

LISTINGLOGIC WHITE PAPER

What the Data Really Reveals

Consumer Behaviour in Property Marketing, Decoded from Real Campaign Data

Prepared by ListingLogic · August 2026

Based on real campaign performance data from Belle Property and Bayleys, distributed via ListingLogic's Meta and Google advertising platform

A ListingLogic White Paper on Consumer Behaviour in Property Marketing

Executive Summary

ListingLogic runs paid property marketing campaigns across Meta and Google for real estate agencies throughout Australia and New Zealand. Every campaign generates a granular record of who looked, who clicked, who watched, and who enquired — broken down by age, gender, and device. This paper analyses that record across a sample of individual property campaigns, a 175-campaign industry benchmark portfolio, and a 122-campaign client performance report, to ask a simple question: what does this data actually tell us about how New Zealanders and Australians shop for property?

The answer complicates the tidy story that is often told about "buyer personas." Property marketing data does not cleanly sort itself by demographic-to-property-type matching — young professional to apartment, family to suburb. Instead, six structural patterns dominate, set out in full below: mobile is no longer a channel preference but the default mode of property research, regardless of price point or dwelling type; the gender composition of an audience is driven far more by campaign objective than by the property itself; different age cohorts move through the buying funnel at fundamentally different speeds, meaning single-touch engagement data systematically undercounts the interest of older buyers; gender composition can be read as a real-time diagnostic of campaign health rather than just a demographic fact; regional and office-level variation reflects genuine local market dynamics rather than execution noise; and reach and lead generation are separate outcomes that respond to different creative levers and are frequently won by different listings within the same portfolio. Each of these findings is grounded in real campaign performance data, not illustrative hypotheticals, and each carries direct implications for how agencies should brief, target, and measure their campaigns — implications that are set out in a dedicated section following the data itself.

Introduction: What the Data Really Reveals

ListingLogic's dataset is fundamentally a record of consumer intent in action. Every click, impression, and engagement signal represents a moment when a potential buyer or renter made a choice — to look or not look, to engage or scroll past, to enquire or abandon. Aggregated across thousands of campaigns, properties, and demographics, this data becomes a microscope into how New Zealanders and Australians think about home, lifestyle, value, and investment. The insights available extend far beyond "which ads work" — they reveal deeper truths about who we are, what we want, and how our preferences are shifting.

To ground this analysis in fact rather than assumption, this paper draws on three real data sources. The first is a set of nineteen individual property campaign performance reports, sampled deliberately across price tiers (from AU\$610,000 apartments to AU\$4,000,000 estate homes), geographies (from Cairns to Adelaide to suburban Sydney), and creative formats (video-for-awareness, image-for-traffic, and carousel-for-lead-generation). The second is a 175-campaign industry benchmark report covering fifteen Belle Property offices and twelve marketing modules across May and June 2026. The third is a 122-campaign client performance report covering 103 unique listings marketed on behalf of Bayleys, spanning 6.27 million impressions and \$51,205 in media spend. Together these sources let us move from anecdote to pattern — and in several places, the patterns overturn the more intuitive story we might otherwise tell.

Fundamental Consumer Insights Embedded in the Data

1. The "Lifestyle Portfolio" Hypothesis: What Buyers Are Actually Seeking

The Insight. Engagement patterns reveal that buyers aren't just shopping for properties — they're shopping for lifestyle bundles. By analysing which demographic groups engage most with specific property types and features, ListingLogic can reverse-engineer the implicit lifestyle narratives buyers respond to.

The Data Story. When 25–34-year-old professionals on mobile devices show high engagement with central urban apartments featuring modern photos and minimal descriptions, while the same age group in suburban areas responds to family-oriented messaging about schools and commute times, the data is telling a story about identity and aspiration.

What the campaign data actually shows. This is an appealing hypothesis, and it is partly true — but the real data suggests the dominant driver of who shows up in a campaign's audience is not the property's location or type at all. It is the creative format and campaign objective

chosen for that property. Across the nineteen properties sampled — a mix of coastal houses, city-fringe apartments, and suburban family homes — video-for-awareness campaigns skewed male in every single instance, averaging 64% male engagement, regardless of whether the listing was a \$610,000 apartment in Cairns or a \$1.6 million family home in Carlingford. Image-based traffic and carousel-based lead-generation campaigns for the very same kinds of properties skewed the opposite direction, averaging 65% female engagement. The property did not change. The creative format did — and the audience flipped by roughly thirty points in the opposite direction.

This means the "lifestyle bundle" a buyer is responding to is shaped as much by the ad format ListingLogic chooses to run as by the property itself. A 35–44-year-old woman who barely engages with a video-for-awareness ad for a family home may respond strongly to an image carousel for the identical listing a week later. The lifestyle narrative is real, but it is being co-created between the property and the campaign mechanics — which means agencies have more control over who they reach than the "buyers self-select by lifestyle fit" framing implies. Rather than reverse-engineering a fixed lifestyle-to-demographic map, the more accurate model is a sequencing model: different creative formats recruit different segments of the same underlying buyer pool at different stages, and a full-funnel campaign captures far more of the true audience than any single format run in isolation.




Figure 1. Average audience gender split by campaign objective, across nineteen sampled property campaigns.

The practical implication for agents is significant: if a listing has run only a single awareness video and produced a male-skewed, younger audience, this is not evidence that the property "isn't resonating with families" — it is evidence that the format has only reached the top of the funnel. The lifestyle story only becomes visible once a property has been run through a fuller creative

sequence.

There is a second, subtler implication worth drawing out. Because the same underlying pool of local buyers is being recruited by different formats at different moments, the total addressable audience for a listing is larger than any single campaign report suggests, and it is composed of overlapping rather than distinct segments. A buyer who scrolled past an awareness video without engaging has not necessarily lost interest in the property — they may simply be waiting for the version of the ad that matches how they prefer to evaluate a decision at that moment, whether that is a fast visual scan (video), a considered comparison of photos (image/carousel), or a direct call to action (lead generation). Treating each format as though it recruits a wholly separate audience, rather than as different entry points into the same pool, understates how much of the market a well-sequenced campaign is actually reaching.

2. Mobile Is No Longer a Segment — It Is the Default

The Insight. Device preference used to be treated as a demographic signal: younger buyers on mobile, older buyers on desktop, with different creative and copy strategies recommended for each. The data no longer supports treating device as a meaningful segmentation variable at all.

The Data Story. Across all nineteen sampled properties — spanning a AU\$610,000 Cairns apartment to a AU\$4,000,000 Pymble estate, spanning houses, apartments, and penthouses, spanning video, image, and carousel formats — mobile device usage ranged from 92.67% to 98.43%, with an average of 95.3%. Desktop engagement never exceeded 2.5% for any property in the sample, and tablet engagement never exceeded 5.5%. There was no meaningful correlation between price tier and device split: the \$4,000,000 King Edward Street listing in Pymble showed 92.95% mobile engagement, statistically indistinguishable from the \$610,000 apartment in Cairns City at 95.54%.

Mobile dominates engagement across every price tier and property type sampled (range: 92.7%–98.4%)




Figure 2. Average device split across nineteen sampled property campaigns, spanning AU\$610,000–AU\$4,000,000 listings.

This finding should reshape how agencies think about creative production. Property marketing has historically treated desktop as the "serious buyer" surface — the assumption being that someone considering a multi-million-dollar purchase would eventually sit down at a laptop to review a listing in detail. The data does not support this. Even at the top of the price range sampled, the buyer's primary point of contact with a listing is a phone screen, almost certainly held in one hand while scrolling a feed. This has direct creative consequences: hero images need to read clearly at thumbnail size before a viewer taps through; video content needs to hook attention in the first two seconds of a vertical, sound-off feed; and long-form property descriptions that read well on a 27-inch monitor are frequently invisible to 95% of the audience who will never see them in that format. The mobile-first buyer is not a segment to design for — it is simply the buyer.

This has knock-on consequences for how ListingLogic and its agency clients should think about video specifically. Across the Belle Property benchmark portfolio, video formats generated 280,238 video views; across the Bayleys client portfolio, Meta Video alone generated 774,282 video views from \$44,430 of the total \$51,205 spend — by far the largest allocation of any channel. Given that this volume of viewing is happening almost entirely on mobile, video creative that is edited for a widescreen, desktop-style viewing experience is being watched, almost without exception, in a vertical, one-handed, frequently sound-off context. The gap between how video is typically

produced for property marketing — wide establishing shots, slow pans, ambient soundtrack — and how it is actually consumed represents one of the more straightforward, low-cost opportunities in this dataset: reformatting existing footage for vertical-first, caption-led, fast-cut delivery would not require new production budget, only a different edit of assets agencies are already generating.

3. The Consideration Curve — Why Different Age Cohorts Enter the Funnel at Different Speeds

The Insight. Age-based engagement data is usually read as a static preference — "this age group likes this kind of property." But when the same property is tracked across a full multi-touch campaign funnel, rather than a single ad, the age composition of the engaged audience changes dramatically as the funnel progresses. This suggests age does not predict whether someone will engage, but when in the funnel they will.

The Data Story. Across the nineteen properties sampled, the 35–44 age bracket was the single most responsive group in twelve of nineteen campaigns — nearly two-thirds of the sample — with 45–54 leading in five more. The 18–24 bracket was the least responsive age group in eighteen of the nineteen campaigns, essentially without exception.

The 35–44 bracket leads first-touch engagement in nearly two-thirds of sampled campaigns




Figure 3. Count of campaigns (of 19 sampled) in which each age bracket was recorded as the single most responsive.

That pattern held consistently for single-touch snapshots. But a case study from 36A Bristol Crescent, Davoren Park — a property run through a full multi-stage funnel spanning brand awareness (Google), engagement (video), traffic and conversions (image), and lead generation (image) — tells a more layered story. Looked at through a single image-for-traffic campaign in isolation, the engaged audience skewed toward 35–44 (30% of engagement) and 25–34 (24%), with the 65+ bracket registering essentially zero engagement. But aggregated across the full funnel

— every touchpoint, from the first awareness impression through to the final lead-generation ad — the 65+ bracket became the single largest segment of the entire engaged audience, at 34%, roughly matching the combined weight of every bracket under 45.




Figure 4. 36A Bristol Crescent, Davoren Park — age distribution of engaged audience, single-format campaign vs. full multi-touch funnel.

The gender composition of the same funnel tells a consistent story: a single traffic-focused campaign for this property showed 70% female engagement, but the full aggregated funnel — pulling in every stage — pushed to nearly 77% female. Older, more female-skewed segments of the audience were present throughout the campaign, but they simply took longer, and more exposures, before they clicked, watched, or enquired. A campaign measured only on its first ad, or only on its best-performing single format, will systematically undercount exactly the buyers who need the most nurturing before they act — which, on this evidence, tend to be older and more likely to be female. For agencies running single-format campaigns to save budget, this is a direct warning: the buyers most likely to eventually convert may be the ones a short campaign never gives enough time to show up.

It is worth being precise about what this case study can and cannot claim. A single property, run once, is a case study rather than a statistical proof, and the shift observed at 36A Bristol Crescent — younger and more male at first touch, older and more female by the time the funnel completes — should be read as a documented example of a pattern worth testing at scale, not as an established law of the New Zealand and Australian property market. What makes it a useful example rather than an isolated curiosity is that it is directionally consistent with two other findings in this dataset: the broader eighteen-of-nineteen dominance of "least responsive: 18–24" across single-touch campaigns, and the systematic gender-by-objective pattern documented in Section

1. Taken together, these three independent observations point in the same direction — that a meaningful share of genuine buyer interest, concentrated among older and more female-skewed segments of the audience, is simply invisible to campaigns that are not run for long enough, or through enough creative stages, to let it surface. The methodology section at the end of this paper sets out explicitly where a larger sample would be needed to move this from a documented pattern to a confirmed rule.

4. Gender Composition as a Diagnostic Signal, Not Just a Demographic Fact

The Insight. Once it is established that gender skew tracks campaign objective rather than property type, gender composition becomes useful in a different way: as a real-time diagnostic for whether a campaign is actually reaching its intended stage of the funnel, rather than simply a description of who happened to see an ad.

The Data Story. The 175-campaign Belle Property benchmark portfolio recorded a weighted click-through rate of 2.74% against an industry benchmark of 2.4% — a 14% uplift — but that headline number masks enormous variation by module. Facebook and Instagram Image for Lead Generation was the standout performer at 4.20% CTR and 530 leads from 21,003 clicks, while Website Traffic and Conversion campaigns achieved 6.72% CTR — nearly triple the benchmark — but Awareness-focused modules across the same portfolio underperformed at just 0.13% CTR.

Read alongside the gender data, this is not a coincidence. The highest-converting modules in the portfolio — lead generation and traffic and conversions — are precisely the formats that recruit predominantly female audiences in the property-level sample (averaging 64–65% female engagement). The lowest-converting module — awareness — is precisely the format that recruits predominantly male audiences (averaging 64% male engagement). This does not mean women are inherently more likely to convert into property leads and men are not; it means the creative objective that reaches men most efficiently (broad-reach video awareness) is structurally the objective least associated with direct conversion, while the creative objective that reaches women most efficiently (targeted image and carousel formats) is structurally the one most associated with conversion. A campaign manager checking mid-flight gender composition against expected objective can therefore use it as an early warning signal: if a lead-generation campaign is unexpectedly skewing 60%+ male, that is a signal the algorithm may be delivering the wrong creative to the wrong audience segment before a single lead is lost, not after.

This diagnostic use of gender data also offers a partial explanation for the wide performance gap between modules that on paper target similar buyers. The 6.72% CTR achieved by Website Traffic and Conversion campaigns in the Belle Property benchmark, against just 0.13% for Awareness campaigns, is a 51-fold difference in click efficiency within the same portfolio, run across the same offices, in the same two-month window. Some of that gap is explained by objective mechanics —

traffic campaigns are optimised by the platform for clicks, awareness campaigns are optimised for reach — but the consistency of the gender split across every property sampled suggests the platform's own optimisation is also selecting for a different, more click-prone audience segment when the objective changes, not simply changing how it bids for the same audience. In practical terms, this means the choice of campaign objective is doing double duty: it is simultaneously a bidding instruction to the ad platform and an audience-selection instruction, and the two effects compound each other. Agencies that want to reach a broad audience and later remarket to convert them cannot assume the remarketing audience will resemble the awareness audience — the data here suggests it may look substantially different, skewing considerably more female and slightly older.

5. Regional and Office-Level Variation Reveals Distinct Local Market Dynamics

The Insight. Aggregate national or portfolio-level averages hide meaningful regional variation in how audiences respond, and that variation appears to reflect genuinely different local market conditions rather than noise.

The Data Story. Within the same 175-campaign Belle Property benchmark, the Beecroft | Carlingford office led the entire portfolio in total clicks (10,542) with a 3.30% CTR, while the Cairns office generated fewer clicks (9,800) but nearly double the lead volume (140 leads) at a lower 1.98% CTR, and the Pymble office recorded the highest CTR of the three lead offices at 3.72% but converted only 16 leads from its clicks. Average campaign frequency across the portfolio was a modest 1.26x, suggesting audiences were not being over-exposed to the same creative, which makes the divergence between offices harder to explain away as fatigue and more likely a genuine reflection of different local buyer behaviour: Cairns audiences appear more enquiry-ready per click, Pymble audiences appear more click-happy but slower to convert to a lead, and Beecroft | Carlingford sits in between on both dimensions.

This matters operationally because it argues against a single, portfolio-wide target CTR or CPL benchmark. An office converting fewer clicks into leads is not automatically underperforming — it may be operating in a market where buyers browse more before committing, which the data suggests is the case in Pymble, a market that also skews toward higher-value, longer-consideration-cycle listings such as the AU\$4,000,000 King Edward Street estate sampled in this analysis.

The practical consequence for account management is that a single office-level scorecard, applied uniformly, will consistently misjudge at least one of these three markets. Judged purely on CTR, Cairns looks like the weakest of the three offices; judged purely on lead volume, it is the strongest. Judged purely on CTR, Pymble looks like the strongest office in the portfolio; judged on cost per lead, it is the least efficient of the three. The offices are not performing inconsistently — they are each optimising against a different underlying buyer behaviour that the market itself imposes, and

a benchmark that does not account for that will either set unrealistic lead targets for slower-moving, higher-value markets like Pymble, or undervalue an office like Cairns that is quietly the most efficient converter in the portfolio. A more useful benchmarking approach, supported directly by this data, would set separate CTR and CPL expectations by price tier and market type rather than by office alone, since price tier appears to be at least as strong a driver of the CTR-to-lead relationship as any difference in local execution quality.

6. From Reach to Revenue: What the Bayleys Case Shows About Lead Quality

The Insight. Scale and lead quality are not automatically in tension — but the data shows they require deliberately different campaign structures to achieve simultaneously.

The Data Story. The 122-campaign client performance report run for Bayleys delivered 6.27 million impressions, 3.29 million reach, 133,414 clicks and 774,282 video views from \$51,205 in total media spend, at an average click-through rate of 5.34% — more than double the industry benchmark identified in the Belle Property portfolio, and comparable to high-intent search advertising rather than typical social distribution. Meta Video was the standout channel by scale, accounting for \$44,430 of spend and the overwhelming majority of impressions and clicks. But when the objective shifted specifically to lead generation, results changed character rather than simply scaling down: across three dedicated lead-generation campaigns, spend of just \$450 produced 19 leads at an average cost per lead of \$23.68.

Individual listings varied widely by which outcome they delivered. 96 Tapu Rise Heights led the entire portfolio in raw engagement with 6,930 clicks from 153,770 impressions, followed by 99 Liardet Street (4,592 clicks) and 23 Pacific Drive (3,720 clicks) — but none of these top-engagement listings appear among the strongest lead-generation performers. That distinction belonged to 5 Ross Place, which generated 34 leads at a highly efficient \$16.91 cost per lead, followed by 104 Beverley Terrace (29 leads at \$11.90 CPL) and 37 Fishermans Bend (20 leads at \$28.75 CPL). The listings that captured the most attention and the listings that captured the most leads were, in nearly every case, different properties — reinforcing the finding elsewhere in this data that reach and conversion respond to different creative levers and cannot be optimised by the same campaign structure at the same time.

This split is consistent with, and helps explain, the module-level pattern already observed in the Belle Property benchmark: high-reach listings such as Tapu Rise Heights and Liardet Street were driven overwhelmingly by Meta Video spend, the same objective shown throughout this paper to recruit a broad, male-skewed, top-of-funnel audience efficiently but to convert that audience into leads only rarely. The strongest lead performers, by contrast, sit in the profile of the traffic-and-conversion and lead-generation objectives — lower reach, more targeted delivery, and a female-skewed, more consideration-stage audience. The fact that only \$450 of the \$51,205 total Bayleys spend was allocated specifically to lead generation, and that this small allocation still

produced 19 verified leads, is itself a signal that lead generation as an objective is being meaningfully under-resourced relative to its efficiency: at an average cost per lead of \$23.68, scaling the lead-generation budget by even a modest multiple, without touching the video allocation that is already doing its job of building reach, appears to be one of the more directly actionable findings in this entire dataset. It is also worth noting that, following the introduction of two-factor authentication on enquiry capture, 100% of lead data in this reporting period was verified as genuine — removing a source of uncertainty that has historically made lead-generation ROI harder to assess with confidence.

Strategic Implications

Read together, these six findings argue for a specific, evidence-based approach to property marketing rather than a persona-driven one:

Design creative mobile-first, without exception. With mobile engagement consistently above 92% regardless of price tier, there is no property value or type in this dataset that justifies a desktop-first creative approach. Hero images, video hooks, and page load performance should be evaluated primarily on a phone screen.

Sequence campaign objectives deliberately, not just run "the best-performing format." Because gender and age composition shift predictably with campaign objective, a single-format campaign — however well it performs on its own metrics — only ever shows an agency one slice of the true audience. Properties run through a full awareness-to-lead-generation sequence surface materially different, and often more valuable, segments of the buyer pool than any single stage in isolation, including older and more female-skewed buyers who take longer to convert.

Treat gender composition as a mid-flight quality signal. Because conversion-oriented formats and awareness-oriented formats recruit structurally different audiences, unexpected gender skew partway through a campaign is a useful, fast diagnostic that a lead-generation campaign may be misdelivering — actionable well before lead volume confirms the problem.

Benchmark regionally, not just portfolio-wide. Office and market-level variation in CTR-to-lead conversion appears to reflect genuine differences in local buyer behaviour rather than execution quality, and portfolio-wide targets should be set with that variation in mind rather than penalising markets with naturally longer consideration cycles.

Separate the reach objective from the lead objective operationally. The Bayleys data shows the properties that generate the most clicks and the properties that generate the most leads are frequently not the same listings. Campaigns optimised purely for scale should not be judged on lead volume, and campaigns optimised for lead generation should be resourced and measured independently of reach metrics.

Conclusion

The original hypothesis this paper set out to test — that ListingLogic's data reveals fixed "lifestyle bundles" buyers self-select into based on demographic identity — turns out to be too simple, but not because the underlying intuition is wrong. Buyers absolutely are shopping for more than square metres; the data confirms that identity, lifestyle stage, and aspiration shape engagement. What the data adds, and what a purely conceptual analysis could not have surfaced, is that campaign mechanics are doing at least as much work as buyer identity in determining who shows up in the numbers. Format determines gender composition more than property type does. Funnel stage determines age composition more than fixed preference does. Mobile has stopped being a proxy for anything at all, because it is now simply how the entire market behaves.

For ListingLogic and the agencies it serves, the practical takeaway is one of agency rather than fatalism: because campaign structure shapes who engages as much as the property itself does, agencies have more influence over who they reach — and how completely they understand their true buyer pool — than a purely demographic view of the market would suggest. The next stage of this analysis should extend the funnel-aggregation approach used in the Bristol Crescent case study across a much larger set of properties, to test whether the age and gender shift observed there — younger and male at first touch, older and more female by the time a lead is captured — holds as a general rule of the New Zealand and Australian property market, or whether it varies meaningfully by region, price tier, or property type once a larger sample is analysed.

There is also a broader point here about how ListingLogic's data should be used internally. Every finding in this paper came from reports that already exist, generated as a routine by-product of campaigns that were run for entirely operational reasons — to sell a specific property, for a specific agent, in a specific week. None of this data was collected with a research question in mind. That is precisely what makes it valuable: it is a record of real behaviour under real market conditions, at a scale that deliberate survey research would struggle to match, and it is being generated continuously, across every campaign ListingLogic runs, whether or not anyone stops to analyse it. The six patterns identified in this paper are very unlikely to be the only ones present in that record. A modest, recurring investment in aggregating and interrogating this data — in the way this paper has done manually, across a sample of nineteen properties and two portfolio reports — would very likely surface further patterns of comparable commercial value, at a fraction of the cost of commissioning new market research to answer the same questions from scratch. In that sense, this paper is best read not as a final answer about how New Zealanders and Australians shop for property, but as a demonstration of what ListingLogic's existing data is already capable of revealing, once someone asks it the right questions.

Methodology & Data Sources

This paper is based on ListingLogic's own campaign performance data, drawn from three sources:

1. Nineteen individual property campaign performance reports, sampled across price tiers, geographies, and creative formats, accessed via ListingLogic's live reporting platform at prople.report. Individual reports referenced by name in this paper include 6/6 Cedar Road, Palm Cove (prople.report/r/2YIY9QhgKO); 36A Bristol Crescent, Davoren Park, tracked as both a single-format campaign (prople.report/r/eqN8mCOYTr) and a full aggregated funnel (prople.report/r/TGVtVlonQE); and 15 King Edward Street, Pymble (prople.report/r/9_CsWlp-8t).
2. The Belle Property Industry Benchmark Performance Report, May–June 2026, covering 175 campaigns across 15 offices and 12 marketing modules, sourced from ListingLogic's internal reporting.
3. The Listing Logic Campaign Performance Report prepared for Bayleys, covering 122 paid campaigns and 103 unique listings between May and June 2026, sourced from ListingLogic's internal reporting.

All figures reported are as recorded in these primary reports; no figures in this paper have been simulated or estimated. Averages across the nineteen-property sample are simple means unless otherwise noted, and sample sizes are stated alongside each finding to allow readers to weigh statistical confidence appropriately — in particular, the image-campaign and carousel-campaign averages in Figure 1 are drawn from four and two campaigns respectively, and should be read as directional rather than definitive pending a larger sample.