

Dividing Leads Fairly Across 3–5 Travel Agents



Meet Maya. She runs a 5-person travel agency that books around \$2M a year in leisure and honeymoon trips. Business is good. But something is quietly breaking inside her team, and she can feel it before she can name it. This is her playbook for fixing it, told the way it actually unfolded.

The Problem That Finally Surfaced



On a Tuesday morning, Maya walked past two of her agents and noticed something strange. Both were typing a reply to the same Instagram DM. The prospect was shopping for a \$40K Maldives honeymoon. By lunch, she had received two different quotes and two different tones. By the next day, she had ghosted the agency entirely.

The Problem That Finally Surfaced



That afternoon, Maya spent two hours piecing the month's lead distribution together by hand. A spreadsheet her senior agent half-maintained. Screenshots from everyone's WhatsApp. Notes scribbled in the margins of a planner. When she finally had the numbers, they were ugly. One agent had claimed 60% of inbound leads. The other four were picking through whatever was left. Her team wasn't slacking. There was no system. "Whoever sees it first replies" had stopped being a rule and started being a fight, and nothing in her scattered tools could tell her who saw what, when.

Sitting Down With the Four Models



That night, after her agents went home, Maya sat down with the four lead-routing models she'd been meaning to research. She drew each one on the whiteboard and thought through how each would play out with her specific team.

Round-Robin




The simplest option. Leads rotate through agents in strict order. Lead 1 goes to Alex, lead 2 to Bea, lead 3 to Carlos, then back to Alex. It's provably fair and removes cherry-picking completely. Where it breaks: Maya's senior agent closes 40% of her deals, and her newest agent closes 10%. Strict rotation wastes leads on the junior. The fix is weighted rotation. Her senior gets 2 leads for every 1 the junior gets.

Territory




Assign leads by destination. Maya's senior agent genuinely knew Italy inside-out. Another had lived in Japan. Territory would route leads to real expertise. Where it breaks: half of her inbound said "maybe Europe, maybe Asia, we're not sure yet." Those leads still need a rule, and someone would end up with the pile of undecideds.

Destination-Expertise



A narrower specialization than territory. Instead of regions, agents own trip types. Honeymoons, corporate, family, luxury. Maya loved the close-rate math. A honeymoon specialist who books 3 a week will always outperform a generalist on honeymoons. Where it breaks: her family-travel agent would have starved through Q1's corporate-heavy season.

Score-Based



Biggest-scoring leads go to senior agents; smaller ones go to juniors. Makes structural sense at 6+ agents. At Maya's size, it almost always creates resentment. Juniors notice they only get leftovers. And a "low-score" lead that turns out to be a big booking makes the senior agent annoyed anyway.

Maya's Choice



After weighing all four, Maya picked weighted round-robin. Her senior agent would get 2 leads for every 1 that went to the newest. Her team was similar enough that territory and expertise were overkill for now. She could revisit in 12 months once she knew who was naturally good at what.

Maya's Choice



Tuesday afternoon, she opened her [lead management software](#) and configured the rule. Every new lead from Instagram, WhatsApp, web form, or Google Profile would auto-assign to an agent the moment it was created. No clicks required. The system made the decision, not the agent who happened to see the notification first.

Maya's Choice



Monday morning, she walked the team through it. Two questions came up. First, what about VIP repeat clients? Those stay with whoever handled them last time. Second, what if a prospect asks for a specific agent by name? The requested agent keeps the lead. She wrote both override rules into the team operating agreement so nobody would have to ask again next month.

Four Weeks Later



In the weekly pipeline review, Maya pulled up the distribution numbers. Variance was exactly where she'd expected. Her senior at roughly twice the junior's count. Close rates varied too, but in predictable ways. Her newest agent closed Instagram leads at 8% and web-form leads at 22%. That wasn't a fairness problem. It was a training opportunity.

Four Weeks Later



She also caught one agent gaming the system. Her newest hire had been claiming leads with a quick "Hi!" to mark them as contacted, then actually responding hours later. The software's SLA tracking flagged the pattern. One coaching conversation, and it was fixed. No policy changes needed. The structure did the enforcement.

Six Months Later



Six months after the change, bookings were up 40%. Zero inter-agent arguments about lead ownership had made it to her desk. Her senior agent had asked to mentor the junior, so Maya added a weekly deal-review pairing. And for the first time in two years, Maya had the bandwidth to focus on winning bigger corporate contracts instead of refereeing her own team.

Takeaways

Three things to take from Maya's story, in the order they matter:

1. **Replace "whoever grabs it" with a rule.** Any rule, consistently applied, beats no rule. But not all rules fit all teams. Pick the one that matches the shape your team is in today.
1. **Encode the rule in your CRM, not in memory.** The system should make the assignment at lead creation. If you rely on agent discipline, the discipline fails the first busy Monday.
1. **Review distribution weekly.** Ten minutes in your [weekly pipeline review](#) catches drift early. Look at both lead counts per agent and close rates per source-agent pair.

The goal isn't perfect fairness. It's the absence of arguments about fairness. When the system is the decider, your agents stop competing with each other. They start competing with the market, which is where your best sellers thrive.