stirista

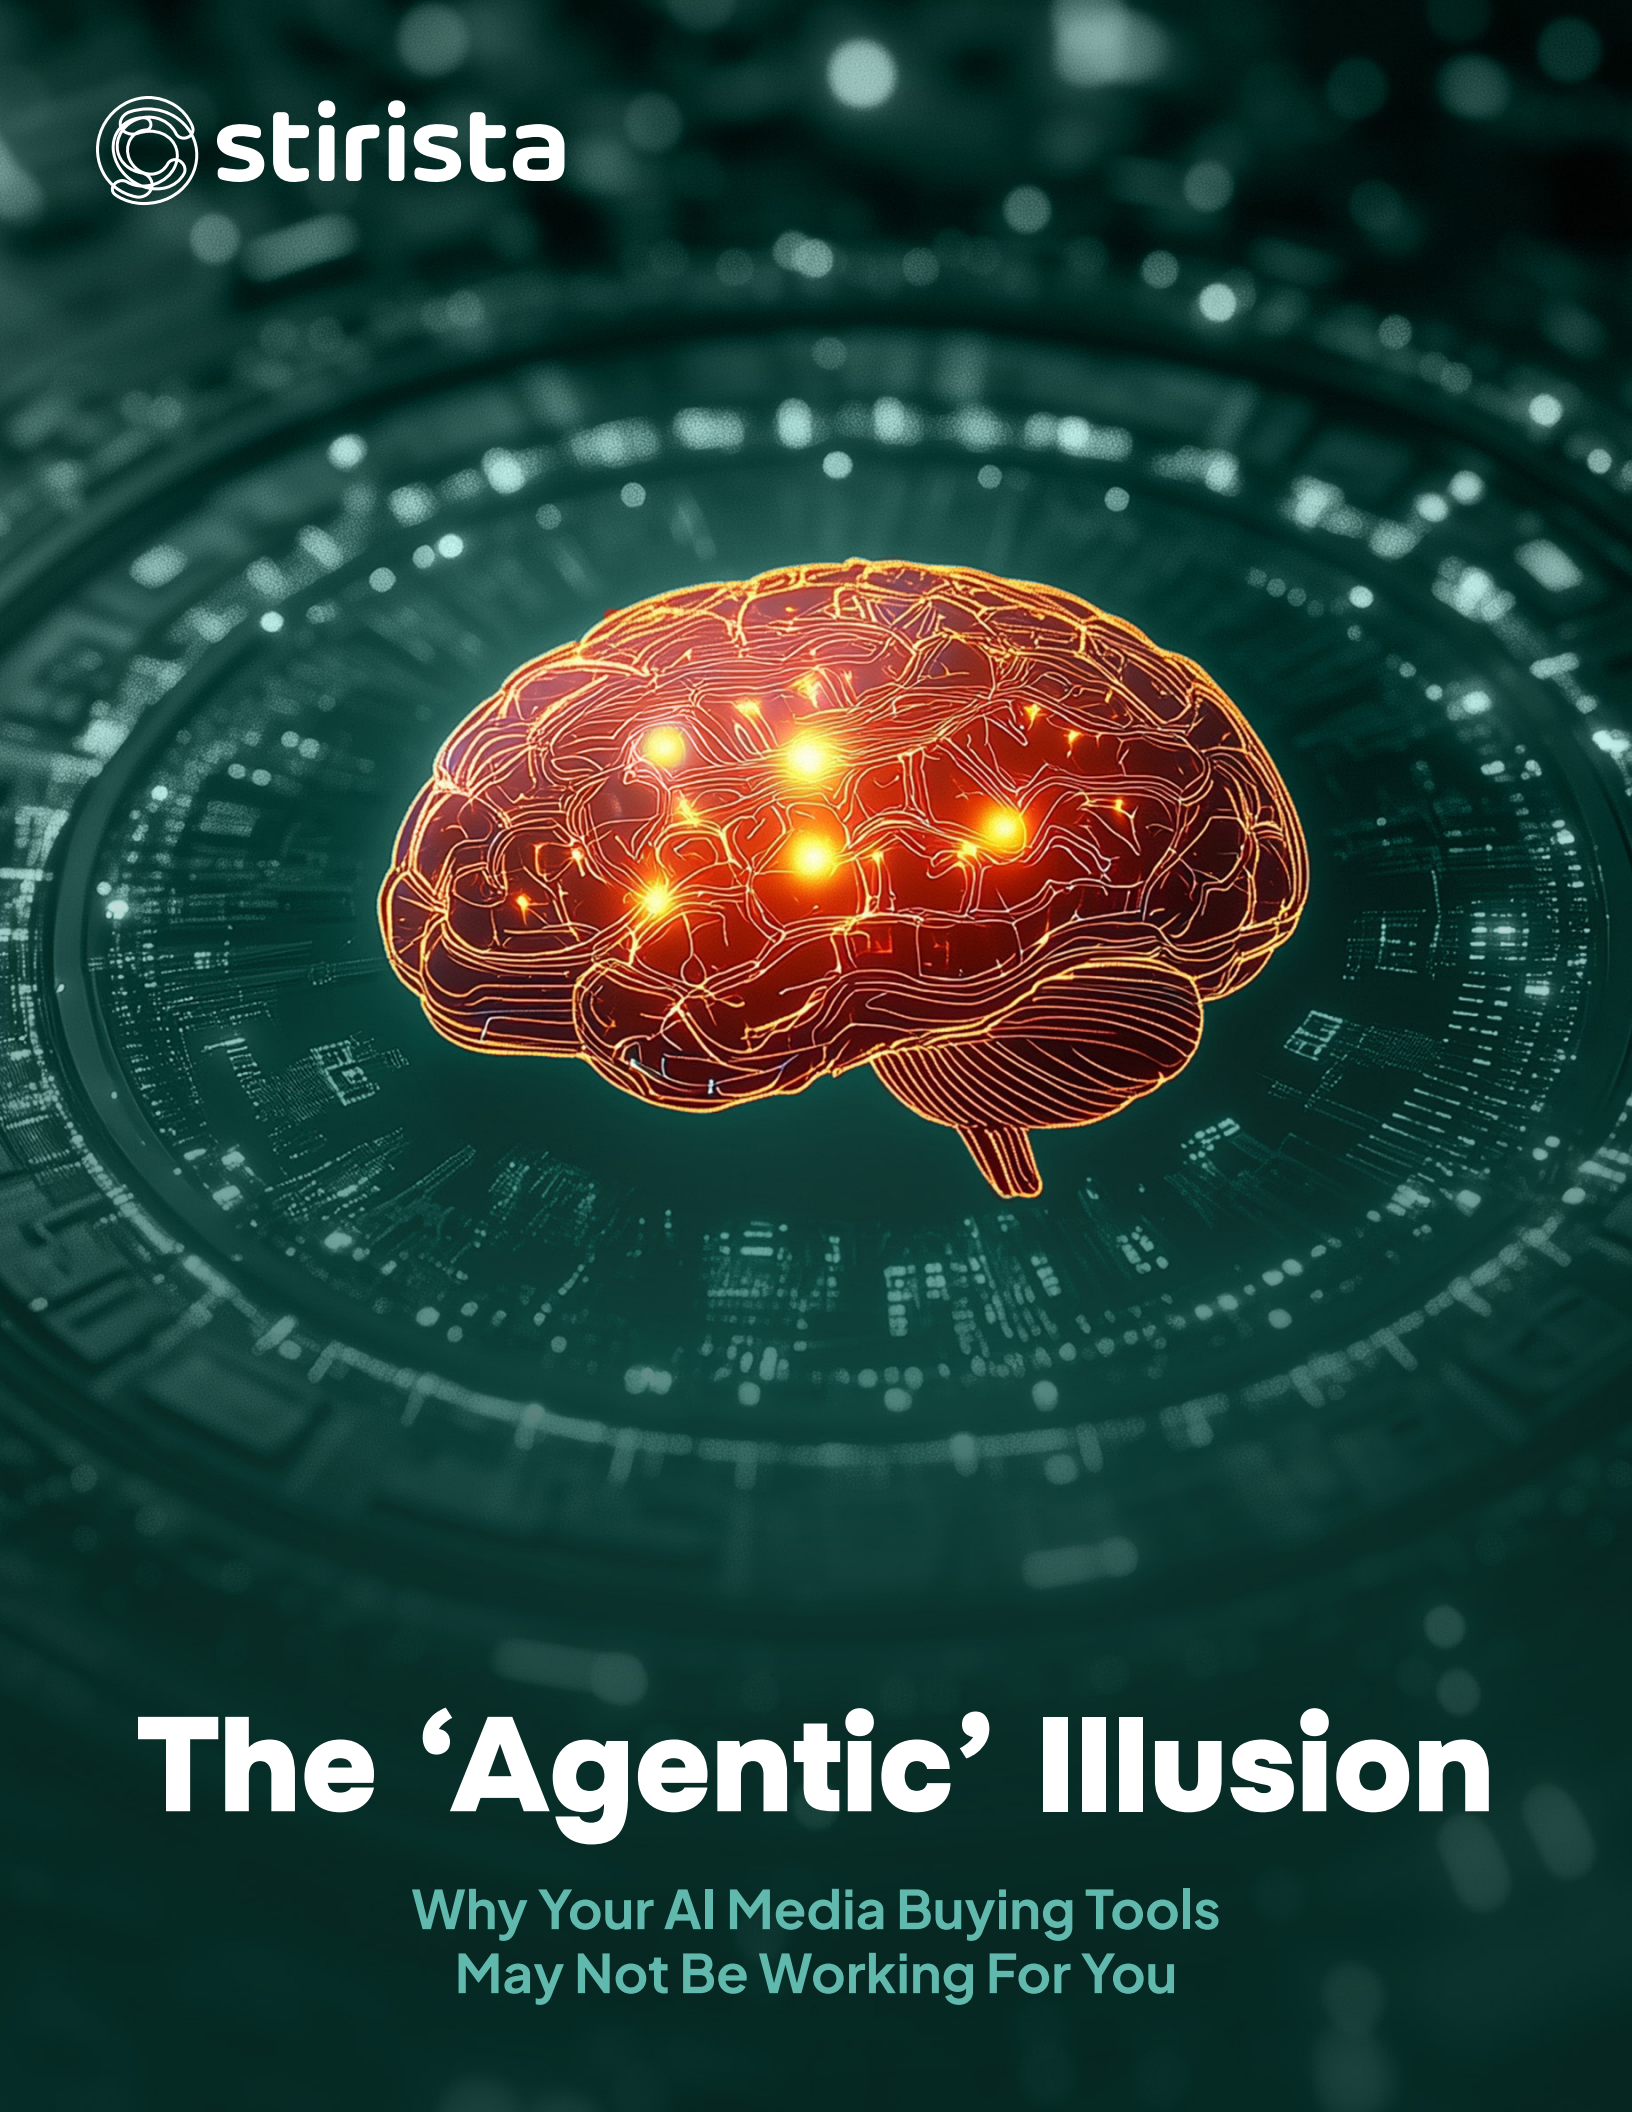

The ‘Agentic’ Illusion

Why Your AI Media Buying Tools May Not Be Working For You

AI tools promise easier media buying. But like principal media, they often hide where your ads go and why. Despite their efficiency, using "agentic" AI media buying platforms may not be worth the risk.

Tools like Google's Performance Max and Meta's Advantage+ (both released to the public in 2022) are not only commanding more attention, but more ad dollars: rapidly gaining share from marketers seeking quicker media solutions.

Yet not all that attention and budget-shifting is good.

While these AI tools ask buyers for parameters for campaigns like budget limits and sales goals, they then rely on algorithms to decide where ads will run, who they target, and how they appear. Essentially, they're a form of "agentic" AI: artificial intelligence tools that work autonomously in pursuit of a defined goal, with little user input.

Buyers often can't know exactly what decisions the AI agent implements, but they can essentially press a button for campaigns to run on their own. It's a trade-off that many buyers are willing to take.

Yet that trade-off includes significant downsides – like low transparency and low incrementality – that lead to even bigger downsides the more you rely on agentic AI media buying. It's a snowball effect.

While "agent" in traditional media buying signifies a relationship of direct advocacy and transparency (buying on behalf of the client), **AI tools, by their very nature, operate more like principal media – acting on their own internal logic and often with opaque methodologies.**

And this lack of true "agency" – like principal media before AI – creates significant risks for advertisers, leading to a loss of control, compromised performance, and a fundamental misunderstanding of how their media dollars are being spent.

Let's break it down further.

The Historical Context of "Agent" vs. "Principal" in Media Buying

AI buying agents may direct over 80% of digital media buys by 2030. At least that's what one prominent chief media officer thinks, according to a Wall Street Journal article.

The data may support that conclusion: AI-powered media buying tools are among the fastest-growing segments in adtech. According to an August 2024 study by eMarketer, about 60% of ad buyers had tried either Google's Performance Max or Meta's Advantage+. In the face of privacy-related signal loss, the AI-powered tools seemed like a good solution for advertisers strapped for effective solutions.

And Google and Meta weren't the only companies to release AI-powered media buying solutions. Amazon, Pinterest, LinkedIn, TikTok, and Microsoft all followed with their own agentic-AI media buying models.

But these AI "agents" aren't agents at all. At least, not in the traditional sense.

Rather, they operate like traditional principal media buying partners – not true agentic media buying, where an agent purchases media on your behalf. "Agentic" is a misleading way to frame these tools.

But before we get into the nitty-gritty, let's do a definition rundown.

Principal Media Buying

Principal media serves as a catch-all term for inventory media, proprietary media, and all other media types that stick to the following model.

In principal media buying, the agency buys the ad inventory (TV spots, digital display, programmatic, and the like) directly from the media owner (whether that's Google, Facebook, or some other platform) – often in bulk at discounted rates – and then resells it to their clients, typically at a marked-up price. The agency is the "principal" in this transaction – both the buyer and the seller of the media.

While principal media buyers take on the financial risk of buying the ad inventory (and the potential margin from selling it), there are multiple cons to this form of media buying, including risks of arbitrage and obfuscation – due to the inherent conflicts of interest within principal media buying.

Arbitrage refers to the practice of principal media partners who sell inventory simply to profit on it, without taking client objectives in mind; obfuscation refers to practices like not revealing the original purchase price or margins of their sales.

So while clients may get premium inventory at lower costs, it's usually bundled with low-value ad placements that provide little transparency to the client on where their ads are being placed, alongside what content.

80% | **AI buying agents may direct over 80% of digital media buys by 2030**

Agent-Based Media Buying

Agent-based media buying is a model in which an agency purchases advertising space or time directly from media vendors on behalf of the client, with the client maintaining ownership and financial responsibility for the media spend.

In this form of ad buying, the agency does not take a financial risk – they simply advise and execute a plan on behalf of the client, keeping the client's best interests in mind. Essentially, the agency acts as a fiduciary – and is typically paid a transparent fee or commission, for example, 15% of media spend.

In this model, the agency knows full details on media cost, placement, and any markups. Clients usually also retain audit rights and can verify media invoices and placements.

The pros to this form of advertising are both increased transparency and trustworthy advocacy on behalf of the client – knowing the agency has no ulterior motives in the transaction, clients can rest assured that their goals are kept top-of-mind.

Why Principal Media Buying Falls Short

Principal-based models come with an inherent conflict of interest – they need to offload loads of pre-purchased inventory onto unsuspecting clients.

A number of media agency groups acting as principals use their client deals to increase their own profits, rather than the other way around. Principal media can undermine media owners and vendors, especially those who are already vulnerable.

Altogether, principal media, along with other opaque practices in the ad industry, are reducing advertising effectiveness.

Unfortunately, marketers often don't realize they're working with a media agency operating on a principal model until it's too late.

Why Agent-Based Media Buying Has Long Won

Agent-based media has long been the gold standard for client-centric media relationships. Agent-buying puts the client first, promising transparency and good pricing in the process.

This is ultimately the result because there's an alignment of interest. **Agent-based media partners and their clients are working toward the same goal, with full transparency present during every step of the process.**

Deconstructing the "Agentic" AI Narrative

It's interesting that AI-media buying tools use the term "agentic AI" when publishers can use the practice to surreptitiously clear large amounts of their unsold inventory. While advertisers may appear to be getting improved performance in return, they lose out majorly on transparency and control.

To understand how AI-driven media buying tools conveniently co-opted the term "agentic" to insinuate an agent-based media buying model, we have to understand where the term "agentic AI" came from in the first place.

The Definition of "Agentic" AI

AI tool developers and researchers settled on the term "agentic" to describe autonomous LLM and machine-learning systems because they brought to light a key shift in how the tools are used: from passive tools to active "agents." Essentially, LLM tools were endowed with agency: the ability to make decisions and take action independently.

Intentionally or not, the term brings to mind a few key concepts in the marketing world: trust, control, and advocacy.

In the context of media planning, the idea of an AI "agent" doing all your media thinking, planning, and buying for you evokes the same idea of an actual agent working on your behalf.

Like agentic media buying, AI "agents" are sold, image-wise, as smart, trustworthy tools that operate on the buyer's behalf.

But the idea of an AI agent doing a task for you doesn't just fall short – it's completely wrong. It's a misnomer given the current realities of AI functionality.

AI agents – even the best of them – are still, fundamentally, "black boxes" of coding and intention. It's impossible to tease out all the decisions that get made within an AI agent, in the tangle of training data, developer input, and more.

This gets even more tricky with AI media buying tools, where inputs that aren't tightly defined (like audience target), decisions on the developer side may lead the tool to optimize for the easiest conversions, without the buyer knowing at all.

The way these agentic tools are built, it's hard to know why decisions are made the way they're made.

The Opacity of AI Media Buying Tools

Reports from Google will sometimes show higher returns for Performance Max campaigns than non-AI. But oftentimes, teams can't tell if their ads reached their core audience, and they can't target specific websites or apps that their target audience may visit.

A big complaint about the tools is that they're extremely opaque, with transparency and control rarely factoring into the equation.

It's easy to know where principal media buyers' interests lie – with themselves. But how can you know what an AI tool has been trained to do?

Prompt Injections & Training Data

Creator biases, objectives, and even unknown errors may be embedded into an AI tool's core programming. This training data can alter the output of these tools, unbeknownst to the buyer.

Additionally, prompt injection, a practice in which individuals manipulate inputs to LLMs to alter their behavior, is a real risk with any AI agent. While the practice isn't highly publicized among AI-media buying tools yet (after all, these tools only started appearing a few years ago), AI buying tools are still vulnerable by design.

> **"Prompt injection, a practice in which individuals manipulate inputs to LLMs to alter their behavior, is a real risk with any AI agent"**

Algorithmic Black Boxes

Users are blind to the internal decision-making processes of AI agents. Some refer to AI-media buying tools like Performance Max and Advantage+ as black boxes.

Not only do marketers get no insight into why these platforms make the decisions they do, they have little insight into what these decisions actually are themselves – what audiences are being targeted, where ads are being run, and how ads look. Marketers also have limited insight into how specific campaigns may have performed.

With key factors outside of teams' control, and results obfuscated, marketers can't tell which channels or audiences are performing well in their efforts. This leads to a collection of other issues: decreased incrementality, platform dependence, and ethical risks among them.

Lack of User Control

While users may interact with their AI media-buying platforms and set out campaign goals and desired outcomes, the ultimate "agency" resides with the AI's creators and their underlying programming. That includes whatever goals the company behind the platform has.

When users can't control who they target, where ads are running, and sometimes even how many (with little visibility in the process to boot), the AI agents behind the platforms are free to make whatever obscure decision they would like.

Again, AI-agent media buying mirrors principal buying, not true agentic media.

Real-World Implications and Risks for Advertisers

Here are a few real-world examples of marketing teams that have encountered problems with their AI media buying tools.

The Washington Post cites one example of the team behind baby products brand Lalo, who use Google's Performance Max. Their team explains they can't tell if their ads are reaching their intended audience, and they can't specifically target websites or apps their audience may frequent.

Or take this example on Adexchanger from James & James, a high-end furniture company that was spending more than half of its monthly budget on Meta's AI-powered Advantage+. In the end, that use wasn't leading to conversions, because AI-enabled products need careful instructions. And the AI-powered product wasn't delivering. The team ended up switching to Meta's manually controlled ad products, away from Advantage+. **What issues were companies like Lalo and James & James encountering?**

Loss of Strategic Control

Lalo and James & James are just two examples of how AI can veer off course from stated client objectives. **Marketers have, for years, wanted greater control and transparency in allocating budgets, not less.** AI, despite being the latest shiny new toy among many industries, could also be a way platforms are wresting control away from marketers.

Performance Obfuscation

With their campaign performance hidden, companies using AI media buying platforms experience difficulty in truly understanding what drives results (or lack thereof). This prevents incrementality and strategic planning based on prior results. **Rather, media planners are stuck doing what they've always done – with little opportunity to evolve or experiment.**

Ethical Concerns & Brand Safety

The potential for AI to make decisions that are misaligned with brand values or inadvertently expose brands to undesirable content/audiences. Take Grok, for example. In 2023 and 2024, Grok made headlines for producing unfiltered, "edgy" content – content that would alarmingly appear next to brand ads, associating the companies with inflammatory material.This is just one example of brand ads appearing next to unsavory content, unbeknownst to the marketing teams behind them.

Budget Misallocation

Hidden biases and opaque objectives can lead to inefficient spending on behalf of marketing teams. Spam leads are one common problem for marketing teams employing AI-agent media buying solutions, especially when input parameters aren't tightly defined.

Vendor Lock-in and Dependence

When you can't know why your ads are performing or to which audiences, you get locked in with the vendors you've always been using – leading to over-reliance on tools that lack transparency. Incrementality is often referred to as one of the things marketers lose with AI media planning and buying tools. Without insight into incrementality, marketers can't know the true impact of their marketing campaigns, beyond their baseline return.

Recommendations for Advertisers

So how do you get the most out of your media planning tools? How do you avoid being stuck in a loop with an AI media buying tool that doesn't allow you to expand on your marketing goals and pursue true understanding? Asking sharp questions can help you uncover whether a media buying tool – AI or otherwise – really has your best interest in mind and aligns with your goals.

We've seen demanding transparancy work with Google's Performance Max. When advertisers demanded transparency from the tool, Google responded by giving advertisers channel-level reporting in the platform.

Demand Transparency

- What level of visibility do I get into where my ads are served?
- Can I audit the list of publishers or inventory sources in real-time and retroactively?
- Do I get access to raw log-level data or only aggregated performance summaries?
- How are decisions made about where to place my ads – what signals or data are used?
- Can I see exactly which data points are used for targeting each ad placement?
- Are there any incentives (financial or otherwise) for your system to favor certain inventory sources?
- How does your AI make optimization decisions, and can we override them?
- What kind of user control do I have over budget allocation, frequency, or pacing?
- Can I cap spend per source or exclude specific inventory manually?
- How do you detect and mitigate unsafe or non-brand-aligned placements?

Understand the "Why"

Consider doing a deeper dive into the underlying mechanics of the AI tools you're using. Figure out how they function and what they're meant to accomplish. By unpacking the tool's architecture and assumptions, you'll be in a better position to identify potential blind spots or mismatches with your campaign goals. A deeper technical and strategic understanding also equips your team to ask smarter questions, make better-informed decisions, and hold vendors accountable.

Maintain Human Oversight

Human strategists and media buyers play a critical role in navigating the complexities of AI. Without human oversight, AI media tools can spin off the tracks. Left unchecked, they can chase performance metrics into undesirable territory, like targeting junk traffic or saturating narrow audience segments. That's where human media strategists come in. We recommend establishing regular checkpoints to audit performance and step in when brand integrity or ROI appears at risk.

Prioritize True Partnerships

Seek out vendors who operate on an agent-based model. Avoid principal-based models and those whose interests conflict with your goals. The right partner should feel like an extension of your team — one that shares your KPIs, adapts to your strategy, and gives you full control over how your budget is spent.

Consider Hybrid Approaches

The best outcomes occur when human experts manage ad campaigns and AI is used as support, not the other way around. When paid campaigns do work, it's never just because of AI. AI can be used supportively rather than as a fully autonomous "agent." For example, AI can help take control of the more monotonous parts of tasks within the marketing sphere. So automate those repetitive tasks that are eating up your time and focus on the thing that marketers do best – the strategy.

Why Stirista is Different
(or How to Embrace True Agency in a Data-Driven World)

We're bullish on new solutions, but we also weigh the true pros and cons of tools without hopping on the bandwagon too soon. When AI tools fall short, we will be the first to let you know whether you should believe the hype.

Stirista has long been committed to transparency, data ownership, and client advocacy. We value genuine control and advertising, in contrast with the "agentic" illusion of AI tools.

We empower advertisers with true agency – one that comes from a team with your goals in mind, operating on your behalf.

As a leader in transparent, data-driven, media buying Stirista believes in empowering advertisers with true control and understanding. We are committed to exposing misleading industry terminology and providing actionable insights that lead to better results.

stirista

stirista.com | info@stirista.com | 210-293-0029