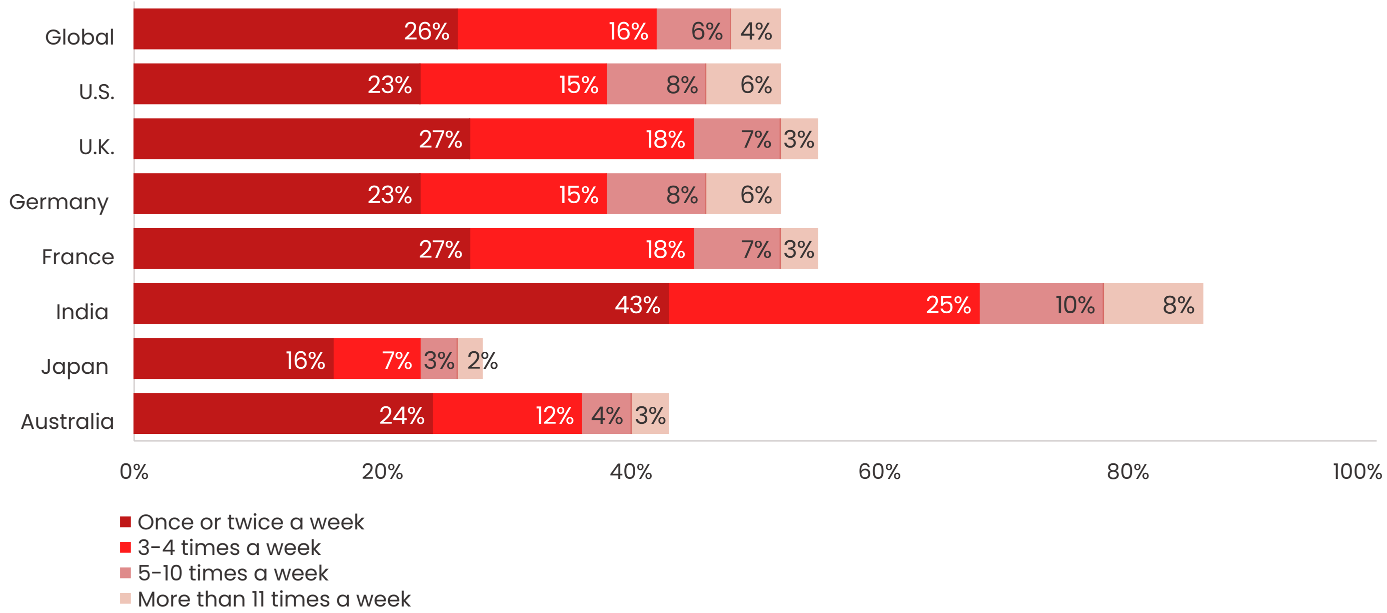

> "Targeted imposter scams are not new, but the availability and access to advanced artificial intelligence tools is, and that's changing the game for cybercriminals. Instead of just making phone calls or sending emails or text messages, with very little effort a cybercriminal can now impersonate someone using AI voice-cloning technology, which plays on your emotional connection and a sense of urgency to increase the likelihood of you falling for the scam."
>
> – Steve Grobman, McAfee CTO

While this may seem harmless, our digital footprint and what we share online can arm cybercriminals with the information they need to target your friends and family. With just a few seconds of audio taken from an Instagram Live video, a TikTok post, or even a voice note, fraudsters can create a believable clone that can be manipulated to suit their needs.

The rise of deepfakes and AI voice clones

While on the surface it might feel like a distant, far-off reality, the fast adoption and availability of artificial intelligence has rapidly changed the cybersecurity landscape. The ability to develop deepfakes, where AI is used to create images and videos of fake events, is not new. In fact, McAfee's CTO, Steve Grobman, demonstrated[4] the rise of this technology more than four years ago at the RSA 2019 conference.

However, those wishing to create these deepfake assets needed to have both the time and technical ability to do so, but now applications and tools that produce near instantaneous, believable results are just a few clicks away. And while we might see entertaining use cases featuring public figures—the Pope in a puffer jacket, for example, or former President Donald Trump being arrested—the same technology is also being used maliciously to replicate the voices of everyday citizens.

With the vast proliferation of AI tools, it's now easier than ever to manipulate content. Whether it's animating a picture, enriching the pixel count of images, or cloning a voice, the technical skill or investment needed in the past is no longer a barrier. With advanced AI tools becoming more pervasive, McAfee Labs expects scams, social engineering attacks, and false propaganda to be on the rise.

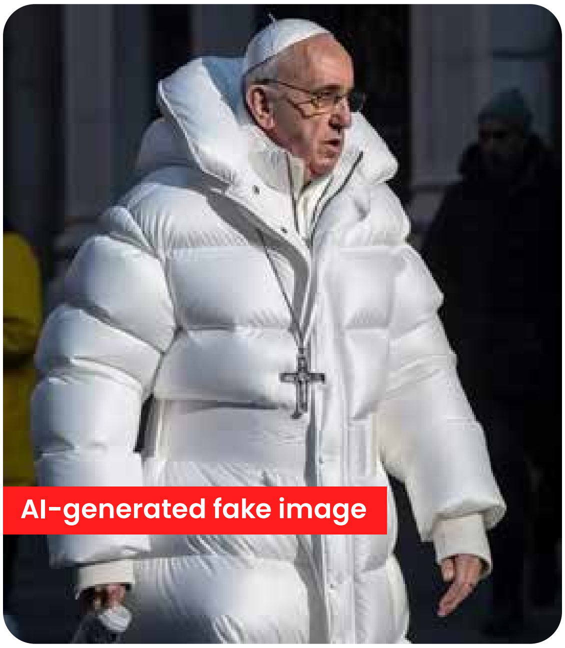

Figure 1. AI-generated deepfake image of the Pope from March 2023. Creator: Pablo Xavier

> "Artificial Intelligence brings incredible opportunities, but with any technology there is always the potential for it to be used maliciously in the wrong hands. This is what we're seeing today with the access and ease of use of AI tools helping cybercriminals to scale their efforts in increasingly convincing ways."
>
> – Steve Grobman, McAfee CTO

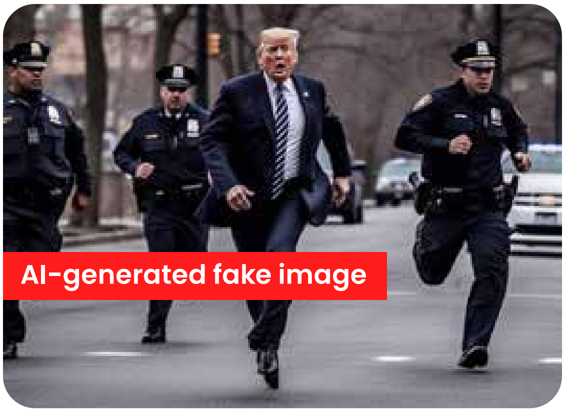

Figure 2. Deepfake image of former President Donald Trump from March 2023. Creator: Eliot Higgins

Topic Two

AI voice scams and their impact

Topic Two: AI voice scams and their impact

How common are AI voice scams?

A quarter of adults surveyed globally have experience of an AI voice scam, with one in 10 targeted personally, and 15% saying somebody they know has been targeted. When you break it down by country, it's most common in India, with 47% of respondents saying they had either been a victim themselves (20%) or knew somebody else who had (27%). The United States is second, with 14% saying it happened to them and 18% to a friend or relative. The U.K. comes third with 8% saying it happened to them directly and 16% to someone they know.

Overall, 36% of all adults questioned said they'd never heard of AI voice scams indicating a need for greater education and awareness with this new threat on the rise.

Have you or someone you know experienced an AI Voice Scam?

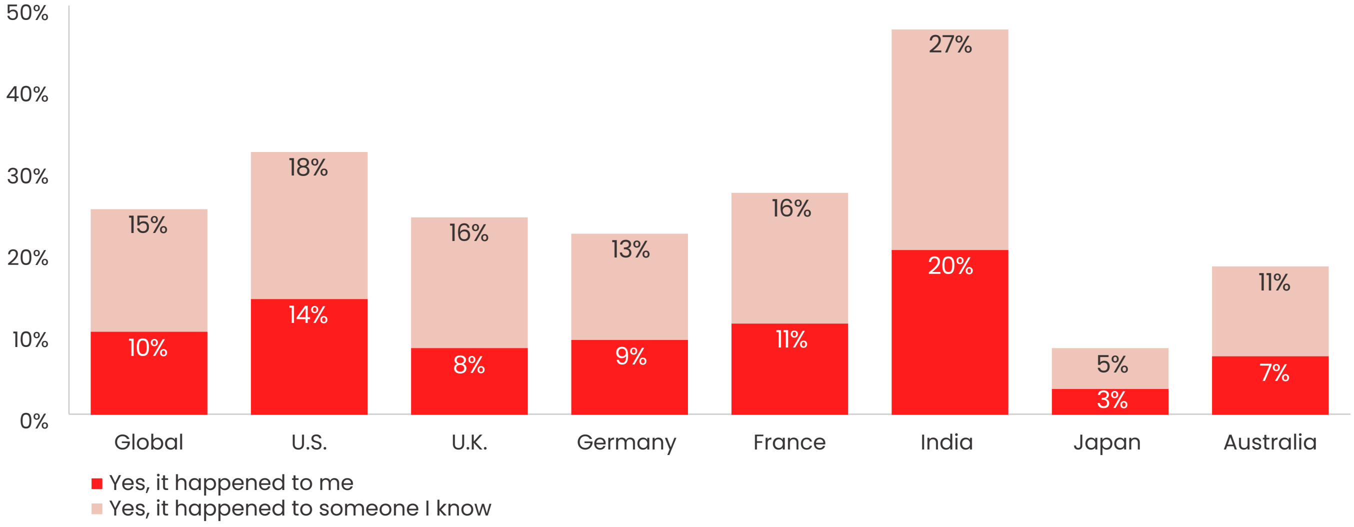

Which voice scams are people most likely to respond to?

Scams are almost always designed to convince the victim to part with their money, and while 45% of people said they'd respond to a request for cash from a friend or loved one, some of the scenarios invented by scammers are more likely to be successful than others.

Top of the list is a car crash or breakdown at 48%, with a robbery just behind at 47%. A lost phone or wallet would work with 43% of people, while 41% would reply to someone who said they were traveling abroad and needed help.

Who the message purportedly comes from also has a big impact on how likely the recipient would be to respond, with 40% saying they would be most likely to reply to a partner or spouse, followed by their mother at 24%. Parents aged 50 or over are most likely to respond to a child at 41%. It's worth noting, most known cases are of parents or grandparents reporting a cybercriminal cloning a child or grandchild voice and impersonating them.

Would you respond and share money if you received a voicemail or voicenote from a family member or friend, based on the following situations?

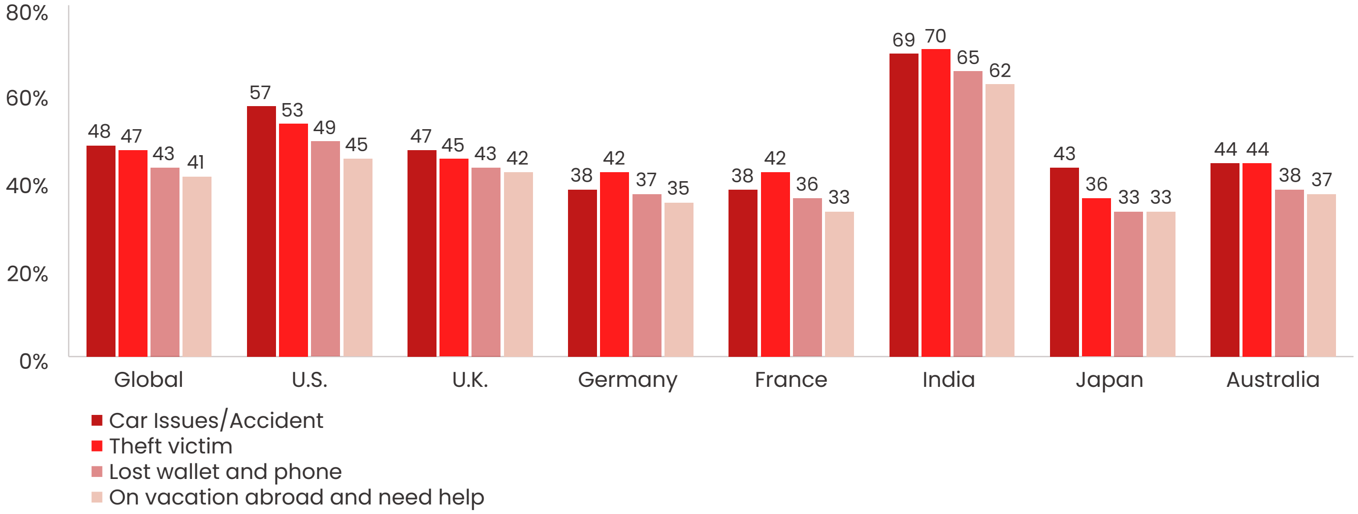

Can people tell the difference between real and fake?

Research by McAfee has found that voice-cloning tools are capable of replicating how a person speaks with up to 95%[5] accuracy, so telling the difference between real and fake certainly isn't easy. In fact, 70% of people said they were either unsure if they would be able to tell (35%) or believe they wouldn't be able to (35%).

Germany has the highest percentage of people who don't think they could tell the difference (48%), with 29% not sure, while 61% of Australians either don't know if they could tell (34%) or believe they can't (28%).

In the publicly reported cases of AI voice-cloning scams, victims have recounted how the voice sounded "just like" the person being cloned. In one particularly egregious case, where a cybercriminal demanded a ransom for a fake kidnapping, the mother said[6] it was "completely her voice" and that "it was her inflection." It's now harder than ever to tell real from fake, so people will need to assume they can't always believe what they see and hear.

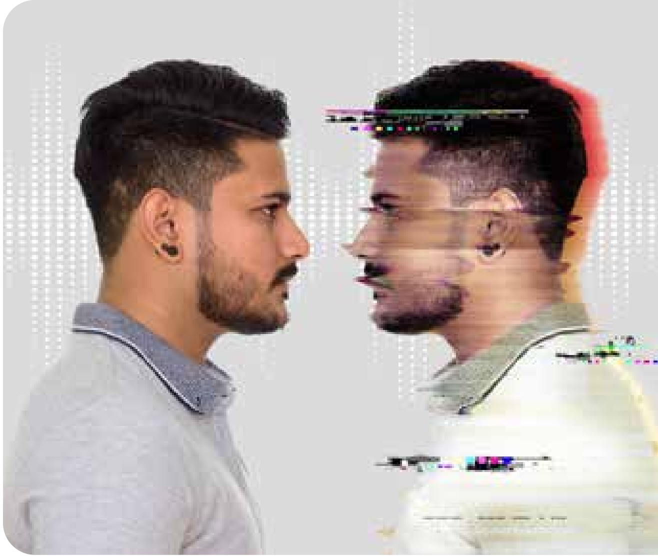

Could you tell the difference between a voicemail from a loved one and one that is AI generated?

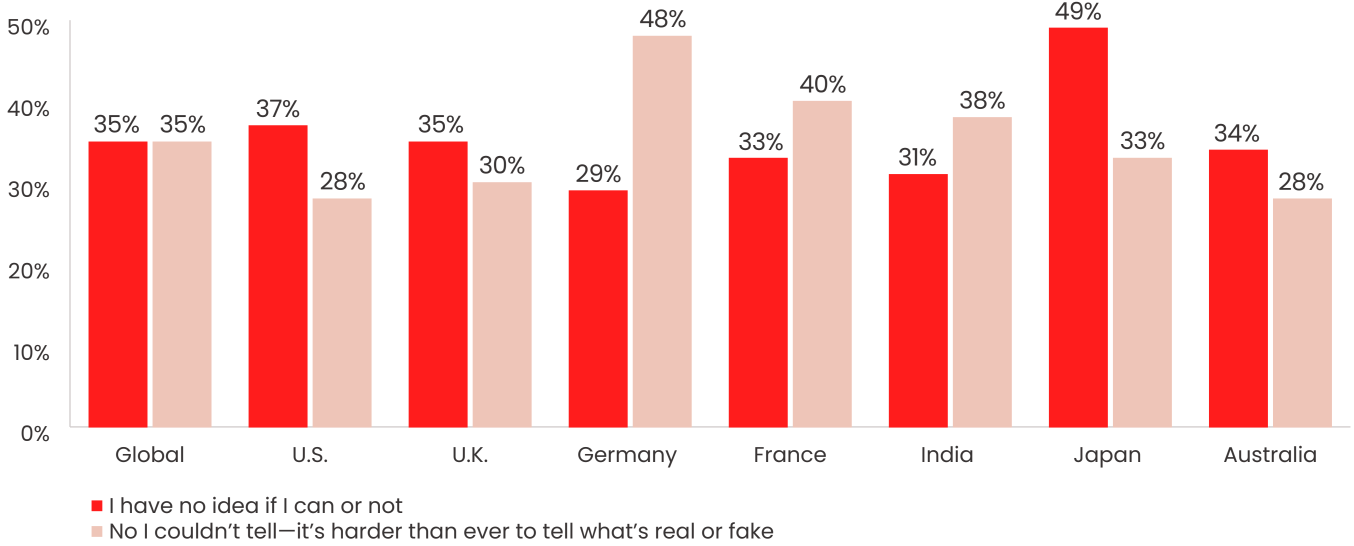

	I have no idea if I can or not	No I couldn't tell—it's harder than ever to tell what's real or fake
Global	35%	35%
U.S.	37%	28%
U.K.	35%	30%
Germany	29%	48%
France	33%	40%
India	31%	38%
Japan	49%	33%
Australia	34%	28%

Testimonials of those who have experienced an AI voice scam

With one in four adults globally having been impacted by an AI voice scam, a number of people have come forward to share their experience, what it sounded like, and how they were able to avoid becoming victim.

Phyllis

"After receiving several 'Grandma' calls, in conversations with my grandsons, I asked if they would call me. Each one said they would not want to upset me and if they were in trouble they would contact their parents. But, we also established a code sentence. Now when I ask the question to the person who is supposed to be my grandson, they hang up."

Jo

"It happened to me but I never fell for it. Sounded just like my granddaughter."

T.G.

"Yes this happened to my mom her grandson called said he had been in a wreck he was hurt and he needed money. Sounded just like him she said. She got scared and called him and he was okay."

Joaquin

"I've had this call twice this past year. It frightened the hell out of me. I asked a few questions. They didn't want me to tell grandma, I asked when they last spoke to grandma, they said 2 weeks ago. Grandma has been deceased a few years."

Joey

"Yes, this happened to a friend of mine who was traveling in Mexico. The caller knew she had been in Mexico, in reality, she had already returned from Cancan [sic] and was already back at work. The caller said she was locked up for a misunderstanding between herself and the Mexican state police. They needed money for a Mexican lawyer...Fake and fabricated with a voice that was left on voicemail so there was no way to call back and confirm. So I called Jessica and she was at work, just as I suspected."

Source: FTC[7]

What's the cost of falling for a voice scam?

AI voice scams are designed to be as convincing as possible, so it's no real surprise that 77% of all victims lost money as a result. More than a third lost over $1,000, while 7% were duped out of between $5,000 and $15,000. This is highest in the U.S., where more than one in 10 victims (11%) lost between $5,000–$15,000.

In the case of AI voice scams, fraudsters are counting on their target's desire to want to help a loved one. Humans make mental shortcuts daily for problem-solving and to reduce cognitive overload—known as heuristics in psychology. This comes into play in this scam, where our brain is more likely to cut corners and believe the voice we're hearing is, in fact, that of the loved one, as it's claiming to be. Because of this, a near-perfect match may not even be required, as our brain will automatically make the shortcut and often motivate the target to act and, in most cases, send money. This is why greater awareness and vigilance is needed to help protect against such scams.

McAfee Labs insights and assessment

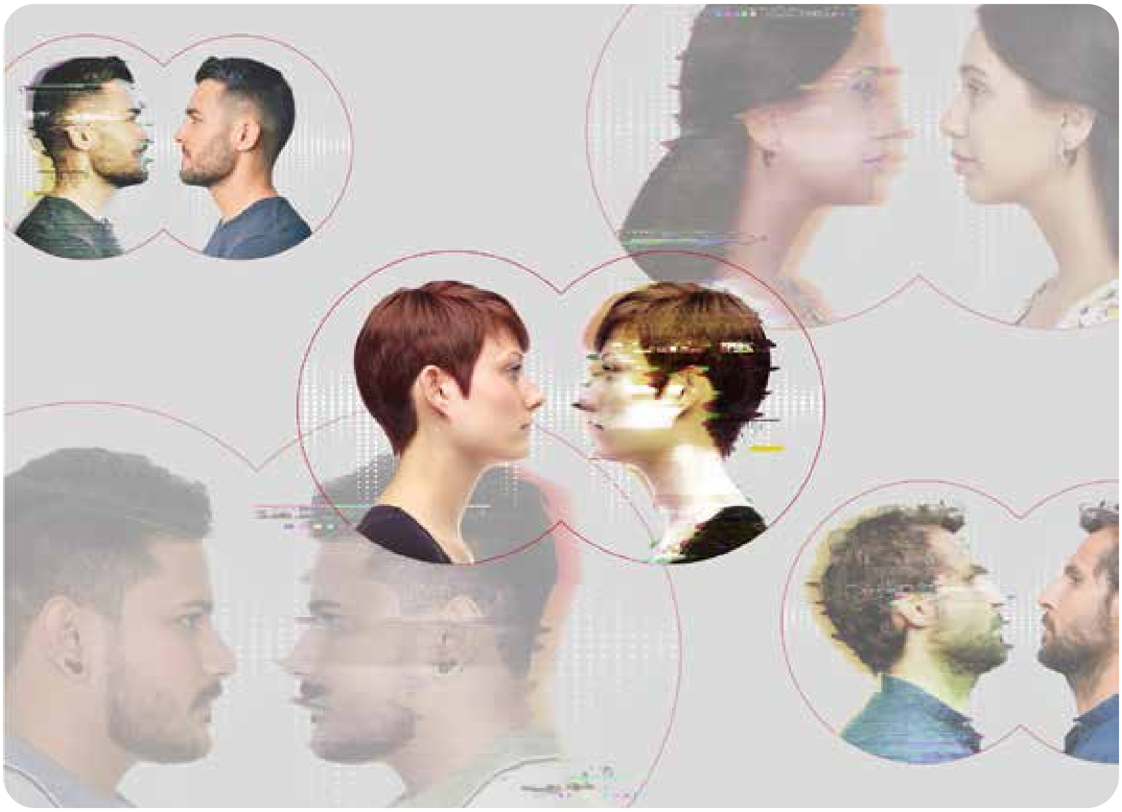

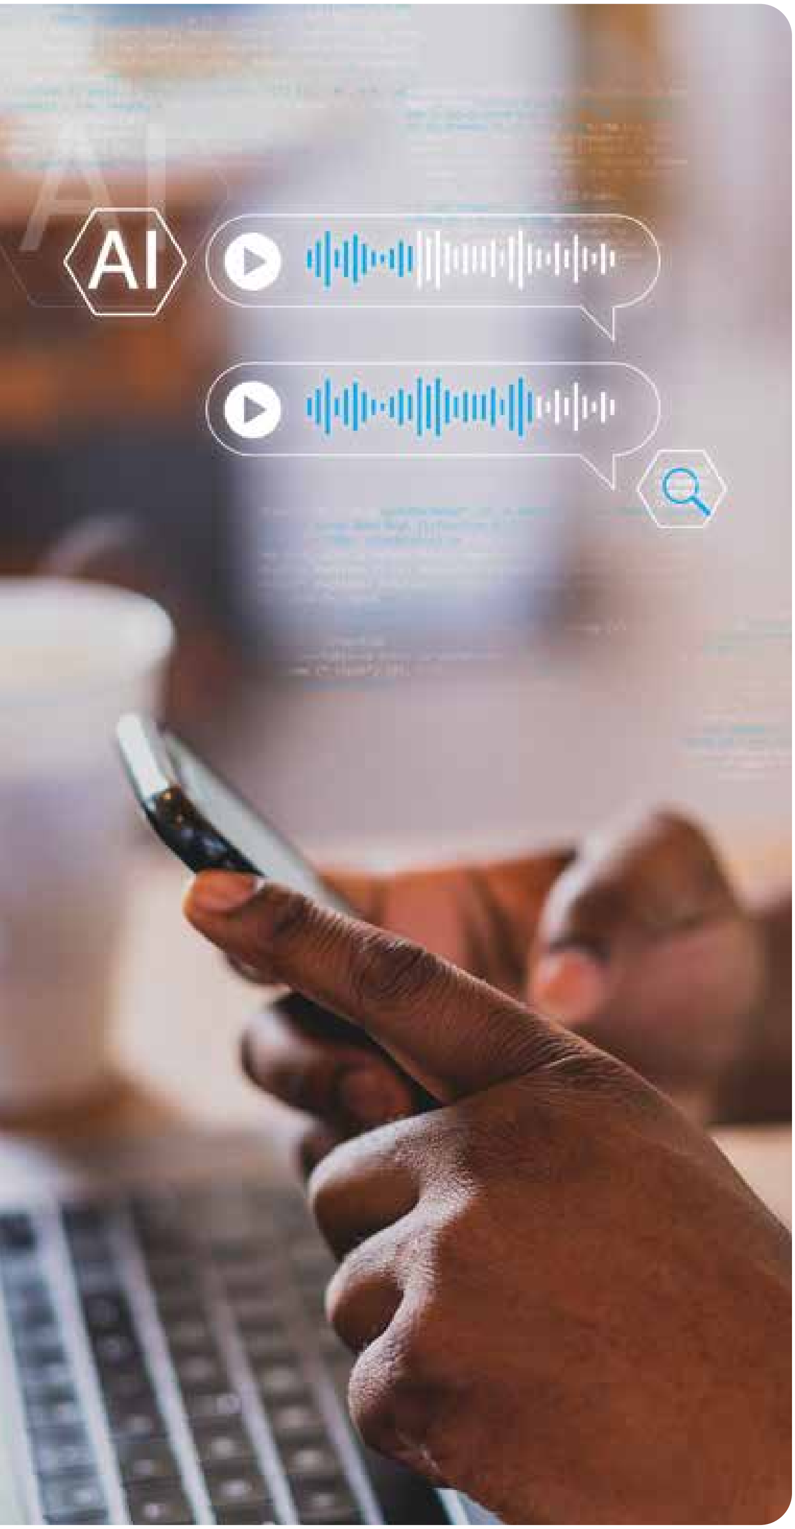

McAfee Labs insights and assessment

As part of the McAfee Labs team's evaluation of voice-cloning tools, security researchers discovered upwards of a dozen free and paid tools on the internet. They even found toolkits available for those wanting to create more sophisticated clones.

McAfee used one of these tools to replicate a researcher's voice. With just three to four seconds of voice recording, the free tool was able to create a convincing clone of her voice at an estimated 85%[8] match. For just $.0006 per second of audio produced, the team was able to record 100 prompts, which resulted in a higher-quality outcome. With the next paid tier, they were able to add things like emotion and inflection, making the voice clone almost indistinguishable from the real thing.

With more investment and effort, more accurate clones are possible, with the researchers able to train the data models to create an estimated 95% voice match.

McAfee researchers also tested the efficacy of different accents and were able to replicate many accents, including U.S., U.K., Indian, and Australian. In one instance, researchers were able to produce and then finetune an Australian voice to a near exact match using freely available tools.

One shortcoming of the technologies that the McAfee Labs researchers discovered was that the more distinctive the voice, the more challenging it was to clone. For example, copying the voice of a person who speaks with an unusual pace, rhythm or style is more challenging and means they're less likely to be targeted as a result. In text to speech conversions, some tools struggled with reading out acronyms or taking pauses in between sentences or depicting accurate emotion of a sentence. On the contrary, some tools were very natural sounding—they took short pauses in between long sentences to catch a breath or added stutter to make it sound more natural. It's clear that with directed training, these tools will only sound more and more natural.

Overall, the researchers determined AI has changed the game for cybercriminals and the barrier to entry has never been lower.

How to stay protected against AI voice cloning scams

How to stay protected against AI voice cloning scams

There are two areas to consider when it comes to staying safe from AI voice-cloning scams. The first is how to avoid being cloned in the first place. The more information we make available about ourselves through the internet or on social media, the more at risk we become to identity theft and other cybercriminal activity. Being proactive in both limiting and understanding what personal data is available is an important consideration.

Two ways to limit the likelihood of being cloned include:

1. **Think before you click and share**—who is in your social media network? Do you really know and trust your connections? Be thoughtful about what you are sharing on Facebook, YouTube, Instagram, and TikTok. Consider limiting your posts to just friends and family through the privacy settings. The wider your connections, the more risk you may be opening yourself up to when sharing content about yourself.
2. Identity monitoring services can help to alert you if your personally identifiable information is available on the Dark Web. Identity theft is often where AI voice and other targeted scams start. Take control of your personal data to avoid a cybercriminal being able to pose as you. Identity monitoring services provide a layer of protection that can safeguard your identity.

And four ways to avoid falling for the AI voice scam directly, include:

1. **Set a 'codeword'** with kids, family members, or trusted close friends that only they could know. Make a plan to always ask for it if they call, text, or email to ask for help, particularly if they're older or more vulnerable.
2. **Always question the source**—If it's a call, text, or email from an unknown sender, or even if it's from a number you recognize, stop, pause, and think. Asking directed questions can throw off a scammer. For instance, "Can you confirm my son's name?" or, "When is your father's birthday?" Not only can this take the scammer by surprise, but they may also need to regenerate a new response, which can add unnatural pauses into the conversation and create suspicion.
3. **Don't let your emotions take over.** Cybercriminals are counting on your emotional connection to the person they're impersonating to spur you into action. Take a step back before responding. Does that really sound like them? Is this something they'd ask of you? Hang up and call the person directly or try to verify the information before responding.
4. **Consider whether to answer unexpected calls from unknown phone numbers.** It is generally good advice not to answer calls from strangers. If they leave a voicemail, this gives you time to reflect and contact loved ones independently to confirm their safety.

> "It's important to remain vigilant and to take proactive steps to keep you and your loved ones safe. Should you receive a call from your spouse or a family member in distress and asking for money, verify the caller—use a previously agreed codeword, or ask a question only they would know. Identity and privacy protection services will also help limit the digital footprint of personal information that a criminal can use to develop a compelling narrative when creating a voice clone."
>
> – Steve Grobman, McAfee CTO

Survey methodology

The survey was conducted online between April 13 and April 19, 2023, by market research company MSI-ACI via email inviting people 18 years and older to complete an online questionnaire. In total, 7,054 completed the survey from seven countries. The sample size completed per country is as follows: 1,009 respondents in the U.S.; 1,009 respondents in the U.K.; 1,007 respondents in France; 1,007 respondents in Germany; 1,004 respondents in Japan; 1,008 respondents in Australia; 1,010 respondents in India.

Voice cloning assessment methodology

McAfee researchers leveraged free and paid tools during the month of April to assess the ease-of-use, access, and efficacy of cloned voices using artificial intelligence. To determine a confidence or match level, researchers ranked the clone against a scale of Low = less than 60% likeness; Medium = 60–85% likeness; High: 85% or more likeness. McAfee researchers reviewed for likeness of voice, as well as indicators of authenticity such as natural pace and rhythm, sounds of breathing, and even in some instances a natural stutter or pause in speech.

About McAfee

McAfee is a global leader in online protection. We're focused on protecting people, not devices. Our solutions adapt to our customers' needs and empower them to confidently experience life online through integrated, easy-to-use solutions.

For more information about
online protection, visit us at
mcafee.com/blogs

[1] https://www.ftc.gov/news-events/news/press-releases/2023/02/new-ftc-data-show-consumers-reported-losing-nearly-88-billion-scams-2022

[2] https://www.which.co.uk/news/article/notorious-hi-mum-and-dad-scam-spreads-from-whatsapp-to-text-message-an7N34c0gVbP

[3] https://www.accc.gov.au/system/files/Targeting%20scams%202022.pdf

[4] RSA 2019 Keynote: Lightning in a Bottle, or Burning Down the House?

[5] Voice match accuracy levels indicated are based on the benchmarking and assessment of McAfee security researchers. See assessment methodology.

[6] https://nypost.com/2023/04/12/ai-clones-teen-girls-voice-in-1m-kidnapping-scam/

[7] Testimonials published here are from the comments section of the FTC article "Scammers use AI to enhance their family emergency schemes" published on March 20, 2023

[8] Estimations based on research and analysis from McAfee researchers based on clones produced using widely accessible free and paid tools.

McAfee™

6220 America Center Drive
San Jose, CA 95002
888.847.8766
www.mcafee.com